

Predistortion for High-Speed Lumped Silicon Photonic Mach-Zehnder Modulators

Brian Murray , *Member, IEEE*, Cleitus Antony , *Member, IEEE*, Giuseppe Talli , *Member, IEEE*, and Paul D. Townsend , *Member, IEEE*

Abstract—The use of electrical precompensation for high-speed operation of lumped MZMs is demonstrated for 50 Gb/s PAM4 modulation through experiment and simulation. An accurate equivalent circuit model is fitted to S11 measurements and then used to analyse design trade-offs in terms of system performance, energy consumption and drive voltage requirements. If precompensation is not used, careful selection of source resistance with respect to electrode and packaging inductance is shown to be crucial for high-speed modulation. Increasing the electrode inductance is shown to be beneficial as it introduces inductive peaking, which allows optimum performance with a higher source impedance. Furthermore, combining precompensation with a slight inductive peaking and optimised source resistance is demonstrated to reduce drive voltage requirements and further reduce the energy dissipated in the equivalent circuit.

Index Terms—Electro-optic modulators, intensity modulation, predistortion, silicon photonics.

I. INTRODUCTION

THE growth of media streaming and cloud storage is driving data centre expansion, placing demands of high speed and low power on the next generation of short-reach optical transceivers. Recent interconnect standards have targeted the 100G and 400G data rates through multi-fibre or multi-wavelength channels operating at 25 Gbps and 50 Gbps [1], [2]. Of the photonic technologies available to meet these requirements, silicon photonics (SiP) notably offers the potential for high-density, low-cost components using mature CMOS foundries. Typical SiP modulators use depletion-mode, plasma-dispersion-effect phase shifters combined in interferometric structures such as micro-ring modulators (MRMs) or Mach-Zehnder modulators (MZMs).

MZMs can be designed to be much more insensitive to temperature variations than MRMs, making them attractive for data centre applications. To keep the half-wave voltage (V_π)

of SiP MZMs low, the phase shifter length must be relatively large, typically in the order of a few millimetres. Keeping the electro-optic bandwidth high at these lengths means MZM electrodes are normally designed as transmission lines, minimising travelling wave effects such as RF attenuation, RF reflections from impedance mismatch, and the velocity mismatch between the optical and electrical waves in the device. These travelling-wave MZMs (TW-MZMs) require a terminating resistor at the end of the electrodes to achieve impedance matching, typically of resistance $R = 50 \Omega$ (although certain designs benefit from freeing their source and/or termination impedance from a 50 Ω environment [3]–[6]). Recent demonstrations of such devices have reached bit rates of over 100 Gb/s using advanced modulation formats and digital signal processing (DSP) [7]–[10]. However, there are two drawbacks to having this terminating resistor. Firstly, the energy consumption of the MZM will be dominated by the resistor in each electrode according to $V_{pp}^2/2R$, where V_{pp} is the peak-to-peak voltage across the phase shifter, and assuming the modulation is a simple, two-level, non-return-to-zero (NRZ) format. Secondly, the voltage available from the output stage of the driver amplifier is now dropped across two circuit elements in series: the internal matching resistor of the driver and the terminating resistor of the MZM. Assuming as a first approximation that the MZM electrodes do not introduce any loss, the effective voltage across the phase shifter is halved, reducing the phase shift produced per drive voltage. In contrast, lumped MZMs, which do not have terminating resistors, act as a capacitive load, and so avoid these two issues. The effective voltage across the phase shifters is equal to the drive voltage, and their energy consumption for NRZ modulation follows the charging and discharging of a capacitor as $CV_{pp}^2/4$, where C is their capacitance, where the energy is dissipated through series resistive elements in the circuit such as the source resistance and junction resistance. However, the lack of a terminating resistor means that lumped MZMs must be kept very short (typically <1 mm) or suffer heavy bandwidth impairment due to reflections from the open-ended electrodes. As mentioned, shorter electrodes require higher V_π driving voltages and hence there would be a clear advantage in terms of drive conditions if lumped MZMs with long electrodes could be used for high-speed modulation. Several papers in recent years have focussed on designing lumped MZMs for operation with high bandwidths through optimisation of doping levels and patterns [11]–[13], silicon-organic hybrid (SOH) phase shifters [14], [15], meandered electrodes [16]–[19], passive equalisation [20], [21] and

Manuscript received November 15, 2021; revised February 21, 2022; accepted March 6, 2022. Date of publication March 10, 2022; date of current version April 5, 2022. This work was supported by Science Foundation Ireland under Grants 12/IA/1270 and 12/RC/2276. (*Corresponding author: Brian Murray.*)

Giuseppe Talli was with the Photonic Systems Group, Tyndall National Institute, T12 R5CP Cork, Ireland. He is now with the Munich Research Centre, Huawei Technologies Dusseldorf GmbH, 80992 Munich, Germany (e-mail: giuseppe.talli@huawei.com).

Brian Murray, Cleitus Antony, and Paul D. Townsend are with the Photonic Systems Group, Tyndall National Institute, T12 R5CP Cork, Ireland (e-mail: brian.murray@tyndall.ie; cleitus.antony@tyndall.ie; paul.townsend@tyndall.ie).

Digital Object Identifier 10.1109/JPHOT.2022.3158255



Fig. 1. Experiment setup for precompensated 50 Gb/s PAM4 modulation with the lumped MZM.

novel driver designs that avoid $50\ \Omega$ environments entirely by using drivers with lower output resistance [11], [13], [19], switched-capacitor [22] or open-collector [23] schemes.

We have previously shown how electrical precompensation and receiver-side equalisation can be used to enable 50 Gb/s modulation with an unterminated, lumped MZM of length 1.5 mm [24]. In this paper, we expand on this work by removing the need for receiver-side equalisation and compare through equivalent circuit modelling the trends in increased drive voltage and energy requirements when using precompensation and inductive peaking for various driver source impedances. The paper is organised as follows: firstly, in Section II, the use of electrical precompensation is shown through experiment to enable 50 Gb/s PAM4 signalling despite the low bandwidth of the MZM. Next, in Section III, an equivalent circuit model of the lumped MZM is fitted to S11 measurements to allow accurate modelling of the lumped MZM's high-frequency performance and energy requirements. In Section IV, the performance of the MZM with and without a typical $50\ \Omega$ termination is compared. Next, in Section V, various methods of boosting the lumped MZM bandwidth at the driver to reach 50 Gb/s PAM4 modulation are compared. Finally, we come to a conclusion in Section VI.

