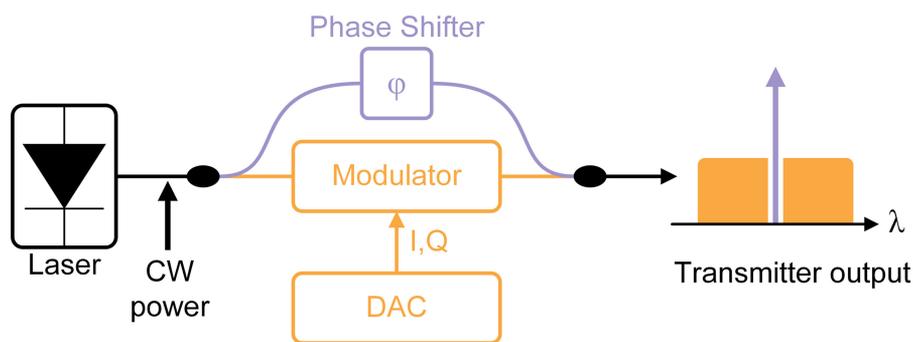


Improved Power Budget of 112 Gb/s/ λ Intra-Datacentre Links Using a Split-Carrier Transmitter Architecture

Volume 11, Number 4, August 2019

Thomas Gerard
Zhixin Liu
M. Sezer Erkılınç
Lidia Galdino
Polina Bayvel
Benn Thomsen
Domañıç Lavery



DOI: 10.1109/JPHOT.2019.2930325

Improved Power Budget of 112 Gb/s/ λ Intra-Datcentre Links Using a Split-Carrier Transmitter Architecture

Thomas Gerard ¹, Zhixin Liu ¹, M. Sezer Erkilinc ²,
Lidia Galdino ¹, Polina Bayvel ¹, Benn Thomsen ³,
and Domaniç Lavery ¹

¹Optical Networks Group, Department of Electronic & Electrical Engineering, University College London, London WC1E 7JE, U.K.

²Submarine and Core Systems Group, Fraunhofer Institute for Telecommunications, Photonic Networks and Systems Department, Berlin 10587, Germany

³Microsoft Research Cambridge, Cambridge CB1 2FB, U.K.

DOI:10.1109/JPHOT.2019.2930325

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see <https://creativecommons.org/licenses/by/4.0/>

Manuscript received May 22, 2019; revised July 8, 2019; accepted July 17, 2019. Date of publication July 23, 2019; date of current version August 7, 2019. This work was supported by the EPSRC TRANSNET programme under Grant EP/R035342/1. The work of T. Gerard was supported by Microsoft Research through its Ph.D. scholarship programme. The work of D. Lavery and L. Galdino was supported by RAEng Research Fellowships. Corresponding author: Thomas Gerard (e-mail: ucceetmh@ucl.ac.uk).

Abstract: A split-carrier transmitter is proposed to improve loss budget for short distance 100 G-per-wavelength intra-datcentre links. By applying a low-complexity hardware modification to the transmitter, the proposed architecture reduces transmitter loss without sacrificing modulation quality. These benefits are demonstrated by implementing intensity modulated, direct-detection signals using a single-ended photodiode receiver. The architecture is evaluated using two formats: Nyquist-shaped 56 GBd PAM-4 and Nyquist-shaped double sideband 28 GBd 16-QAM. The results are compared against a conventional transmitter modulated with Nyquist-shaped 56 GBd PAM-4. Through simulation, the split-carrier transmitter is shown to achieve hard decision forward error correction ready performance after 2 km of transmission using a laser output power of just 0 dBm; a 5.2 dB improvement over the conventional transmitter.

Index Terms: Optical transmitters, data centres, optical interconnects.

1. Introduction

The development of cloud data services is expected to see the internal bandwidth requirements of datacentres rise dramatically in the near future. To accommodate this, datacentre intraconnects will be upgraded from the current standard of 100G (4×28 GBd on-off keying channels) to 400 Gb/s. A favoured solution for reaching this higher rate is to use 56 GBd 4-level pulse amplitude modulation (PAM-4), achieving 112 Gb/s/ λ , multiplexed across four wavelengths or fibres [1]. This provides 400 Gb/s with 12% overhead. Although this approach is considered to be both low-cost and low complexity, it introduces numerous technical challenges that potentially increase system cost. The high symbol rate comes at the expense of requiring wide bandwidth electronics, external modulators, on-chip retiming and is susceptible to dispersion [2]–[4]. Furthermore, PAM-4 is a low-sensitivity format relative to its spectral efficiency [5], leaving little margin for system variation and future improvements in data rate. These factors lead to a low power budget transmission format

which requires high continuous wave (CW) laser power, defined as

$$P_{CW,req} = P_{mod,loss} + P_{fibre,loss} + P_{Rx,req}, \quad (1)$$

where $P_{mod,loss}$ is the modulator loss in dB, $P_{fibre,loss}$ is the fibre transmission loss in dB, and $P_{Rx,req}$ is the power required at the receiver, in dBm, for a specified bit error rate (BER). By expressing $P_{CW,req}$ in this way, Eq. (1) is observed to be analogous to system loss budget, which is simply the difference between the CW power deployed and $P_{CW,req}$.

A low $P_{CW,req}$ is necessary for 400G links in order to keep power consumption as low as possible; this has seen the specifications elsewhere in the system rise accordingly. For example, 56 GBd PAM-4 has been demonstrated to achieve a bit error rate (BER) of 3.8×10^{-3} over 2 km transmission without using digital signal processing (DSP) at the receiver [6]; however, this demonstration required a $P_{Rx,req}$ of +8 dBm, and a $P_{CW,req} > 21$ dBm. Given this, 400G links are expected to include an analogue-to-digital converter (ADC) and modest DSP to reduce $P_{Rx,req}$ to more practical levels [7]. 3-tap and 23-tap equalisers at the receiver have been shown to reduce $P_{Rx,req}$ to -5.9 dBm and -7.2 dBm, respectively [8]. Furthermore, transmitter-side DSP has been shown to alleviate modulation bandwidth requirements through a combination of signal pre-emphasis [9], [10] and spectral shaping [11], [12]. These methods have been used to demonstrated 56 GBd PAM-4 with a 3 dB modulator bandwidth as low as 16 GHz, though using a 161-tap least mean squares (LMS) equaliser as well as a 3-tap maximum likelihood sequence estimation receiver [13].

Other approaches have shown it is possible to directly modulate the CW laser at 56 GBd and compensate for chirp-related impairments using nonlinear adaptive filtering at the receiver [14], [15]; in this way $P_{mod,loss}$ is removed and system hardware costs are reduced. Currently these nonlinear filters are considered excessively complex for practical implementation, though more recent results have shown promising reductions [16]–[18].

Alternatively, both the symbol rate and DSP requirements can be reduced by using a high spectral efficiency modulation format, such as quadrature amplitude modulation (QAM). When used in combination with Nyquist-shaped subcarrier modulation (N-SCM), the phase-encoded data can be recovered using a single photodiode [19]–[21]. This technique allows for future improvements in both symbol rate and spectral efficiency, though demonstrations typically use IQ modulators, increasing system complexity and, crucially, $P_{mod,loss}$. Furthermore, the carrier-to-signal power ratio (CSPR) of these systems must be carefully balanced to optimise the trade off between modulation quality and signal-signal beating interference (SSBI), which typically needs to be further corrected for at the receiver [21].

In this paper, the recently proposed split-carrier transmitter (SCT) is investigated in the context of datacentre intraconnects. The SCT has been used previously as a method of CSPR control while maximising modulation quality in amplified links [12], [22]; however, more recent work has shown that the SCT can also be used to bypass modulator loss, significantly improving loss budget and outperforming mature systems [23]. This paper expands this principle by emulating the SCT as an integrated device. Simulation and experimental results are presented for both Nyquist-shaped 56 GBd PAM-4 and N-SCM 28 GBd 16-ary QAM (16-QAM), quantifying the SCT performance at 112 Gb/s in terms of received optical power as well as CW power. A characterisation of a quadrature-biased, Nyquist-shaped 56 GBd PAM-4 transmission is also provided, such that a side-by-side comparison can be made with a conventional transmitter using a similar setup and DSP to the SCT. The results show that the SCT can reduce $P_{mod,loss}$ without significantly impacting $P_{Rx,req}$, thereby outperforming the conventional transmitter in terms of required laser power.