II. TRANSMISSION EXPERIMENT

The experimental setup used to characterise the RF performance of the lumped SiP MZM is shown in Fig. 1. Differential electrical drive signals were generated by a 100 GSa/s arbitrary waveform generator and digital-analog converter (AWG/DAC) with 33 GHz analog bandwidth, then amplified by a commercial 32 GHz linear RF amplifier before being applied to the MZM in dual-drive, push-pull mode via a high-speed GSGGSG probe. Identical DC bias of $-3\ \text{V}$ was added to each RF drive signal via an external bias tee. A deliberate, small asymmetry in the MZM arm lengths allowed quadrature operation to be reached by finely tuning the commercial C-band laser in a 2 nm range around 1550 nm, avoiding the need for DC or thermal bias sections in the arms. The MZM is from IMEC's iSIPP25G process and uses depletion-mode phase shifters based on rib waveguides of rib width 500 nm, rib height 220 nm, slab height 90 nm, and length 1.5 mm, with a lateral PN junction placed at the centre of the rib. Due to the high loss of the non-optimised grating couplers on this device (approx. 6 dB each), the output of the MZM was boosted by an erbium-doped fibre amplifier (EDFA). The optical signal was filtered using a 0.8 nm bandpass filter (BPF) to remove the out-of-band EDFA noise before being detected by a commercial 40 GHz photoreceiver (Rx). The resulting electrical signal was sampled by a 33 GHz, 100 GSa/s real-time oscilloscope (RTO). Before connecting to the MZM, the RF components and cables (excluding the RF probe) had

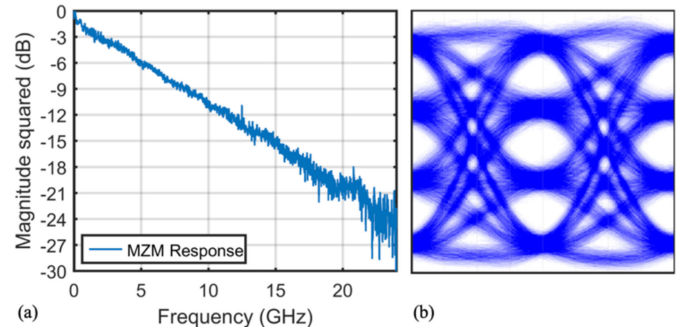


Fig. 2. (a) Measured frequency response of the lumped MZM at a bias of $-3\ \text{V}$. (b) Measured 50 Gb/s PAM4 eye diagram from the lumped MZM using a matched electrical precompensation filter in the AWG/DAC to obtain a 25 GHz fourth-order Bessel response. The outer ER is 3 dB.

been connected directly to the RTO. Their frequency roll-off was measured using flat-spectrum, Schroeder-phased harmonic signals (SPHS) [25], and then compensated out of the system for subsequent high-speed measurements by applying suitable zero-forcing filters at the AWG/DAC. The frequency response of the MZM was also measured using SPHS and is plotted in Fig. 2(a) at a bias voltage of $-3\ \text{V}$, showing a steep roll-off due to its relatively long length for a lumped design.

To measure the ER of the MZM's output when no precompensation is applied, a lower-speed 5 Gbaud NRZ signal was transmitted. The ER was calculated from the resulting optical signal's rails and found to be 9 dB for a peak-to-peak drive voltage of 3 V in each arm. A 50 Gb/s PAM4 pattern was formed from two PRBS15 patterns, with the most-significant bit (MSB) and least-significant bit (LSB) shifted by half a pattern, and then transmitted using the AWG. Due to the very low bandwidth of the MZM, the measured eye diagram was completely closed. Using this sampled signal, a 41-tap, fractionally spaced (two taps per symbol) FIR filter was generated offline in Matlab to compensate for the frequency response impairments in the system. This relatively high number of taps was used to fully cover the width of the MZM's impulse response as a general proof of concept that low-bandwidth lumped MZMs can be precompensated to work at high bit rates, taking advantage of the AWG's ability to apply as many taps as needed to maximally compensate for the MZM's performance. To reduce the energy overhead in the predistortion, a lower number of taps would be used in a driver with a feed-forward equaliser (FFE) circuit, at which point optimisation of the number of taps would be necessary for a particular device. The matched filter was adapted using a least-mean square (LMS) algorithm and designed to produce an ideal fourth-order Bessel response with a 3 dB bandwidth equal to the symbol rate when convolved with the frequency content of the transmitted signal. The filter was then subsequently applied at the AWG/DAC as an electrical precompensation filter to produce an ideal 50 Gb/s PAM4 waveform. The peak-to-peak drive voltage into each arm of the MZM was 3 V. The resulting measured eye diagram is shown in Fig. 2(b) as obtained from the RTO. The optical input power to the receiver is $-5\ \text{dBm}$. The signal exhibits clear eye opening, although the precompensation causes the outer ER to reduce to 3 dB.

III. LUMPED MZM MODEL

A. Simulation Testbed

Although 50 Gb/s PAM4 modulation is achievable with the lumped MZM in this setup, the low bandwidth of the modulator means that there is a significant drop in the ER of the signal after precompensation. To investigate more general lumped MZMs than this device and have the freedom to sweep device parameters such as the driver source impedance, it is convenient to use a simulation testbed. This also allows investigation of the energy consumption trends of the modulator by accurately knowing the time-domain voltage and current waveforms in the device. This section describes the model used. The code for the testbed was internally developed and written entirely in Matlab. The pattern generation and analysis used is identical to those in the experimental setup of Fig. 1; the simulation testbed simply replaces the DAC, laser source, MZM, and photodiode with suitable models. The electrical bandwidths of the DAC and photoreceiver are modelled via fourth-order Bessel responses with a 3 dB point equal to the symbol rate. The modelled photodiode has a noise-equivalent power (NEP) of 24 pW/ $\sqrt{\text{Hz}}$ and an ideal responsivity of 1.

B. SiP MZM Structure

To model the SiP MZM structure, the depletion-mode phase shifters in each arm are first modelled after the method described by [26]. The phase shifters are based on rib waveguides of rib width 500 nm, rib height 220 nm, and slab height 90 nm, with a length (L) of 1.5 mm. Optical mode simulations with *FEMSim* software give an effective index (n_{eff}) of 2.57 and a group index (n_g) of 3.84. A lateral doping profile is used, with the junction placed in the centre of the waveguide. The choice of p-type and n-type doping levels then determines the free carrier absorption α and effective refractive index n_{eff} of the phase shifter, providing them from the model as a function of applied voltage. The electric field of the outgoing light E_{out} from the MZM is then related to the that of the incoming light E_{in} through $E_{\text{out}} = E_1 + E_2$, where

$$E_1(t) = \frac{1}{2} E_{\text{in}} A_1(t) \exp(i\phi_o + i\phi_1(t)) \quad (1)$$

$$E_2(t) = \frac{1}{2} E_{\text{in}} A_2(t) \exp(i\phi_2) \quad (2)$$

are the respective electric fields in each arm of the MZM. The terms $A_1(t) = \sqrt{\exp(-\alpha(V_1(t))L)}$ and $A_2(t) = \sqrt{\exp(-\alpha(V_2(t))L)}$ account for the propagation loss of the electrical field in the waveguides, which is here due to the voltage-dependent free carrier absorption. The terms $\phi_1(t) = (2\pi/\lambda)\Delta n_{\text{eff}}(V_1(t))L$ and $\phi_2(t) = (2\pi/\lambda)\Delta n_{\text{eff}}(V_2(t))L$ represent the phase shift from the voltage-dependent refractive index change. The phase term ϕ_o is a phase offset introduced by length asymmetry to allow tuning of the phase without adjusting electrical bias. It is set to $\pi/2$ to operate the MZM at its quadrature point.




Fig. 3. (a) Equivalent circuit model of the driver and phase shifter for the lumped MZM, including parasitics. (b) Thévenin equivalent circuit with respect to the junction capacitance C_j .