2. Principle of Operation

2.1 Split-Carrier Transmitter

The proposed design of the SCT is shown in Fig. 1. The CW laser is split into an upper carrier branch and a lower signal branch. The signal branch is passed through a modulator operating in the linear regime, creating a high signal-to-noise ratio (SNR) carrier-suppressed signal. After

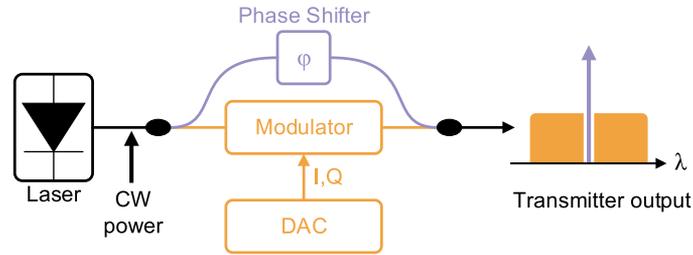


Fig. 1. Principle of the split-carrier transmitter. The laser is split into two branches: an upper carrier branch (purple) and a lower modulated signal branch (orange). When recombined, the branches form a self-homodyne signal with reduced loss. CW: continuous wave. DAC: digital-to-analogue converter.

modulation, the carrier and signal branches are coherently recombined to form a self-homodyne transmitter. This design is well suited for photonic integration, which ensures polarisation alignment between the carrier and signal branches. A phase controller, such as a heating element, may also be required to ensure the carrier and signal tones are in phase when recombined.

The power preserved in the carrier tone is transferred to the signal when received through direct-detection (DD). The amount of power saved is determined by the coupling ratio used when the CW laser is split and recombined, which also sets the CSPR. A coupling ratio that favours the carrier branch will decrease the transmitter power lost to the modulator, improving the loss budget. The power of the carrier tone, P_{carrier} , and the signal tone, P_{signal} , can be expressed as

$$\begin{aligned} P_{\text{carrier}} &= P_{\text{CW}} - 20 \log_{10}(C_{cc}), \\ P_{\text{signal}} &= P_{\text{CW}} - 20 \log_{10}(1 - C_{cc}) - P_{\text{mod,loss}}, \end{aligned} \quad (2)$$

where P_{CW} is the CW power in dBm, and C_{cc} is the carrier coupling coefficient. This determines the fraction of power routed through the carrier branch, and takes a value between 0 and 1 (with 0.5 representing a 3-dB coupler). Eq (2) therefore varies with $20 \log_{10}(C_{cc})$ as the coupler loss is applied twice, both before and after the modulator. Taking the difference of the terms in Eq. (2), the CSPR ratio is defined in dB as

$$\begin{aligned} \text{CSPR} &= P_{\text{carrier}} - P_{\text{signal}}, \\ &= 20 \log_{10}(C_{cc}) - 20 \log_{10}(1 - C_{cc}) + P_{\text{mod,loss}}. \end{aligned} \quad (3)$$

Also using Eq. (2), the total loss across the SCT can be defined as

$$\begin{aligned} P_{\text{SCT,loss}} &= P_{\text{CW}} - 10 \log_{10}(10^{P_{\text{carrier}}/10} + 10^{P_{\text{signal}}/10}) \\ &= -10 \log_{10}(10^{-2 \log_{10}(C_{cc})} + 10^{-2 \log_{10}(1 - C_{cc}) - P_{\text{mod,loss}}/10}). \end{aligned} \quad (4)$$

For typical $P_{\text{mod,loss}}$ values, the first term of Eq. (4) dominates, such that the loss across the SCT is dictated by C_{cc} . To visualise this information, Eq. (3) and Eq. (4) are plotted in Fig. 2 as C_{cc} is varied from 0.5 to 1 for a SCT with a typical modulator loss of 10 dB.

Fig. 2 shows how increasing the carrier coupling coefficient reduces the overall transmitter loss and increases the CSPR. A high CSPR reduces the effect of SSBI; however, if the signal branch power is too low, the signal is corrupted by transmitter noise. The carrier coupling coefficient must, therefore, be balanced to jointly optimise the power at the receiver and the CSPR. This relationship is investigated in Section 4.2.

2.2 Bias Procedure

A further conclusion that can be drawn from Fig. 2 is that the SCT brings the most benefit to a modulation scheme that requires a high CSPR, as it is in this regime that the transmitter loss is lowest. High CSPR requirements are frequently found in DD Nyquist-shaped modulation formats,

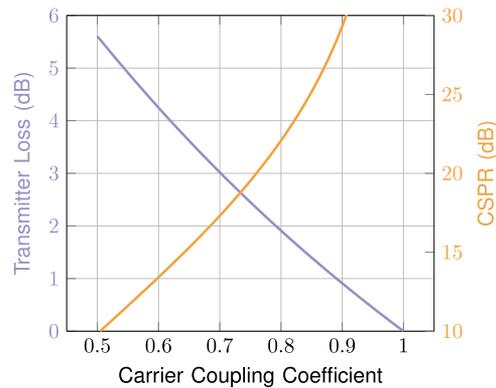


Fig. 2. Relationship between the carrier coupling coefficient, the CSPR and the loss across the transmitter, for a modulator loss of 10 dB.

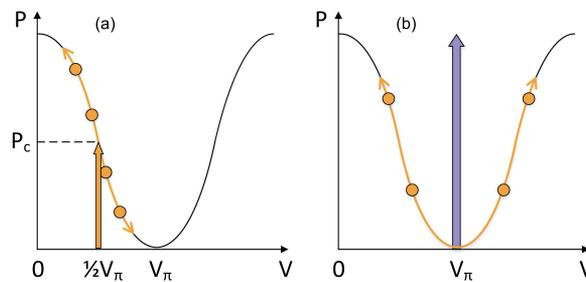


Fig. 3. Illustration of a PAM-4 signal modulated using a MZM biased at the quadrature point, (a), and the null point, (b). In (a), carrier power P_c is allowed through the MZM as part of the signal waveform, all shown here in orange. In (b), the MZM attenuates the intrinsic carrier and instead a strong carrier tone is provided by the SCT, shown in purple.

where a highly-linear modulation response is required in order to reconstruct the time-domain sinc function [24]. To achieve this response in a DD system, a simple intensity modulator is biased at its quadrature point. This is illustrated in Fig. 3(a), where a PAM-4 signal is overlaid on a Mach Zehnder modulator (MZM) transfer function. By biasing at the quadrature, carrier power P_c is permitted to pass through the modulator. The signal itself is modulated uniformly around this bias point, and for best performance should not extend into the nonlinear transfer regions close to 0 and V_π volts [25], [26]. Typical demonstrations of Nyquist-shaped, quad-biased transmission require CSPRs of 9 dB and greater before any SSBI correction is applied [25]–[29]. As the carrier power is fixed at P_c , these high CSPRs can only be reached by limiting the signal driving voltage.

This problem can be avoided by biasing the MZM at its null point. In this case, the modulator heavily suppresses the optical carrier and the signal waveform is modulated symmetrically around V_π . This is illustrated in Fig. 3(b), where the signal waveform, shown in orange, can be strongly driven around V_π while remaining linear. However, without an intrinsic optical carrier inserted at the transmitter, the PAM-4 symbols with equal amplitude cannot be distinguished when detected using a single-ended photodiode. This is solved by using the SCT, which provides a strong, phase-synchronous carrier term separate from the modulation procedure, shown in purple in Fig. 3(b). Furthermore, the power of this carrier tone can be varied separately from the signal, meaning the driving signal amplitude does not need to be compromised to reach a target CSPR. For the MZM-based demonstrations in this paper, the SCT is set to operate at the null point. In contrast, the comparison work performed using the 56 Gb/s Nyquist-shaped PAM-4 signal is generated with the MZM biased at the quadrature point.