C. S11 Fitting

To accurately model the bandwidth and energy consumption of a lumped MZM, a suitable equivalent circuit model must be obtained. Typically, some variety of LRC circuit with additional parasitics is used to model lumped modulators of this type [13], [27]. Fig. 3(a) shows such a circuit, based on a model provided by the IMEC foundry for this lumped MZM which was produced in their iSIPP25G process. Here, R_s is the source resistance of the driver (typically 50 Ω), C_m is the parasitic capacitance of the RF bond pads, R_p is the resistance of the bulk silicon, C_p is the capacitance through the oxide layer, and L_{elec} is the inductance of the electrodes. The parameters R_j and C_j are the junction resistance and capacitance, respectively, which are dependent on the doping levels of the silicon phase shifter. For a given angular frequency ω , the complex impedances of these parameters can be obtained in the usual fashion: the complex impedance for a resistance R is given simply by R , while the complex impedances of a capacitance C and inductance L are given by $1/i\omega C$ and $i\omega L$, respectively. The complex impedances resulting from the circuit parameters in Fig. 3 will be respectively labelled as Z_{R_s} , Z_{C_m} , Z_{R_p} , Z_{C_p} , Z_{R_j} , and Z_{C_j} . As the voltage dropped across the junction capacitance is the crucial load voltage for obtaining modulation with the MZM, this has been labelled $V_{\text{load}}(\omega)$.

While R_j and C_j can be modelled directly by knowing the geometry and doping of the lateral PN junction [28], knowledge of the additional parasitic elements is best obtained by fitting the circuit model to the S11 parameter (reflection coefficient) of the device [27]. This is a better choice for fitting than the S21 parameter (transfer function) as it is a one-port measurement on a VNA, removing extraneous factors from the measurement such as a photoreceiver or the optical transit time of the MZM (discussed later in this section). In terms of impedances, the S11 parameter can be expressed as

$$\Gamma = \frac{Z_{\text{mod}} - Z_{\text{source}}}{Z_{\text{mod}} + Z_{\text{source}}}, \quad (3)$$

where $Z_{\text{mod}} = Z_{C_m} \parallel (Z_{C_p} + Z_{R_p}) \parallel (Z_{R_j} + Z_{C_j} + Z_{L_{\text{elec}}})$ and $Z_{\text{source}} = Z_{R_s}$ are the complex impedances of the MZM chip and voltage source, respectively. The S11 parameter of the lumped MZM is measured at various bias voltages over frequencies from 10 MHz to 40 GHz by connecting each arm in turn to a VNA using an RF cable and a high-speed GS probe. Appropriate short-open-load (SOL) calibration is used to de-embed the effects of the cables and probe from the measurement.

TABLE I
CIRCUIT VALUES OBTAINED FROM FITTING OF THE MEASURED S11
PARAMETER FROM THE LUMPED MZM FOR VARIOUS BIAS VOLTAGES

Bias	R_j	C_j	L_{elec}	C_m	R_p	C_p
[V]	[Ω]	[fF]	[pH]	[fF]	[Ω]	[fF]
0	10.7	770	75	100	1000	200
-1	10.6	685	75	100	1000	200
-2	10.4	612	75	100	1000	200
-3	10.3	564	75	100	1000	200
-4	10.2	529	75	100	1000	200
-5	10.0	501	75	100	1000	200
-6	9.9	477	75	100	1000	200




Fig. 4. Fit of the equivalent circuit model to the (a) Real and (b) Imaginary part of the measured S11 parameter at various bias voltages. Solid lines are measured data, dotted lines are the fits.

Equation (3) is then fitted to the measured S11 data at each bias voltage using Matlab's *fmincon* function. As Γ is complex valued, its real and imaginary parts are separated during the fitting; the respective differences between the real and imaginary parts of the measured data and analytical expression are calculated individually before being combined into one mean squared error. Appropriate bounds for the fit are decided upon using sample circuit values provided by IMEC for a shorter 0.5 mm MZM in the iSIPP25G process, combined with coarse brute searches over the parameter space to identify trends in the fit. For the final fittings across the range of bias voltages, the values of L_{elec} , C_m , R_p and C_p are held fixed at sensible mean values obtained from trial fits to reflect the fact that the electrode, bond pads and bulk silicon structure should be unaffected by changes in voltage.

The resulting fitted circuit parameters are shown in Table I for the various reverse bias voltages. The measured S11 parameter and corresponding fitted curves are shown in Fig. 4 for a selection of bias voltages. Close agreement is seen between the measured and fitted data, sufficient to use the equivalent circuit model to identify key trends in behaviour.

D. Circuit Transfer Function

With an accurate equivalent circuit model determined for the lumped MZM, the load voltage V_{load} that falls across the junction capacitance C_j at various frequencies can now be determined. The transfer function that determines this voltage with respect to the source voltage is the S21 parameter. Due to the multiple circuit elements present, it is convenient to use the Thévenin equivalent of the circuit to simplify the treatment. By doing

so, the voltage source and various complex impedances in the model for the lumped MZM are replaced with an equivalent combination of a single voltage source V_{TH} in series with a single impedance Z_{TH} as shown in Fig. 3(b). The Thévenin voltage is given by

$$V_{TH}(\omega) = V_s(\omega) \frac{Z_1}{Z_2} \frac{Z_{Rt}}{Z_{Lelec} + Z_{Rt}}, \quad (4)$$

where, for ease of notation, we have set $Z_1 = Z_{Cm} \parallel (Z_{Rp} + Z_{Cp} \parallel (Z_{Lelec} + Z_{Rt}))$ and $Z_2 = Z_{Rs} + Z_1$. The frequency-dependent Thévenin impedance is given by

$$Z_{TH} = ((Z_{Rs} \parallel Z_{Cm} \parallel (Z_{Rp} + Z_{Cp}) + Z_{Lelec}) \parallel Z_{Rt}) + Z_{Rj}. \quad (5)$$

The load voltage across the junction capacitance C_j is then given by

$$V_{load}(\omega) = V_{TH}(\omega) \frac{Z_{Cj}}{Z_{TH}(\omega) + Z_{Cj}}. \quad (6)$$

The crucial transfer function of the circuit that determines the MZM frequency response is then

$$H_{vload}(\omega) = \frac{V_{load}(\omega)}{V_s(\omega)}. \quad (7)$$

The current drawn is given by dividing the source voltage by the net impedance of the source and the modulator:

$$I(\omega) = \frac{V_s(\omega)}{Z_s + (Z_1 \parallel Z_2)}. \quad (8)$$

For the PRBS waveforms used later in this paper, the power can then be calculated in the time domain by taking a dot product with $P(t) = \frac{1}{2} \text{Re}\{V_s(t) \cdot I^*(t)\}$, where I^* represents the complex conjugate of the current.

E. Effect of Optical Transit Time on Frequency Response

Once the voltage transferred to the capacitor C_j has been determined from (6), the final frequency-dependent effect that must be considered at high frequencies is the finite transit time of the light from one end of the phase shifter to the other. For electrodes below the lumped element cut-off length, the voltage present in an MZM arm at a given instant in time can be assumed to be identical regardless of the position on the electrode. Optical phase shift occurs as the light propagating in the MZM arm sees a change in effective index due to applied voltage. This effective index change is then a function of voltage, which is in turn a function of time, i.e., $\Delta n_{eff}(V(t))$. The transit time of the light through the phase shifter can be expressed as $\tau_g = L/v_g$, where v_g is the group velocity of the light, given by $v_g = c/n_g$. Here, c is the speed of light in vacuum and n_g is the optical group index of the silicon waveguide. Crucially, if τ_g is significant with respect to the time variation of the voltage, the instantaneous voltage present along the electrode at the time the light enters the phase shifter will have changed in value by the time light exits the phase shifter, weakening the overall modulation effect. This effect is similar to the well-known velocity mismatch between optical and electrical waves in TW-MZMs [29], but is treated differently due to the inherent lumped-element assumption that the RF signal is uniform over the electrode at any point in time. As derived in [30], the effect the optical transit time has on the