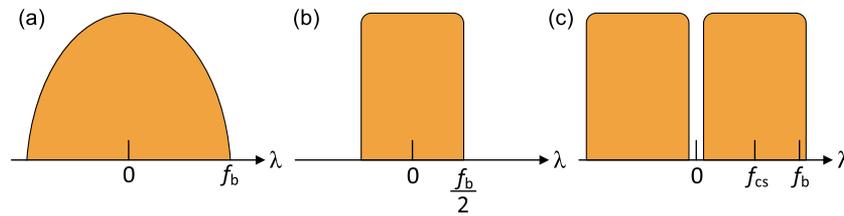


Fig. 4. Generation procedure for double sideband Nyquist-shaped subcarrier modulation. The baseband unshaped QAM signal at a symbol rate of f_b , (a), is Nyquist shaped using an RRC filter, (b), then upconverted using subcarrier frequency f_{sc} to a DSB signal, (c). The two sidebands in (c) are Hermitian symmetric, forming a real signal that can be generated using an intensity modulator.

2.3 Double Sideband Nyquist-Shaped 16-QAM Subcarrier Modulation

Just like a conventional modulator biased at its quadrature point, the SCT generates a self-homodyne signal. Therefore, it is possible to implement advanced modulation formats such as QAM, while using low complexity optoelectronic hardware for the receiver, in both cases. When received through DD, beating between the carrier and signal tones allows for the recovery of the phase-encoded information. In this paper, N-SCM 16-ary QAM (16-QAM) is used to compare against Nyquist 56 GBd PAM-4 and better understand the impact of CSPR on performance. For transmission links in the C-band, over which chromatic dispersion can be significant, single sideband modulation has been proposed as it allows for chromatic dispersion compensation [20]–[22], [25], [26], [29]. However, this is not necessary for intra-datacentre links, for which transmission distances of up to a few kilometres are usually considered. For example, dispersion-related penalties over 2 km transmission links for Nyquist-shaped, 56 GBd signals can be kept below 1 dB [4]. As dispersion compensation is not required over these distances, it is possible to use double sideband (DSB) modulation with N-SCM QAM signals, which offers an SNR improvement over single sideband modulation in terms of back to back (BTB) performance [30]. Furthermore, DSB signals can be generated using a single MZM, keeping system hardware requirements comparable with Nyquist 56 GBd PAM-4. The generation procedure in the frequency domain for DSB N-SCM is shown in Fig 4. A root raised cosine (RRC) filter is applied to a digital QAM signal, improving the spectral efficiency. The signal is then upconverted to a subcarrier frequency f_{sc} , creating a DSB signal which can be generated using a single MZM. The upconversion can be performed using either a digital subcarrier tone before generation by the digital-to-analogue converter (DAC), or an analogue radio frequency (RF) mixer afterwards.

3. Experimental Setup

The ideal implementation of the SCT is on an integrated photonic chip. This would ensure polarisation and phase matching between the carrier and signal branches using the simplest hardware. In this paper, a dual polarisation IQ MZM (DP-IQ MZM) was used to match the desired design as closely as possible. The experimental setup used to achieve this is shown in Fig. 5. An external cavity laser (ECL) with 16 dBm output power and linewidth of approximately 100 kHz, emitting at 1552 nm, was directed through a variable optical attenuator (VOA) and into a DP-IQ MZM. The Y-polarisation branch of the modulator was driven using a 92 GS/s DAC, generating the signal waveform under test, biased at the null. The signal branch and data are shown in orange in Fig. 5. The X-polarisation of the DP-IQ MZM was biased at the IQ-modulator's transmission maximum to allow as much of the carrier tone through as possible, shown in purple. The X- and Y-polarisations were then recombined using a polarisation beam combiner (PBC) co-packaged with the DP-IQ MZM. In this way, an integrated SCT was emulated that ensures polarisation and phase stability, contrasting with previous works which use discrete components [12], [22], [23].

A side-effect of this design is that the carrier and signal tones exit the DP-IQ MZM on orthogonal polarisations; this is illustrated isometrically in Fig. 5. To correct this, the light was passed through

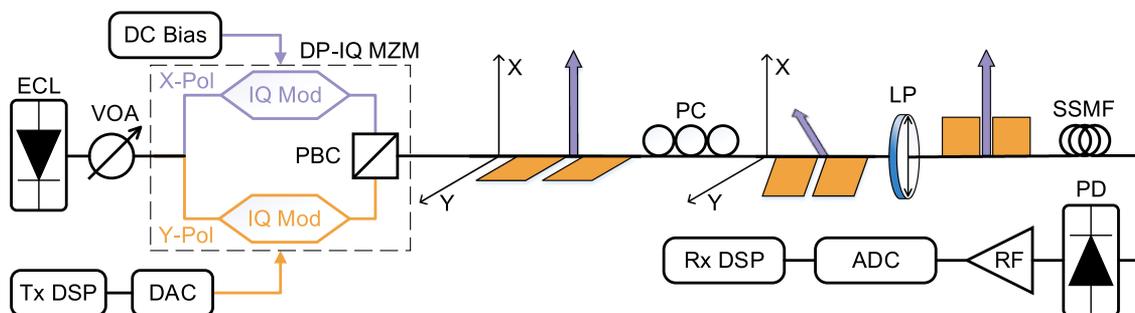


Fig. 5. Experimental setup used to demonstrate the functionality of the SCT. A dual-polarisation IQ MZM (DP-IQ MZM) was used to emulate the integrated photonic chip design required by the SCT. A polarisation controller (PC) and linear polariser (LP) were used to recombine the carrier tone (purple) and signal tone (orange) while simultaneously controlling the CSPR.

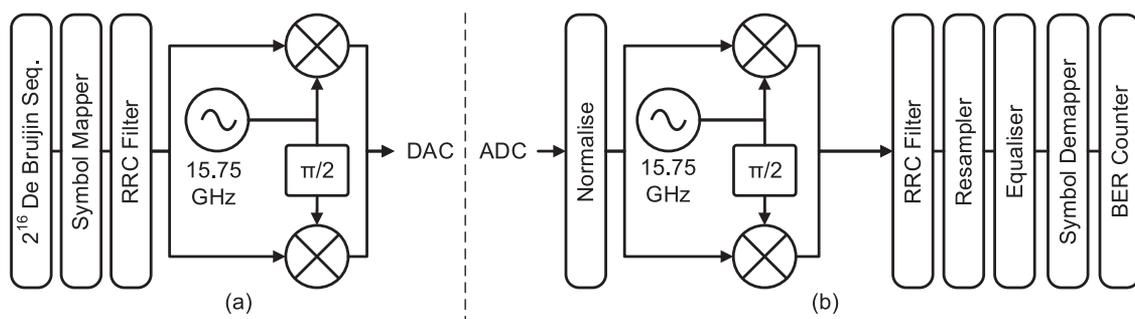


Fig. 6. Transmitter (a) and Receiver (b) DSP used for the SCT-PAM-4, the SCT-16-QAM, and the QB-PAM-4. A digital clock tone was used to up- and down-convert the data to the subcarrier for the SCT-16-QAM before demapping. These steps were not required by the PAM-4.

a polarisation controller (PC) followed by a linear polariser (LP) with an extinction ratio of 30 dB. By manually rotating the PC, a proportion of the carrier and signal tones were both allowed through the LP, after which they are polarisation matched. This setup also allowed control of the CSPR, measured using an optical spectrum analyser (OSA). In a purpose built, integrated SCT, this process would not be required as the carrier and signal tones would be directly recombined in the same polarisation. It should be noted that the same result could be achieved more simply using a single polarisation IQ MZM and modulating only one arm while strongly biasing the other; however a MZM of this type with sufficient bandwidth was not available at the time of the experiment.

The single polarisation SCT output was then passed through 2 km of standard single mode fibre (SSMF) prior to detection with a 40 GHz bandwidth PIN photodiode. Following photodetection, the signal was amplified using an RF amplifier with 17 dB gain and 6 dB noise figure. Finally, the signal was digitised using a 80 GS/s, 33 GHz ADC before the offline DSP was performed.

Two modulation formats were tested with the SCT. The first was a 56 GBd Nyquist-shaped PAM-4 waveform, henceforth referred to as SCT-PAM-4. This waveform was selected for direct comparison against the conventional quad-biased 56 GBd PAM-4 format (see Fig. 3). the DSP used to generate this format is shown in Fig. 6(a). A 2^{16} De Bruijn sequence was mapped to PAM-4 symbols and subsequently shaped using a Nyquist RRC filter with a roll off factor of 0.01. This was then uploaded to the DAC for transmission. Both arms of the Y-polarisation IQ modulator were driven with the same DAC signal, ensuring all light passing through the signal branch was modulated identically. The two Y-polarisation signal bias ports were biased at their null points, with the phase bias set to maximise constructive interference. At reception, a 2^{18} bit capture was saved for offline processing. The DSP used to achieve this is shown in Fig. 6(b). The received signal was normalised then passed through

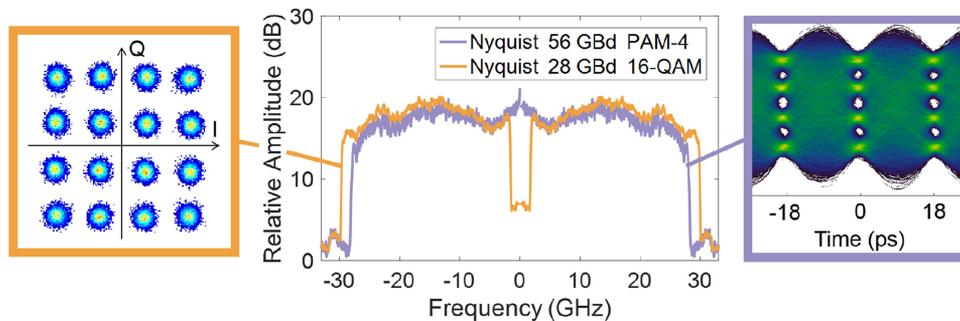


Fig. 7. Received signal spectra under test with corresponding modulation formats, received at 0 dBm. The post-equaliser Nyquist 28 GBd DSB 16-QAM (orange) occupies 60 GHz of bandwidth, while the post-equaliser Nyquist 56 GBd PAM-4 (purple) occupies 56 GHz. Both signals were generated using similar equipment and DSP. The PAM-4 was tested using the SCT as well as a quad-biased MZM.

a matched filter before being resampled to two samples per symbol. A 21-tap decision-directed LMS filter was used to recover the PAM-4 signal and account for frequency roll-off. The resulting received spectrum and eye diagram are shown in Fig. 7. Following this, the symbols were demapped and a bit-to-bit comparison of transmitted and recovered bits was performed to measure the BER.

The second modulation format tested with the SCT was a 28 GBd DSB N-SCM 16-QAM signal (SCT-16-QAM). This was selected to demonstrate the phase stability of the SCT and its usefulness in supporting advanced modulation formats. To keep the systems under test as similar as possible, the DSP used to generate the SCT-16-QAM was kept largely the same as that used for the SCT-PAM-4. However, to encode the complex IQ signal onto a real waveform, the baseband Nyquist signal was digitally upconverted using a subcarrier frequency, f_{sc} , of 15.75 GHz, corresponding to a subcycle ratio (f_{sc}/f_b) of 0.5625. This value was chosen for experimental ease, as it places a 3.5 GHz gap between the sidebands, while ensuring an integer number of subcycles have been performed across the whole of the De Bruijn sequence; that is, pattern length \times subcycle ratio = round(pattern length \times subcycle ratio). This allowed for BER analysis across multiple neighbouring sequences within the large single capture of 2^{18} bits, improving result accuracy. Similarly, the received signal was digitally downconverted using a second 15.75 GHz clock tone. As before, the signal was matched filtered and resampled, then passed through an adaptive constant modulus equaliser. Due to the use of self-homodyne detection there is no residual phase noise; therefore, no carrier phase recovery is required. Instead, as part of the equalisation process, the 16-QAM constellation was rotated by a constant phase offset. This was followed by symbol demapping and finally, the BER estimator. The resulting received spectrum and constellation diagram are shown in Fig. 7.

To contrast the proposed SCT with a conventional transmitter design, a quadrature-biased 56 GBd Nyquist-shaped PAM-4 signal (QB-PAM-4) was also generated. As with the SCT, the PAM-4 signal was Nyquist-shaped with a roll-off factor of 0.01 and was applied to just the Y-polarisation of the DP-IQ MZM. As before, both arms of the Y-polarisation IQ modulator were driven with the same signal, with the phase bias set to ensure full constructive interference. Both bias ports of the signal driving arms were biased at their quadrature points such that a CSRR of 12 dB was set; this value was selected to optimise the received signal in terms of BER. The X-polarisation IQ modulator was biased at the null point. After the modulator, the signal passed through the PC and LP to ensure complete suppression of the X-polarisation carrier, before being passed through the 2 km of SSMF. Following detection, a 21-tap decision-directed LMS filter was used to recover the PAM-4 signal. In this way, the SCT-PAM-4, SCT-16-QAM and the QB-PAM-4 were all generated using the same modulator, with similar bandwidth and DSP requirements, validating side-by-side comparison.

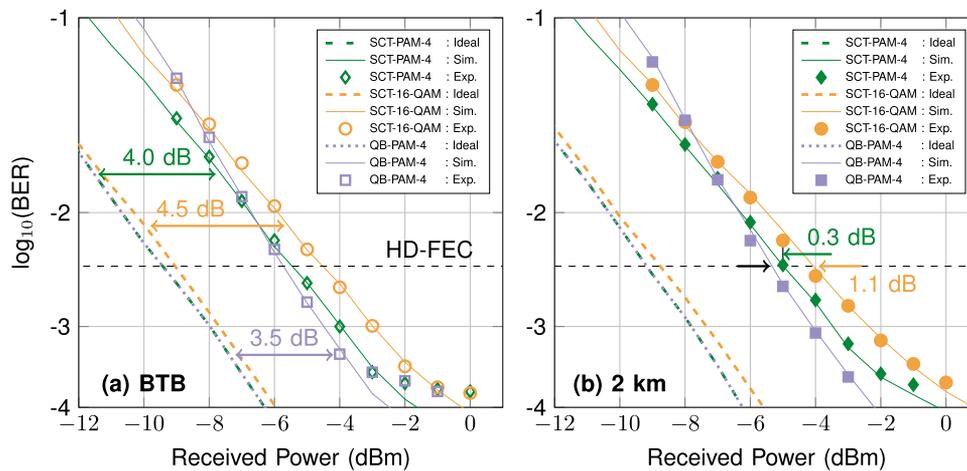


Fig. 8. Ideal simulation, practical simulation, and experimental results for the SCT-PAM-4, SCT-16-QAM and QB-PAM-4 112 Gb/s systems in BTB (a) and over 2 km (b). The ideal case excludes transmitter noise, bandwidth limitations and RF amplifier noise.

4. Results and Discussion

4.1 Performance With Received Optical Power

The performance of the SCT-PAM-4, SCT-16-QAM and QB-PAM-4 systems were compared through Monte Carlo simulation. The DSP steps shown in Fig. 6 were used for signal generation and reception. Shot and thermal noise contributions were modelled at the photodiode assuming a dark current value of 10 nA and device temperature of 298 K [31]. A responsivity of 0.6 A/W was used to match the photodiode used in experiment. In the 2 km transmission case, the effects of dispersion were modelled using a dispersion parameter of 16 ps/(nm.km). For the SCT results, the CSPR was optimised for each received power by sweeping a range of CSPRs and selecting the best result; this is discussed in more detail in Section 4.2. The ideal case performance of each system was considered first by ignoring practical effects such as electronic transmitter noise, frequency roll-off and electronic noise from the RF amplifier. For this ideal case assessment, the QB-PAM-4 was also swept through a range of CSPR values and the best measurement selected. The results for the BTB and 2 km cases including the hard decision forward error correction (HD-FEC) threshold [32] at a BER of 3.8×10^{-3} are shown in Fig. 8(a) and (b), respectively. In both Fig. 8(a) and (b), the SCT-PAM-4 and QB-PAM-4 systems are predicted to perform identically; this is expected behavior as the waveforms generated in each case are, fundamentally, identical. The ideal-case SCT-16-QAM is seen to experience a 0.3 dB penalty compared to the PAM-4 formats; this is attributed to SSBI.