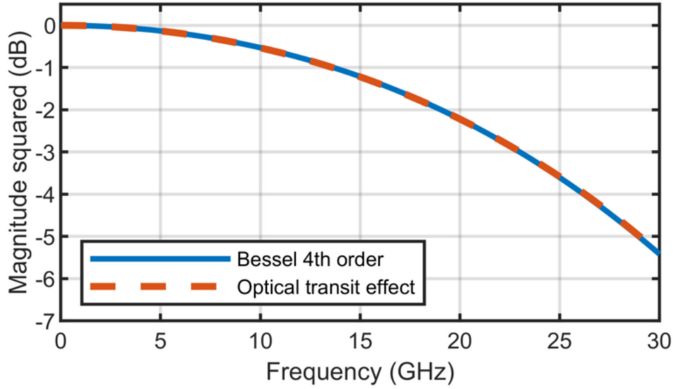


Fig. 5. Similarity of the optical transit time effect and a fourth-order Bessel response.

frequency response can be expressed as

$$H_{\text{op}}(\omega) = \frac{1 - \exp(-i\omega\tau_g)}{i\omega}, \quad (9)$$

the magnitude of which simplifies nicely to a sinc function: $|H_{\text{op}}(\omega)| = \text{sinc}(\tau_g\omega/2)$. From this, we can obtain a 3 dB bandwidth for a given MZM length due to optical transit time as $f_{\tau} \cong 0.44/\tau_g = 0.44c/(Ln_g)$. For the 1.5 mm lumped MZM studied in this work, the corresponding f_{τ} of 22.9 GHz should be investigated for the 25 Gbaud PAM4 signals in question. As a judge of its potential impact, Fig. 5 plots the magnitude squared of the frequency response resulting from (9). Notably, the roll-off is virtually identical over this range of frequencies to that of a fourth-order Bessel filter with the same 3 dB bandwidth. The group delays of both responses are flat (not plotted). Given that the 3 dB bandwidth in question is 91.6% of the 25 Gbaud symbol rate, such a roll-off can be expected to have minimal impact on the signal quality. The effect of the optical transit time on system performance will be investigated further in Section IV.

Finally, it should be noted that this bandwidth description does not include the possible onset of travelling-wave effects (RF attenuation, RF reflections, and velocity mismatch of optical and RF waves) over the length of this 1.5 mm MZM for which the lumped model cannot entirely account. Modelling such effects necessitates a longer simulation flow and knowledge of the electrode characteristics that are not immediately accessible [3]. A typical rule of thumb for the boundary between travelling-wave and lumped operation of a device is to require the electrode length to be acceptably small relative to the wavelength of the RF signal passing through it, such as within a ratio of a tenth [30] or, less stringently, a quarter [22], [31]. A suitable expression for this is $l_{\text{LE}} = (\lambda_{\text{RF0}}/n_{\text{RF}}) \cdot \Delta$, where λ_{RF0} is the free-space RF signal wavelength (given by the ratio of the speed of light in vacuum to the signal frequency), n_{RF} is the RF refractive index of the electrode and Δ is the arbitrary constant that determines the acceptable length ratio. The relevant length for determining lumped operation can be increased by making the RF contact point at the middle of the electrodes halfway along the device length [16], i.e., the signal travels 0.75 mm in each direction along a 1.5 mm MZM. Assuming $n_{\text{RF}} \approx 4$ [17], the maximum length for lumped operation in this case is then 10.5 GHz for

$\Delta = 1/10$ or 26.1 GHz for $\Delta = 1/4$. However, it should be recognised from this discussion that the boundary for a lumped-element regime is relatively loosely defined and is not a hard cut-off; the bandwidth is gradually penalised as the device is lengthened. If needed, to ensure lumped electrical operation in the upper region of frequencies, the MZM could be implemented with a meandered design [16]–[19]. This has been shown to keep the given RF distances electrically short (around 0.25 mm) while maintaining relatively long effective MZM phase shifter lengths (around 1.5 mm). As such, the analysis carried out using the lumped frequency response model in the rest of this paper can be seen as a way of identifying key trends in performance for various driver parameters while using a simple, convenient model.

IV. EFFECT OF REMOVING THE TERMINATION RESISTOR

With a model determined for the frequency response of the lumped MZM, various drive circuit conditions can now be analysed. The analysis will start by examining the effect of a having or not having a 50 Ω termination resistor on a lumped MZM. From here on, the terms *terminated* and *unterminated* will be used respectively for these conditions. As the circuit parameters derived from the S11 fitting in Table I showed a high junction capacitance (attributed to high doping levels), an MZM with lower-doped PN junctions ($5 \times 10^{17} \text{ cm}^{-3}$ for both p and n sides) will be modelled in the following simulations, with the junction capacitance C_j adjusted accordingly. The fitted junction resistance R_j was higher than values of approximately 6–8 Ω obtained using modelling following the method in [28] over doping levels of $5 \times 10^{17} \text{ cm}^{-3}$ to 10^{18} cm^{-3} , which is likely attributable to extra parasitic effects. Its value was therefore left unchanged from the fit. The other elements of the circuit also kept fixed. Hence, the final circuit parameters serve as a typical, representative example of a more moderately doped SiP depletion-mode phase shifter with higher bandwidth than our measured device, while using the additional knowledge of the parasitic elements gained from the S11 fitting. The values used are therefore as follows: $V_{\text{bias}} = -1 \text{ V}$, $R_j = 10.6 \Omega$, $L_{\text{elec}} = 75 \text{ pH}$, $C_j = 350 \text{ fF}$, $C_m = 100 \text{ fF}$, $R_p = 1000 \Omega$, and $C_p = 200 \text{ fF}$. The value of the termination resistor, when used, is $R_t = 50 \Omega$. Otherwise, it is set to a very large impedance to remove its effect from the circuit. A low peak-to-peak RF drive voltage of 1 V in each arm is used here.

Removing the termination resistor from the MZM has a noticeable effect on the magnitude and group delay of the frequency response, as shown in Fig. 6(a) and (b), respectively. The most obvious change is an increase of 6 dB in the power of the modulation at low frequencies, which corresponds to twice the effective voltage V_{load} falling across the junction capacitance. Essentially, the MZM presents an infinite impedance at low frequencies as its behaviour is dominated by the capacitance C_j of the PN junction. It acts similarly to a simple series RC circuit, which results in a high reflection coefficient. The energy from the driver is only consumed in the resistive parts of the equivalent circuit. This is beneficial in reducing the required modulation voltage for a given ER. However, similarly to our results in [24], the 3 dB bandwidth decreases significantly without the




Fig. 6. Comparison of the circuit frequency response in terms of (a) Magnitude squared and (b) Group delay for the LRC model, with and without the 50Ω termination. Also shown is a 25 GHz fourth-order Bessel response, ideal for 25 Gbaud NRZ and PAM4 signals.

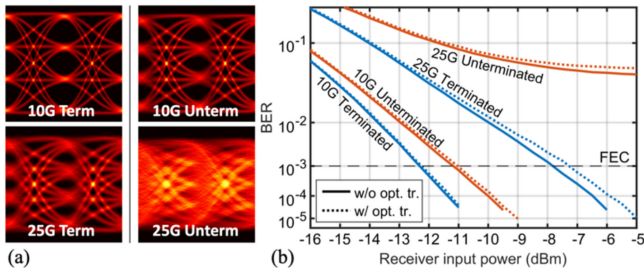


Fig. 7. (a) Simulated PAM4 eye diagrams at 10 Gbaud and 25 Gbaud for the terminated and unterminated MZM. (b) Corresponding BER curves, with/without effects of optical transit time.

termination, dropping from 11.4 GHz to 6.2 GHz. The group delay is also degraded, with the unterminated case reaching a maximum group delay of 33 ps as opposed to the 17 ps of the terminated case.

To get a qualitative sense of the effect of these frequency responses on signals, simulated PAM4 eye diagrams for the terminated and unterminated MZMs in a back-to-back link are shown in Fig. 7(a) for 10 Gbaud and 25 Gbaud. For a fair comparison, the set source voltage is adjusted to ensure that the same V_{load} falls across the junction capacitance in both the terminated and unterminated cases, i.e., the peak-to-peak source voltage from the driver is 1 V in each arm for the unterminated MZM and 2 V in each arm for the terminated MZM. Although the 10 Gbaud eyes are relatively similar, there is a marked difference in eye quality between the 25 Gbaud eyes, with the unterminated MZM showing complete eye closure due to its ~ 6 GHz transmitter bandwidth. This comparison is shown quantitatively in Fig. 7(b), where the corresponding BER curves are shown against optical power at the receiver. While 10 Gbaud PAM4 is not unduly stressed by the absence of the termination, 25 Gbaud PAM4 is less forgiving, with the unterminated MZM flooring at error rates well above a typical forward-error correction (FEC) threshold of 10^{-3} . Also shown on this graph are the equivalent BER curves obtained when the effect of the optical transit time on frequency response (9) is not included in the simulation. The resulting differences between the corresponding curves are very small for all cases due to the relative high bandwidth of this effect ($f_{\tau} = 22.9$ GHz), and so this effect will be omitted in future simulations in this paper. This has the advantage of simplifying

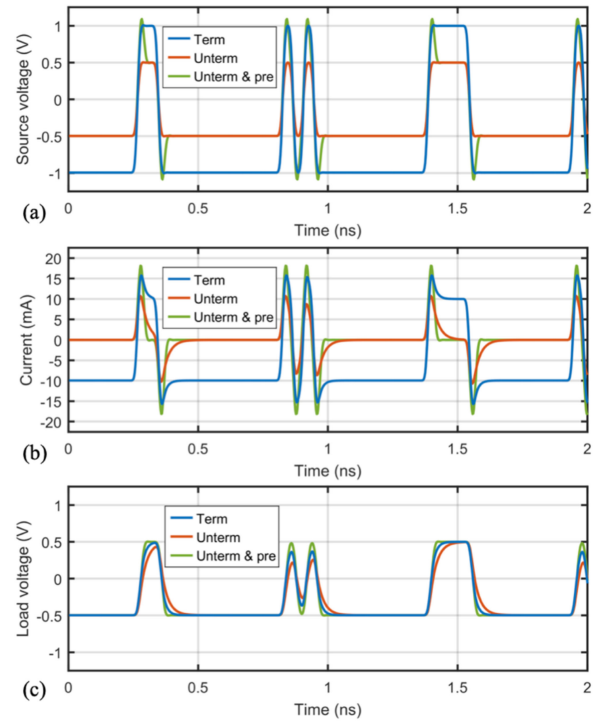


Fig. 8. Comparison of the (a) Source voltage, (b) Drawn current and (c) Load voltage waveforms for 25 Gbaud NRZ with the terminated, unterminated and precompensated unterminated MZM.

the analysis and discussion of the MZM to deal only with the equivalent circuit model.

V. IMPROVING THE LUMPED MZM BANDWIDTH THROUGH DRIVER CONSIDERATIONS

A. Method 1: Use Predistortion

As shown experimentally in Section II, a bandwidth-limited system can be precompensated at the transmitter to produce an ideal signal at the receiver, provided the frequency response of the system is known and sufficient drive voltage is available. In the following sections, it will be assumed that the driver can increase its peak-to-peak source voltage as necessary to boost the high frequencies of the signal. The resulting increase in energy dissipation and drive voltage will be noted. The magnitude and group delay of the frequency response of the modelled terminated MZM and unterminated MZM were shown in Fig. 6, where the 25 GHz Bessel response also shown corresponds to that of the unterminated MZM with precompensation applied. Although such precompensation produces an ideal PAM4 eye, it does require additional drive voltage and, hence, energy consumption, if the ER is to be kept the same after precompensation.

The reason for the increase in energy consumption can be seen in a familiar way in the time domain with PRBS waveforms; Fig. 8 shows the source voltage, drawn current and load voltage waveforms for a 25 Gbaud signal for the three cases. For clarity, NRZ is shown rather than PAM4. As could be expected, the terminated MZM requires a 2 V peak-to-peak source voltage to produce a 1 V peak-to-peak load voltage, and draws current




Fig. 9. Effect of lowering source resistance on the frequency response of the lumped MZM in terms of (a) Magnitude squared and (b) Group delay for a low electrode inductance. (c) and (d) Correspond respectively to a higher electrode inductance.

throughout each symbol duration, while the unterminated MZM needs half the source voltage and draws current only at symbol transitions. Most noticeably, the precompensated-terminated MZM still draws current only at the symbol transitions, but the required peak-to-peak source voltage has increased to 2.1 V, similar to that of the terminated MZM. The spike of drive voltage present at the start of each symbol is typical of a waveform that is precompensated to account for low link bandwidth. It counteracts the slow rise and fall times that are present in the time-domain, giving a “net” ideal load waveform, as shown in Fig. 8(c). For 25 Gbaud PAM4 signals, the terminated MZM consumes 0.42 pJ/bit, the unterminated MZM consumes 0.07 pJ/bit, and the precompensated unterminated MZM consumes 0.19 pJ/bit. Although the precompensated-terminated MZM still uses less energy per bit than the terminated MZM, the similar, doubled source voltage required for each is not ideal. Hence, the next section will examine another method of dealing with the low bandwidth of lumped MZMs.

B. Method 2: Reduce Source Resistance

Although precompensation solves the problem of the lumped MZM’s low bandwidth, it is not an ideal solution in this case as it causes a doubling of the required peak-to-peak source voltage. Alternatively, the bandwidth of a lumped MZM can be extended by reducing the source resistance R_s from the typical 50 Ω to a much smaller value to achieve impedance matching between source and modulator [19]. The resulting effect on the lumped MZM’s frequency response is shown in Fig. 9(a) and (b). As R_s reduces, the magnitude rolls off more gently and the group delay decreases towards zero and flattens out. These two effects are conducive to better eye quality and BER. Towards very low values of R_s (5 Ω , highlighted in red), the magnitude of the frequency response starts to rise above zero. This is caused by the inductance of the electrode (75 pH) taking

effect and causing a small amount of inductive peaking. This small magnitude of this peaking, which shows only at very low source resistances indicates that the lumped MZM essentially follows RC behaviour. This is to be expected, given that the only non-resistive or capacitive element in the circuit is the inductance of the electrode L_{elec} . The fitted value for this (75 pH) can be considered to be quite low, given that such an exercise for a similar lumped MZM in another work [27] gave a much higher value of 375 pH.