Following this, an experimental comparison was performed using the testbed described in Section 3. For the two SCT formats, the CSPR was swept through a range of values to optimise the performance at each received power. The received power was fixed during each of these sweeps using the VOA shown in Fig. 5. The QB-PAM-4 transmission was set to a fixed CSPR of 12 dB as described in Section 3. The results of this experiment are shown in Fig. 8 as markers. Practical simulation results are overlaid for comparison as solid lines. Frequency roll-off was simulated using a 5th-order Bessel electrical low-pass filter with 28 GHz bandwidth at the transmitter and receiver. To model transmitter noise, an optical signal-to-noise ratio (OSNR) of 26 dB was set for all signal data. When modelling the SCT, the carrier optical source was set to a finite OSNR of 40 dB, estimated by measuring the un-modulated signal after the DP-IQ MZM on an OSA at a 0.1 nm resolution. Thermal noise equivalent to the RF amplifier's 6 dB noise figure was also included [31]. A fixed CSPR of 12 dB was set for the QB-PAM-4 system, matching the experiment.

TABLE 1
Received Powers Required to Reach HD-FEC Performance in Fig. 8

Format	BTB: ideal (dBm)	BTB: measured (dBm)	2 km: ideal (dBm)	2 km: measured (dBm)
QB-PAM-4	-9.4	-5.7	-9.1	-5.3
SCT-PAM-4	-9.4	-5.4	-9.1	-5.0
SCT-16-QAM	-9.1	-4.5	-8.8	-4.2

From Fig. 8(a) and (b), it can be seen that the practical simulation results agree well with the experimental data. Comparing the ideal vs experimental SCT formats, a 4.0 dB penalty is measured for the SCT-PAM-4, while the SCT-16-QAM experiences a 4.5 dB penalty. The most significant contributor to these penalties was determined through simulation to be the 6 dB noise figure of the RF amplifier, which accounts for 3 dB of the difference between ideal and practical results in both systems. Both systems suffer a further 0.5 dB penalty due to the finite OSNR of the split-carrier. Transmitter noise and bandwidth roll-off account for the final 0.5 dB of penalty for the SCT-PAM-4, which grows to 1.0 dB for the SCT-16-QAM due to additional attenuation experienced in the 28-30 GHz frequencies.

The QB-PAM-4 is observed to outperform the SCT at received powers above -7 dBm, reaching a minimum penalty of 3.5 dB vs the ideal case at a BER of 4.8×10^{-4} . Common with the SCT results, 3 dB of this is attributed to the RF amplifier and 0.5 dB to transmitter impairments, allowing the QB-PAM-4 to outperform the SCT by avoiding additional noise from the split-carrier. The performance of the QB-PAM-4 system is observed to decrease at lower received powers due to its fixed CSPR, well modelled in the simulation. All the tested formats are observed to reach an error floor at a BER of 1.5×10^{-4} due to the finite maximum SNR achievable from our transceiver subsystems.

From Fig. 8(b), it can be seen that the QB-PAM-4 transmission after 2 km reaches the HD-FEC threshold at -5.3 dBm. Accounting for the 3 dB penalty attributed to the RF amplifier used in this experiment, this result is observed to be in good agreement with 56 GBd PAM-4 results in the literature [8], [9]. The QB-PAM-4 results presented in this paper are therefore considered a valid comparison case for assessing the SCT's performance. Fig. 8(b) also shows the SCT-PAM-4 reaches the HD-FEC threshold at -5.0 dBm, a penalty of just 0.3 dB. This result demonstrates that the SCT can be considered competitive with QB-PAM-4 in terms of received optical power, while making all of the loss budget gains discussed in Section 2.1. These gains are quantified in Section 4.3. The experimental SCT-16-QAM results see a 1.1 dB penalty, consistent with the differences observed in Fig. 8(a). The results presented in Fig. 8 are summarised in Table 1, which records the received powers required by each format to reach the HD-FEC threshold.

4.2 Performance With Carrier-to-Signal Power Ratio

As discussed in Section 2.1, the CSPR must be optimised in order to minimise signal-signal beating interference while preventing transceiver noise from dominating. To quantify this, CSPR was swept through a range of values and BER measured for five different received powers for both of the SCT systems under test. The results for the SCT-PAM-4 are shown in Fig. 9 along with simulated performance. Good agreement is observed between the simulated and experimental results, with the deviation observed at higher CSPRs attributed to inaccurate setting of the CSPR. Experimental eye diagrams are overlaid at a received power of -3 dBm to visualise the effect of CSPR on signal integrity. At the relatively low CSPR of 7 dB, the eye diagram is observed to be squashed, closing the lower eyes. This is attributed to the non-linear response of the receiving photodiode. At the high CSPR of 19 dB, the eye diagram is linear, though the eyes are less open. This is due to insufficient power being passed to the modulator, producing an excess of transmitter noise, uniformly affecting the signal. Between these two extremes the optimum CSPR is shown at 13 dB; here the eye diagram is seen to be mostly linear, with the most open eyes.

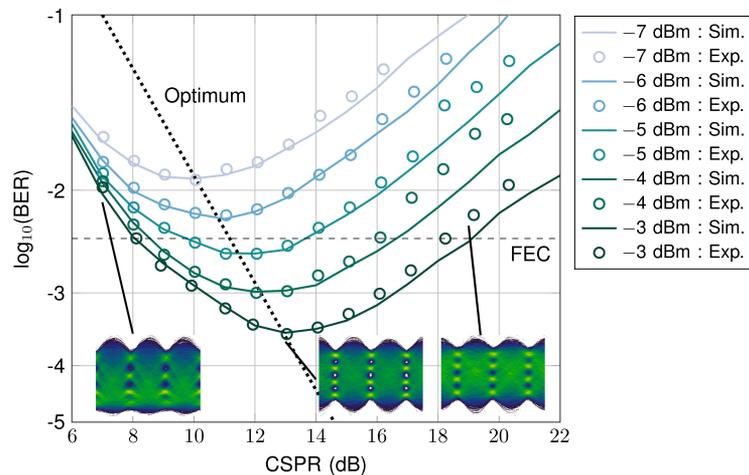


Fig. 9. Simulated and experimental BTB performance of the SCT-PAM-4 with varied CSRR for range of received powers. The dotted black line highlights the trend of optimum CSRRs across the different received powers. Experimental constellation diagrams for the -3 dBm data at CSRRs of 7 dB, 13 dB and 19 dB are also included, showing how CSRR effects signal quality.

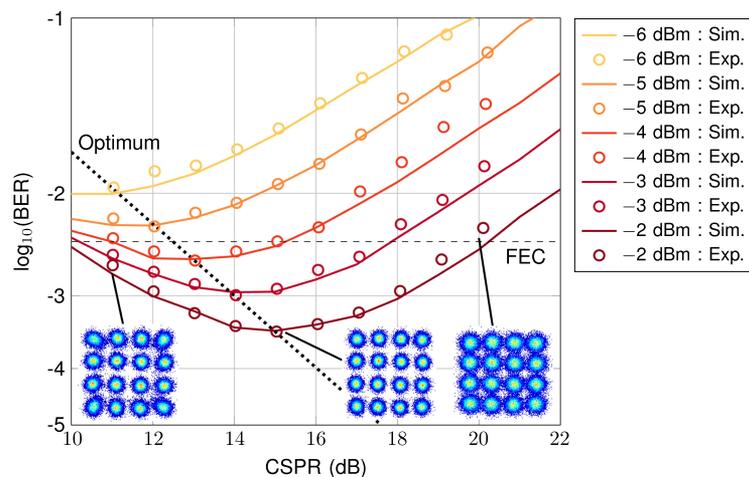


Fig. 10. Simulated and experimental BTB performance of the 28 GBd, 16-QAM SCT with varied CSRR for range of received powers. The dotted black line highlights the optimum CSRR for each received power. Experimental constellation diagrams for the -2 dBm data at CSRRs of 11 dB, 15 dB and 20 dB are also included, showing the effect of SSBI and transmitter noise on the signal.