A potential reason for this difference in inductance is the geometry of MZM electrodes; treating them as two parallel busbars, or long rectangular cuboids, the effective inductance associated with each is approximated by [32]

$$L = \frac{\mu_0 l}{2\pi} \left[\ln \left(\frac{d}{w+t} \right) + \frac{3}{2} \right], \quad w, d \ll l, \quad (10)$$

where l is the length of the electrodes, d is the separation of the electrodes as measured between two corresponding points on each, w is the electrode width and t is the electrode thickness. This equation holds predominantly for low frequencies but can still be used to identify trends in how electrode inductance changes with geometry. Clearly, a larger inductance can be arrived upon for the same length MZM if greater separation occurs between the signal and ground electrodes, or if the electrodes are made thinner in height or narrower. The exact trade-offs in the choice of MZM electrode geometry are beyond the scope of this paper and will not be discussed further. However, given that MZMs can have larger values of electrode inductance than our fitted value, the resulting effect on the bandwidth and system performance is worth analysing. In addition, regardless of the inductance of the electrodes themselves, if wire bonds (as opposed to bump bonds) are used to electrically couple a driver to the MZM, significant inductance can be added, which can be used to inductively peak the frequency response and extend bandwidth [33]–[35]. For example, using the rule of thumb of 1 nH/mm for a typical wire bond of diameter 25 μm [36], the larger inductance value of 375 pH in [27] can be reached from our fitted value of 75 pH with a short wire bond of length 300 μm . As there are many possible values of L_{elec} depending on design, the following analysis will use the two previously mentioned inductance values (75 pH and 375 pH) as representative examples of *low-inductance* and *high-inductance* cases, respectively.

Fig. 9(c) and (d) show the frequency response of the circuit in terms of magnitude and group delay when the inductance has been increased to 375 pH. The effect of the inductance is now far stronger as the source resistance is lowered, with noticeable inductive peaking exhibited below 25 Ω (marked in red). The resonant frequency that determines the location of this peak occurs near $\sqrt{1/(2\pi LC)}$, where C is the net value of the capacitances of the circuit and L is the inductance. Taking $L = L_{elec}$ and summing the parallel capacitances to get $C = C_j + C_m$ (the other capacitance C_p can be ignored due its very large corresponding series resistance R_s), the calculated resonant frequency is then 12.25 GHz, closely matching what is seen in the figures.




Fig. 10. (a) Comparison of high-inductance and low-inductance lumped MZM 3 dB bandwidth as source resistance is lowered. (b) Corresponding energy dissipation in the circuit. (c) Rx sensitivity for 25 Gbaud NRZ and PAM4 from the lumped MZM against source resistance.

Such peaking naturally improves the MZM's 3 dB bandwidth, as demonstrated in Fig. 10(a), where the bandwidths of the low-inductance and high-inductance cases are compared. The high-inductance MZM gives greater 3 dB bandwidth than the low-inductance MZM for source resistances from 50 Ω to 17.5 Ω, after which inductive peaking causes the bandwidth to drop off quickly. However, there is a price to pay due to the extra inductance since it causes greater energy consumption in the circuit, as shown in Fig. 10(b) where the energy consumed per bit is shown against source resistance for 25 Gbaud PAM4 and NRZ. The increase is small (7 fJ/bit for PAM4) when the source resistance is a standard 50 Ω, but grows more significant as the source resistance is decreased, being 17 fJ/bit more energy intensive when the source resistance is 25 Ω, and requiring 57 fJ/bit more for a 5 Ω source resistance. In the figure, for the same symbol rate (25 Gbaud) and voltage swing, PAM4 appears ~ 3.5 times more power efficient than NRZ. This should be viewed simply in terms of the energy dissipated in the equivalent circuit model; it does not factor in any additional energy consumption from more complex PAM4 driver circuitry. It is caused by the RMS values of voltage and current being lower by a factor of 0.75 in the PAM4 signal due to the intermediate symbol levels being of smaller amplitude than the outer symbol levels. With one in every two symbols on average in a PAM4 waveform being an intermediate level, the energy consumption is then reduced by a factor of $1/(0.75 \cdot 0.75 \cdot 2) \approx 3.56$.

Although a viewing of Fig. 10(a) shows that the maximum 3 dB bandwidth occurs when R_s and L_{elec} are lowest, this is not guaranteed to be most suitable design choice for the 25 Gbaud PAM4 symbol rates of interest. There are practical challenges in creating such a circuit; every electrical trace and transistor in the driver adds to the net source impedance, adding difficulty to the circuit design, while the packaging process can easily contribute to inductance. Hence, the effect of lowered source resistance on the required receiver power for a given bit rate should be investigated to evaluate overall system performance and potentially find an optimum value for R_s .

Fig. 10(c) compares the 25 Gbaud PAM4 system performance of the low-inductance MZM and the high-inductance MZM, showing against source resistance the required power at the receiver for a BER of 10^{-3} . The 25 Gbaud NRZ case is also shown for reference. For the low-inductance case, the required power decreases monotonically with lowered source resistance, reaching a minimum when R_s is 5 Ω, i.e., when the 3 dB bandwidth is maximised. Near the minimum point, there is a 5 dB power penalty between the NRZ and PAM4 signals, which is to be expected due to the eye openings decreasing by a factor of three for the four-level signal. For the high-inductance case, the optimal source resistance is 25 Ω rather than the lowest possible value, as could be expected from a reading of Fig. 9(c) and (d). Generally, the optimal overall frequency response for a transmission system is one that remains flat for longest (or rolls off slightly) towards the symbol rate. For values of R_s above 25 Ω, the dominant system impairment is low bandwidth, while below this, over peaking from the inductance starts to take effect. As indicated in the figure, the regions to the left and right of this optimal source resistance can be classified respectively as the *inductive peaking regime* and the *low-bandwidth regime*. There is minimal difference (less than 0.25 dB) penalty in the lowest required Rx power between the low-inductance and high-inductance case; this is important, as it shows that the requirements on lowering source resistance can be significantly relaxed for a lumped MZM that has an appreciable amount of inductance. Furthermore, although the trough-like shape of the Rx power curve for PAM4 in Fig. 10(c) suggests that best performance requires exactly 25 Ω source resistance, the required Rx power degrades by only 1 dB over the range of 18 Ω to 34 Ω. Again, this relaxes the design constraints for an implementation of such a driver. As could be expected, 25 Gbaud NRZ is less sensitive to variations in source resistance than its 25 Gbaud PAM4 counterpart due to its simpler eye shape; the region of over-peaking from inductance does not have noticeable effect until approximately 10 Ω source resistance. In fact, the optimal source resistance for NRZ (20 Ω) is marginally lower than that of PAM4, as slight over-peaking benefits the simpler two-level signal; there is no outer eye impacted by the overshoot in an NRZ signal as there would be in PAM4. This emphasises the need to carefully tune source resistance and inductance against one another if attempting to inductively peak a PAM4 modulator in comparison with the more relaxed needs of an NRZ modulator.

C. Method 3: Reduce Source Resistance & Use Predistortion

With two distinct methods proposed for improving the bandwidth of lumped MZMs, a question regarding a third naturally arises: what happens if both precompensation and a lowered source resistance are implemented?