Overlaid on Fig. 9 is a dotted line marking the progression of optimum CSRRs as received power increases. The experimental data points closest to this optimum are the data shown previously in Fig. 8(a). Across the 5 dB of received powers plotted, the optimum CSRR increases by 4 dB, showing that the SCT mostly benefits by adding power to the carrier branch rather than the signal. This indicates that it is only necessary to have sufficient signal power to overcome transmitter noise; once that has been achieved extra power should be allocated to the carrier in order to improve linearity and benefit from signal-carrier beating.

The CSRR dependence of the SCT-16-QAM system was also investigated through simulation and experiment, using the same method described for the SCT-PAM-4 results above. These results are shown in Fig. 10. Again, good agreement is observed between the simulated and experimental results, with the deviation observed at higher CSRRs attributed to inaccurate setting of the CSRR.

Experimental constellation diagrams for a received optical power of -2 dBm are also shown. At the relatively low CSPR of 11 dB the constellation is observed to be distorted, with the high amplitude points pulled towards the centre. This is attributed to SSBI. At a CSPR of 20 dB, the SSBI distortion is negligible. However, the overall noise present is uniformly higher, representative of the signal tone being lost to transmitter noise. The constellation at 15 dB shows the optimum position.

As with Fig. 9, the dotted line overlaid on Fig. 10 marks the optimum position for each received power; the data along this line are the same points plotted in Fig. 8(a). Noting the linearity of this line, it is observed that as the received power increases, the optimum CSPR also increases correspondingly. As with the data shown in Fig. 9, it is observed that as the available optical power increases, it is always advantageous to place it in the carrier branch in order to keep SSBI to a minimum. From this conclusion the value of the SCT design becomes apparent, as higher CSPRs result in lower transmitter loss, which in turn increases received power (see Fig. 2). It should also be noted from both Fig. 9 and Fig. 10 that the BER of the system varies slowly around the optimum value, meaning exact calibration of the CSPR is not critical. Expanding on this, increasing the CSPR by 1 dB will have a small effect on performance with respect to transmitter noise, but will allow ~ 1 dB more optical power onto the receiver, improving performance. The true optimum performance is therefore not defined by received optical power, but by the CW power available. This is investigated in Section 4.3.

Comparing Fig. 9 and Fig. 10, it is apparent that for a given target performance, the SCT-PAM-4 requires a lower CSPR than the SCT-16-QAM. This is particularly clear when considering the -4 dBm result in Fig. 9 to the -3 dBm result in Fig. 10, both of which just reach a BER of 1×10^{-3} at their optimum CSPRs. The SCT-PAM-4 reaches its optimum at a CSPR of 12 dB, while the SCT-16-QAM requires 14 dB. This difference is important for the same reason discussed above - that, in the case of the SCT, a higher CSPR equates to a lower transmitter loss. Therefore, although it has been established in Section 4.1 that the SCT-PAM-4 outperforms the SCT-16-QAM by 0.8 dB in terms of received power, for a fixed CW power the SCT-16-QAM setup will have close to 2 dB more optical power available at the receiver. This again emphasises the importance of analysing the SCT for a fixed CW power, as performed in Section 4.3.

4.3 Performance With CW Power

In Section 2.1, it was explained how the SCT's carrier coupling coefficient determines both the CSPR and the received optical power; i.e. one cannot be varied without changing the other. Furthermore, Section 4.2 demonstrates that both these terms effect the final BER. It is therefore insufficient to assess the transmitter performance in terms of the traditional metric of received power. Instead, the required metric is how the transmitter performs for a fixed CW launch power, which is of course the conserved quantity in a real system. To reiterate, the CW power is the optical power emitted by the laser before any modulation or splitting, as indicated in Fig. 1. It is this value that must be optimised by varying the carrier coupling ratio. It was not possible to vary the carrier coupling coefficient of the DP-IQ MZM used in this experiment; however, the simulation results presented in Fig. 9 and Fig. 10 show good agreement with experiment, particularly around the optimum performance points. Therefore, further Monte Carlo simulations were carried out for a range of fixed CW powers, using Eq. (3) and Eq. (4) to determine the received optical power corresponding to each CSPR. The results for the SCT-PAM-4 system are shown in Fig. 11. As before, a dotted line is overlaid to mark the optimum CSPR for each CW power. Simulated eye diagrams for a CW power of 1 dBm are also overlaid to show the expected signal integrity as CSPR is varied. Note that compared to Fig. 9, the CSPR for a target performance is considerably higher when the CW power is fixed. This can be appreciated by comparing the -3 dBm result plotted in Fig. 9 to the 1 dBm result shown in Fig. 11, both of which reach an optimum BER at around the 3×10^{-4} level. For a fixed received power (Fig. 9), the SCT-PAM-4 requires a CSPR of 13 dB; however, this rises to 16 dB when the CW is fixed (Fig. 11). This is a significant result, as it shows that using a higher CSPR is actually a better use of a SCT's finite CW power, despite this ratio being non-optimal in terms of received power.

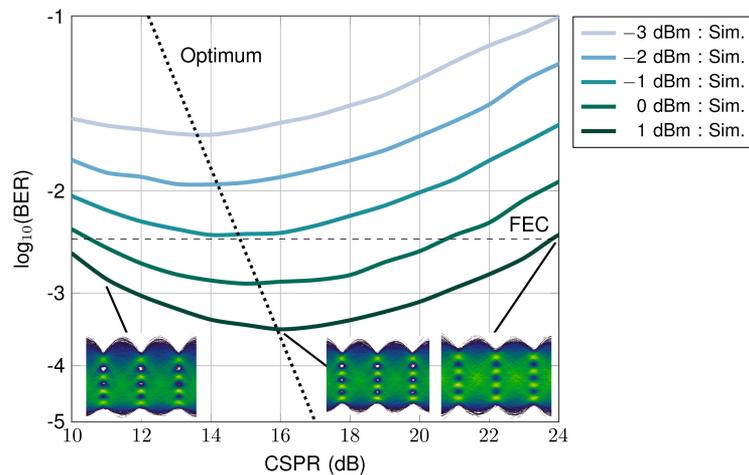


Fig. 11. Simulated BTB performance of SCT-PAM-4 for a range of CW powers. By varying the CSRR (analogous to changing the carrier coupling ratio) the performance can be optimised for received power, nonlinearity and transmitter noise. Simulated eye diagrams for a CW power of 1 dBm are included to visualise the effect of CSRR on signal quality.

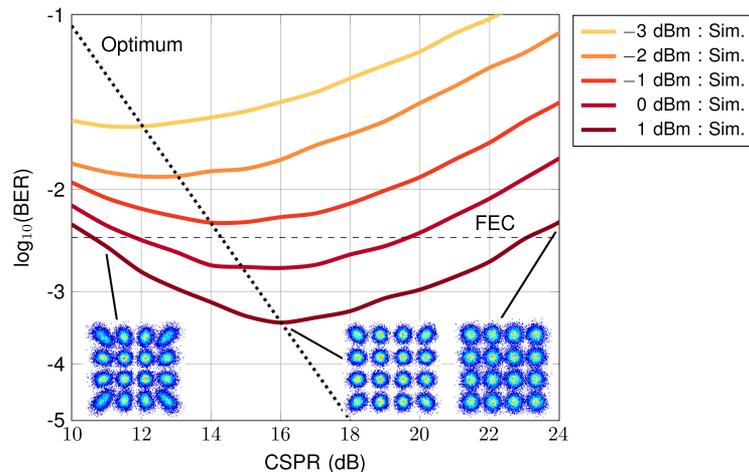


Fig. 12. Simulated BTB performance of the 28 GBd, 16-QAM SCT for a range of CW powers. By varying the CSRR (analogous to changing the carrier coupling ratio) the performance can be optimised against received power, SSBI and transmitter noise. Simulated constellation diagrams for a CW power of 1 dBm are included to visualise the trade off between SSBI and transmitter noise.