Fig. 11(a) shows the simulated required Rx power for the same high-inductance MZM, now with and without precompensation to optimise the signal. The precompensation applied is a matched filter that is again designed to give a fourth-order Bessel response with a 3 dB bandwidth of 25 GHz. As could be expected, precompensation results in the best possible receiver sensitivity regardless of the source resistance; provided sufficient source

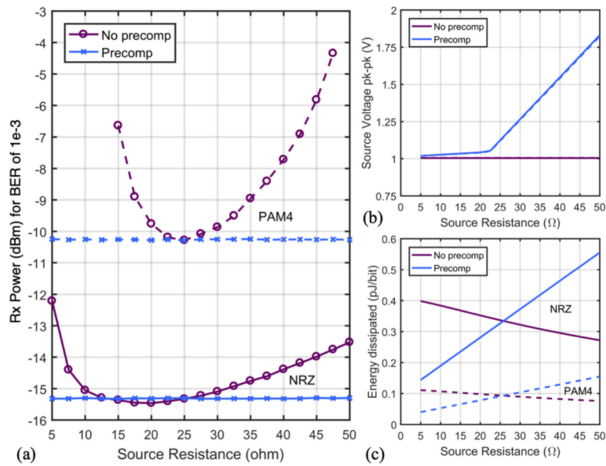


Fig. 11. (a) Rx sensitivity for 25 Gbaud NRZ and PAM4 from the lumped MZM without and with precompensation as source resistance is reduced. (b) Peak-to-peak drive voltage in each case. (c) Corresponding energy dissipated in the circuit.

voltage is available from the driver and the frequency response is well known, the linear impairments of the system can be compensated perfectly. Notably, the required Rx power of the non-precompensated case at optimal source impedance matches that of the precompensated case. Certainly, any real implantation of a MZM and driver will have features in the frequency response that a simple reduction in source impedance and inductive peaking cannot perfectly compensate. However, this result does point towards the potential for a well-designed driver without precompensation to at least approach the performance of a precompensated scenario.

Although precompensation resolves the issue of bandwidth regardless of source resistance, the required source voltage and resulting energy dissipated are heavily affected by the type of bandwidth impairment in question; compensating for the effects of low bandwidth is not the same as compensating for the effects of over peaking. Fig. 11(b) shows the required source voltage for precompensating the signal, plotted against source resistance. The non-precompensated case is shown for reference. The peak-to-peak source voltage is highest for the 50 Ω source resistance, with 1.8 V required in each arm. This is slightly less than the 2 V that was needed to precompensate the low-inductance MZM in Fig. 8(a) as the bandwidth is slightly better with higher inductance at a standard 50 Ω source impedance. The required source voltage decreases with the source resistance, until a knee is reached when $R_s = 22.5 \Omega$, very close to the source resistance for minimum required Rx power. Beyond this point, the voltage needed remains relatively constant at approximately 1 V. Essentially, as the chosen source resistance strays into the low-bandwidth regime, the required drive voltage when using precompensation increases. In contrast, no extra drive voltage is needed in the inductive-peaking regime. The resulting effect of precompensation on the energy consumption of the circuit is shown in Fig. 11(c), where the energy dissipated per bit for 25 Gbaud NRZ and PAM4 waveforms is shown against source resistance. The non-precompensated case is again shown for comparison. The precompensated case dissipates approximately

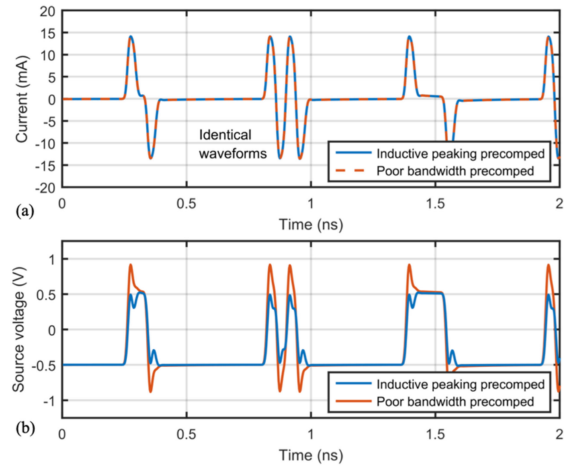


Fig. 12. (a) Current drawn for the precompensated lumped MZM when the source resistance is 50 Ω (poor bandwidth regime) and 20 Ω (inductive peaking regime). (b) the precompensated source voltages for the two cases.

twice the energy of the non-precompensated case at standard 50 Ω source resistance. However, as established earlier, the bandwidth of the non-precompensated lumped MZM is too low at this source resistance for 25 Gbaud PAM4 transmission, so its energy saving is not a significant plus. When the system performance is of both cases is equal (i.e., when $R_s = 25 \Omega$), the energy dissipated is exactly equal. Below this point, in the inductive peaking regime, the precompensated case dissipates less energy per bit than non-precompensated case.

The reason for this transition in energy performance can be seen clearly by looking at the precompensated waveforms for the drawn current and source voltage in the time domain. Fig. 12(a) shows the drawn current for two example source resistances in the low-bandwidth ($R_s = 50 \Omega$) and inductive peaking regimes ($R_s = 20 \Omega$). For ease of analysis, 25 Gbaud NRZ is used rather than PAM4. Due to the precompensation, the current flowing in the circuit is identical for both cases. However, when viewing the corresponding source voltages in Fig. 12(b), a clear difference is seen. In the poor-bandwidth regime, naturally the peak-to-peak voltage has increased, providing a spike of voltage at the symbol transitions to account for the insufficient circuit rise time. In the inductive peaking regime, however, the opposite occurs. The presence of the resonant peak in the frequency response means that certain high frequencies in the signal now have too much power, and the source voltage must be reduced just after the symbol transitions to compensate and flatten the response. The power, i.e., the dot product of the source voltage and current waveforms, is then reduced for the precompensated, inductively peaked MZM. As a result, one potential for a low-power, low-voltage design of lumped MZM driver is to reduce the source impedance as much as possible, while adding inductive peaking and precompensation.

Finally, while Fig. 11(a) showed that the performance of the lumped MZM with optimal inductive peaking matches that of the precompensated scenario in a back-to-back (B2B) PAM4 link, a comparison should also be made of their performance over optical fibre, as subtle differences in their eye quality can


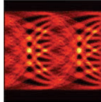
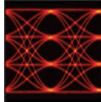
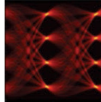
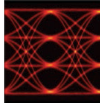


Fig. 13. Comparison of Rx sensitivity against fibre length for 25 Gbaud PAM4 with the lumped MZM in the cases of inductive peaking and precompensation. Here, $R_s = 25\Omega$.

TABLE II
COMPARISON OF KEY METRICS FOR THE DISCUSSED DRIVE OPTIONS FOR 25 GBAUD PAM4 WITH THE UNTERMINATED LUMPED MZM WITH $L_{\text{elec}} = 375 \text{ pH}$. SSMF LENGTH REFERS TO THE MAXIMUM DISTANCE ACHIEVABLE AT AN FEC BER THRESHOLD OF 10^{-3}

	50 Ω driver	Precomp & 50 Ω driver	Low impedance driver	Precomp & lower impedance driver
V_{source}	1 Vpp	2.1 Vpp	1 Vpp	1 Vpp
R_{source}	50 Ω	50 Ω	25 Ω	<25 Ω
$E_{\text{consumption}}$	70 fJ/bit	190 fJ/bit	95 fJ/bit	<95 fJ/bit
B2B eye				
ER	4 dB	4 dB	4 dB	4 dB
Max SSMF length	N/A	13 km	14 km	13 km

result in different link performance. Hence, a sweep of required power at the receiver for a BER of 10^{-3} was performed against optical fibre length for these two cases when $R_s = 25\Omega$. The simulated fibre is standard single mode fibre (SSMF) with a group velocity dispersion of 17 ps/nm-km. Fig. 13 shows that the transmission performance of both drive schemes is similar, with less than 1 dB in sensitivity penalty between the curves up to 12 km of fibre. Noticeably, the drive method using only inductive peaking slightly outperforms the electrical precompensation method over optical fibre, reaching 1 km further before flooring above the FEC threshold. We attribute this to slight over-peaking in the frequency response counteracting the effects of chromatic dispersion [37].