The fixed CW power simulation results for the SCT-16-QAM format are presented in Fig. 12. These results are similar to the SCT-PAM-4 results presented in Fig. 11, with an optimum CSRR of 16 dB used to reach an BER of 4×10^{-4} with 1 dBm of available power. This contrasts with the comparisons made in Section 4.1, where the SCT-PAM-4 outperformed the SCT-16-QAM by 0.8 dB, and also the results presented in Section 4.2, where a 2 dB difference in CSRR was noted. This again shows that the SCT performs best when used with high CSRRs, where the performance in terms of received optical power is compromised to minimise modulator loss.

By selecting the optimum values at each CW power in Fig. 11 and Fig. 12, the optimum performance of the SCT can be plotted as a function of the CW power available. This measure accounts not just for the effects of CSRR and received optical power, but all other losses in the system, as shown in Eq. (1). This is the most useful metric for the SCT, as it can be compared directly against alternative designs, such as QB-PAM-4, which are also constrained by CW power. Fig. 13 plots this

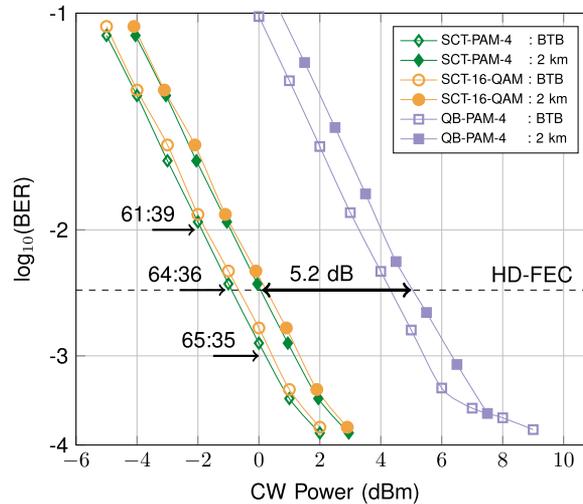


Fig. 13. Simulated BTB and 2 km performance for a given CW power of the SCT-PAM-4, the SCT-16-QAM, and QB-PAM-4 systems. Both SCT formats have been optimised for CSRR at each point, based on the experimentally measured results in Section 4.2. The SCT-PAM-4 is observed to outperform the QB-PAM-4 system by 5.2 dB, and is HD-FEC ready with only 0 dBm CW power. The ideal carrier:signal coupling coefficients required to reach target performance levels are also shown.

TABLE 2

Continuous Wave Powers Required to Reach HD-FEC Performance in Fig. 13

Format	BTB (dBm)	2 km (dBm)
QB-PAM-4	4.3	5.2
SCT-PAM-4	-0.9	0.0
SCT-16-QAM	-0.7	0.2

comparison, where an illustrative modulator loss of 10 dB is set for the SCT-PAM-4, the SCT-16-QAM and the QB-PAM-4 systems. A measured $P_{\text{fibre,loss}}$ of 0.5 dB is also included for all systems in the 2 km case. All curves are informed by the experimental performance values presented in Fig. 8. In fact, the QB-PAM-4 results are the same data as presented in Fig. 8, with the losses simply added to infer the required CW power. The SCT results, meanwhile, use the experimentally-validated simulation results from Section 4.2 to optimise for modulation loss and receiver performance. Ideal coupling losses are assumed, as described by Eq. 4. The CW powers required by each format to reach the HD-FEC BER are recorded in Table 2.

From Fig. 13, both SCT formats are observed to perform nearly identically, with a difference of 0.2 dB. This shows the performance of a well-calibrated SCT to be relatively format independent, even given the differences in received power performance presented in Fig. 8(b). When compared to the QB-PAM-4 system, the SCT-PAM-4 is observed to achieve the same performance over 2 km while using 5.2 dB less CW power. This is equivalent to a 5.2 dB improvement in loss budget for this 112 Gb/s link, a significant improvement given the similar levels of hardware and identical DSP between the two systems.

Furthermore, the SCT-PAM-4 simulation reaches the HD-FEC threshold at a CW power of 0 dBm. This exceptionally low power requirement raises the possibility of using ultra-low complexity laser sources with limited output powers, such as vertical cavity surface emitting lasers, or using a single high-power laser to power multiple modulators. This could be important for future electrical datacentre switches, where on-board co-packaged optical transmitters will be densely packed around an integrated circuit, each requiring an external optical power supply [3], [33].

Fig. 13 also displays the calculated carrier-to-signal coupling ratios required to achieve specific BER performance targets. At the HD-FEC threshold this value is 64:36, where 65:35 results in an optimum BER of 1×10^{-3} . Given the sensitivity of the optimum BER to this ratio, it is likely advisable to aim for a marginally higher carrier coupling coefficient to leave room for manufacturing variation, additional path loss, and unideal coupling losses (not included within this simulation). Overall, this 5.2 dB improvement represents a 70% reduction in the laser power required for a 112 Gb/s link, reducing the power requirements of optical intraconnects that are ubiquitous in the datacentre environment.

5. Summary

A split-carrier transmitter is proposed to improve the power budget of 400G intra-datacentre connections. Through the emulation of an integrated device, a 112 Gb/s self-homodyne signal was generated using a single MZM and detected using a single photodiode. Two modulation formats were evaluated and compared: Nyquist 56 GBd PAM-4 and Nyquist 28 GBd 16-QAM. As a further comparison, a conventionally modulated, quadrature-biased Nyquist 56 GBd PAM-4 signal was also generated using similar equipment and DSP. In both simulations and experiments, the two PAM-4 systems were observed to perform similarly at the HD-FEC threshold, in terms of received optical power, after 2 km of transmission. The QB-PAM-4 outperformed the SCT-16-QAM by 1.1 dB. In terms of CW laser power, the SCT was observed through simulation to reach the HD-FEC threshold using just 0 dBm, outperforming the QB-PAM-4 by 5.2 dB. This represents an equivalent increase in loss budget, offering a 70% reduction in the required CW power to achieve the target performance. The proposed system demonstrates the feasibility of using ultra-low complexity, low energy lasers and laser power sharing in intra-datacentre connections.

Acknowledgment

The provision of the dual polarisation IQ MZM by Oclaro (now Lumentum) is gratefully acknowledged.

References

- [1] IEEE, "P802.3bs 400 Gigabit Ethernet task force," Jul. 2017. [Online]. Available: <http://www.ieee802.org/3/bs/>
- [2] M. Wu, "112Gbps serial transmission over copper PAM4 vs PAM8 signaling", DesignCon 2017, Santa Clara, CA, USA, 2017. [Online]. Available: <https://www.xilinx.com/publications/events/designcon/2017/112gbps-serial-transmission-over-copperpam4-vs-pam8-slides.pdf>
- [3] H. J. S. Dorren, E. H. M. Wittebol, R. de Kluijver, G. G. de Villota, P. Duan, and O. Raz, "Challenges for optically enabled high-radix switches for data center networks," *J. Lightw. Technol.*, vol. 33, no. 5, pp. 1117–1125, Mar. 2015.
- [4] R. C. Figueiredo, A. L. N. Souza, S. M. Ranzini, A. Chiuchiarelli, L. H. H. Carvalho, and J. D. Reis, "Investigation of 56-Gbd PAM4 bandwidth and chromatic dispersion limitations for data center applications," in *Proc. SBMO/IEEE MTT-S Int. Microw. Optoelectron. Conf.*, Aug. 2017, doi: [10.1109/IMOC.2017.8121057](https://doi.org/10.1109/IMOC.2017.8121057).
- [5] C. Cole, "Ideal SNR Penalties," in *Proc. 200 GbE & 400 GbE Task Force IEEE P802.3bs Interim Meeting*, 2014. [Online]. Available: http://www.ieee802.org/3/bs/public/14_09/
- [6] J. Verbist *et al.*, "DAC-less and DSP-free 112 Gb/s PAM-4 transmitter using two parallel electroabsorption modulators," *J. Lightw. Technol.*, vol. 36, no. 5, pp. 1281–1286, Mar. 2018.
- [7] D. Sadot, G. Dorman, A. Gorshtein, E. Sonkin, and O. Vidal, "Single channel 112Gbit/sec PAM4 at 56Gbaud with digital signal processing for data centers applications," *Opt. Express*, vol. 23, no. 2, pp. 991–997, Jan. 2015. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-23-2-991>
- [8] M. A. Mestre *et al.*, "Compact InP-based DFB-EAM enabling PAM-4 112 Gb/s transmission over 2 km," *J. Lightw. Technol.*, vol. 34, no. 7, pp. 1572–1578, Apr. 2016.
- [9] K. Zhong *et al.*, "Experimental study of PAM-4, CAP-16, and DMT for 100 Gb/s short reach optical transmission systems," *Opt. Express*, vol. 23, no. 2, pp. 1176–1189, Jan. 2015. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-23-2-1176>
- [10] N. Eiselt *et al.*, "Performance comparison of 112-Gb/s DMT, Nyquist PAM4, and partial-response PAM4 for future 5G ethernet-based fronthaul architecture," *J. Lightw. Technol.*, vol. 36, no. 10, pp. 1807–1814, May 2018.
- [11] C. Yang *et al.*, "IM/DD-based 112-Gb/s PAM-4 transmission using 18-Gbps DML," *IEEE Photon. J.*, vol. 8, no. 3, Jun. 2016, Art. no. 7903907.
- [12] Z. Liu, T. Xu, G. Saavedra, and P. Bayvel, "448-Gb/s PAM4 transmission over 300-km SMF-28 without dispersion compensation fiber," in *Proc. Opt. Fiber Commun. Conf. Expo.*, Mar. 2018, Paper W1J.6.

- [13] J. Huo *et al.*, "Transmitter and receiver DSP for 112 Gbit/s PAM-4 amplifier-less transmissions using 25G-class EML and APD," *Opt. Express*, vol. 26, no. 18, pp. 22 673–22 686, Sep. 2018. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-26-18-22673>
- [14] Y. Gao, J. C. Cartledge, S. S. H. Yam, A. Rezaia, and Y. Matsui, "112 Gb/s PAM-4 using a directly modulated laser with linear pre-compensation and nonlinear post-compensation," in *Proc. Eur. Conf. Exhib. Opt. Commun.*, 2016, Paper M.2.C.2. [Online]. Available: <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7767500>
- [15] Y. Matsui *et al.*, "55 GHz bandwidth distributed reflector laser," *J. Lightw. Technol.*, vol. 35, no. 3, pp. 397–403, Feb. 2017.
- [16] F. Gao *et al.*, " 2×64 Gb/s PAM-4 transmission over 70 km SSMF using O-band 18G-class directly modulated lasers (DMLs)," *Opt. Express*, vol. 25, no. 7, pp. 7230–7237, Apr. 2017. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-25-7-7230>
- [17] W. Huang *et al.*, "93% complexity reduction of Volterra nonlinear equalizer by ℓ_1 -regularization for 112-Gbps PAM-4 850-nm VCSEL optical interconnect," in *Proc. Opt. Fiber Commun. Conf. Expo.*, Mar. 2018, Paper M2D.7.
- [18] K. Zhang, Q. Zhuge, H. Xin, W. Hu, and D. V. Plant, "Performance comparison of DML, EML and MZM in dispersion-unmanaged short reach transmissions with digital signal processing," *Opt. Express*, vol. 26, no. 26, pp. 34 288–34 304, Dec. 2018. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-26-26-34288>
- [19] A. S. Karar and J. C. Cartledge, "Generation and detection of a 56 Gb/s signal using a DML and half-cycle 16-QAM Nyquist-SCM," *IEEE Photon. Technol. Lett.*, vol. 25, no. 8, pp. 757–760, Apr. 2013.
- [20] M. Erkılınç, S. Pachnicke, H. Griesser, B. C. Thomsen, P. Bayvel, and R. I. Killey, "Performance comparison of single-sideband direct detection Nyquist-subcarrier modulation and OFDM," *J. Lightw. Technol.*, vol. 33, no. 10, pp. 2038–2046, May 2015.
- [21] Z. Li, M. Erkılınç *et al.*, "Digital linearization of direct-detection transceivers for spectrally efficient 100 Gb/s/ λ WDM metro networking," *J. Lightw. Technol.*, vol. 36, no. 1, pp. 27–36, Jan. 2018.
- [22] X. Chen *et al.*, "218-Gb/s single-wavelength, single-polarization, single-photodiode transmission over 125-km of standard singlemode fiber using Kramers-Kronig detection," in *Proc. Opt. Fiber Commun. Conf. Exhib.*, Mar. 2017, Paper Th5B.6.
- [23] T. Gerard, M. Erkılınç, Z. Liu, B. Thomsen, P. Bayvel, and D. Lavery, "A low-loss spit-carrier architecture for datacentre intraconnects," in *Proc. Eur. Conf. Opt. Commun.*, Sep. 2018, doi: [10.1109/ECOC.2018.8535385](https://doi.org/10.1109/ECOC.2018.8535385).
- [24] M. Yang *et al.*, "Optical sinc-shaped Nyquist pulses of exceptional quality," *Nature Commun.*, vol. 4, Dec. 2013, Art. no. 2898.
- [25] L. Shu *et al.*, "Single-lane 112-Gbit/s SSB-PAM4 transmission with dual-drive MZM and KramersKronig detection over 80-km SSMF," *IEEE Photon. J.*, vol. 9, no. 6, Dec. 2017, Art. no. 7204509.
- [26] M. Zhu *et al.*, "Optical single side-band Nyquist PAM-4 transmission using dual-drive MZM modulation and direct detection," *Opt. Express*, vol. 26, no. 6, pp. 6629–6638, Mar. 2018. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-26-6-6629>
- [27] Z. Li *et al.*, "Signal-signal beat interference cancellation in spectrally-efficient WDM direct-detection Nyquist-pulse-shaped 16-QAM subcarrier modulation," *Opt. Express*, vol. 23, no. 18, pp. 23 694–23 709, Sep. 2015. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-23-18-23694>
- [28] A. Mecozzi, C. Antonelli, and M. Shtaif, "Kramers–kronig coherent receiver," *Optica*, vol. 3, no. 11, pp. 1220–1227, Nov. 2016. [Online]. Available: <http://www.osapublishing.org/optica/abstract.cfm?URI=optica-3-11-1220>
- [29] S. van der Heide, A. Albores-Mejia, F. Gomez-Agis, B. Docter, and C. Okonkwo, "112-Gbit/s single side-band PAM-4 transmission over inter-DCI distances without DCF enabled by low-complexity DSP," in *Proc. Eur. Conf. Opt. Commun.*, Sep. 2017, pp. 1–3.
- [30] B. Lin, J. Li, H. Yang, Y. Wan, Y. He, and Z. Chen, "Comparison of DSB and SSB transmission for OFDM-PON [Invited]," *J. Opt. Commun. Netw.*, vol. 4, no. 11, pp. B94–B100, Nov. 2012. [Online]. Available: <http://jocn.osa.org/abstract.cfm?URI=jocn-4-11-B94>
- [31] G. Agrawal, *Fiber-Optic Communication Systems*, 3rd ed. Hoboken, NJ, USA: Wiley, 2002.
- [32] E. Agrell and M. Secondini, "Information-theoretic tools for optical communications engineers," in *Proc. IEEE Photon. Conf.*, Sep. 2018, pp. 1–5.
- [33] A. Ghiasi, "Large data centers interconnect bottlenecks," *Opt. Express*, vol. 23, no. 3, pp. 2085–2090, Feb. 2015. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-23-3-2085>