D. Summary of Design Options

The drive options discussed in this paper to increase the bandwidth of a lumped MZM are summarised in Table II, which compares the required source voltage and energy dissipation per bit to achieve a 25 Gbaud PAM4 signal with 4 dB ER. The listed voltage refers to the peak-to-peak source voltage needed in each arm of the MZM when operating in push-pull fashion. The listed energy per bit refers to the energy dissipated in the equivalent circuit of Fig. 3(a) with the typical circuit values used throughout the simulations in this paper. For a

50 Ω drive scenario, using an unterminated, lumped MZM and applying precompensation causes the drive voltage requirement to increase significantly, exceeding that which would be required for a terminated MZM of the same length. If the constraint of a 50 Ω environment can be relaxed while still employing precompensation, lowering the source resistance until some inductive peaking occurs can offer the best combination of eye quality, power consumption and peak-to-peak voltage. Notably, careful tuning of the source resistance and inductive peaking can result in a precompensation-free modulator with comparable BER performance to a precompensated case.

It should be emphasized that this work looks only at trends in energy consumption for various drive methodologies; the driver is modelled here as ideal voltage source in series with a simple source resistance and does not include any additional energy overheads such as the static power consumption of a driver used to generate the precompensated electrical PAM4 signals. Instead, it focusses on the increased output swing for a given precompensation, which is crucial both in terms of the energy consumption and for the choice of silicon technology used (CMOS or BiCMOS). It should also be recognized that reducing the output impedance of the driver is not a trivial task and requires some trade-offs in design. For example, reducing the output impedance of a driver using current-mode logic (CML) will lead to an increase in power consumption in the driver as more current flows across the output/source resistors [38], making it a non-ideal solution for this case. A less energy-hungry approach is to drive the MZM with the output of a CMOS inverter, which naturally has a low source impedance and benefits from a small footprint [20], [39], [40]. Unfortunately, the binary nature of inverter outputs means that electrical precompensation of the signal is not possible with such a simple drive scheme. In this case, predistortion techniques would be restricted to inductive peaking such as presented here or alternative passive RC equalisation methods [20]. If electrical precompensation is desired, a linear driver amplifier would be required; such a driver with a relatively low source impedance (30 Ω) has been demonstrated previously in the literature for an unterminated MZM [6].

Finally, while this work focusses on 50 Gb/s solutions, further bandwidth enhancements would be possible with a more optimised junction design that provides more phase shift efficiency. Indeed, achieving a lumped MZM of shorter length (and smaller footprint) while maintaining the same V_{π} would fit well with the datacentre industry's recent moves towards co-packaged optics (CPO) [41]. In particular, silicon-organic hybrid (SOH) phase shifter designs have recently shown promise for achieving large phase shifts with very short electrodes operating as unterminated devices [15], in which case predistortion of lumped MZMs could be particularly relevant towards 100 Gb/s and beyond.

VI. CONCLUSION

In this paper, drive methodologies for lumped MZMs are compared systematically in terms of peak-to-peak source voltage requirements, energy consumption, and system performance. After experimentally demonstrating 50 Gb/s PAM4 with a very low-bandwidth lumped MZM using electrical precompensation,

an accurate equivalent circuit model is used to analyze driver design tradeoffs for high-speed PAM4 signaling. The model shows that reducing the source impedance of the driver significantly improves the electrical bandwidth of the lumped MZM from 6 GHz for a 50 Ω case to 38 GHz for a 5 Ω driver. Additionally, a moderate amount of additional electrode inductance (300 pH) introduces inductive peaking, which is shown to relax the requirement of having a very low source impedance and enable the same system performance at a higher driver impedance of 25 Ω . Furthermore, careful selection of source impedance and inductive peaking can allow equivalent performance as for the case of precompensation, showing the potential for a well-designed driver to reduce the need for transmitter-side DSP. Finally, the use of precompensation in combination with a low-source-resistance, inductively peaked lumped MZM is shown to further reduce the energy dissipated in the equivalent circuit.

REFERENCES

- [1] Y. Sun and R. Lingle, "Technical feasibility of new 200 Gb/s and 400 Gb/s links for data centers," in *Proc. IEEE Opt. Interconnects Conf.*, Jun. 2018, pp. 37–38.
- [2] "Standard amendment for media access control (MAC) parameters, physical layers and management parameters for 200 Gb/s and 400 Gb/s operation," IEEE 802.3bs, 2017.
- [3] H. Yu and W. Bogaerts, "An equivalent circuit model of the traveling wave electrode for carrier-depletion-based silicon optical modulators," *J. Lightw. Technol.*, vol. 30, no. 11, pp. 1602–1609, 2012.
- [4] K. Zhu, V. Saxena, X. Wu, and W. Kuang, "Design considerations for traveling-wave modulator-based CMOS photonic transmitters," *IEEE Trans. Circuits Syst., II, Exp. Briefs*, vol. 62, no. 4, pp. 412–416, Jan. 2015.
- [5] C. Xiong, D. M. Gill, J. E. Proesel, J. S. Orcutt, W. Haensch, and W. M. J. Green, "Monolithic 56 Gb/s silicon photonic pulse-amplitude modulation transmitter," *Optica*, vol. 3, no. 10, pp. 1060–1065, 2016.
- [6] A. Zandieh, P. Schvan, and S. P. Voinescu, "Linear large-swing push-pull sige BiCMOS drivers for silicon photonics modulators," *IEEE Trans. Microw. Theory Techn.*, vol. 65, no. 12, pp. 5355–5366, Nov. 2017.
- [7] H. Li *et al.*, "A 112 Gb/s PAM4 silicon photonics transmitter with microring modulator and CMOS driver," *J. Lightw. Technol.*, vol. 38, no. 1, pp. 131–138, 2020.
- [8] M. Li, L. Wang, X. Li, X. Xiao, and S. Yu, "Silicon intensity Mach-Zehnder modulator for single lane 100 Gb/s applications," *Photon. Res.*, vol. 6, no. 2, pp. 109–116, 2018.
- [9] A. Samani *et al.*, "A low-voltage 35-GHz silicon photonic modulator-enabled 112-Gb/s transmission system," *IEEE Photon. J.*, vol. 7, no. 3, Apr. 2015, Art. no. 7901413.
- [10] D. Patel *et al.*, "Design, analysis, and transmission system performance of a 41 GHz silicon photonic modulator," *Opt. Exp.*, vol. 23, no. 11, pp. 14263–14287, 2015.
- [11] D. M. Dourado, M. L. Rocha, J. P. P. Carmo, G. B. Farias, and J. E. Ribeiro, "Modeling and trade-off analysis of a capacitive silicon Mach-Zehnder modulator for telecom applications," in *Proc. SBFoton Int. Opt. Photon. Conf. (SBFoton IOPC)*, 2018, pp. 1–5.
- [12] D. M. Dourado, G. B. de Farias, M. de M. A. C. Moreira, M. de L. Rocha, and J. P. Carmo, "Silicon modulator design using a system-oriented methodology for high-speed data center interconnect PAM-4 applications," *Opt. Commun.*, vol. 492, 2021, Art. no. 126977.
- [13] X. Li, F. Yang, F. Zhong, Q. Deng, J. Michel, and Z. Zhou, "Single-drive high-speed lumped depletion-type modulators toward 10 fJ/bit energy consumption," *Photon. Res.*, vol. 5, no. 2, pp. 134–142, Apr. 2017.
- [14] S. Wolf *et al.*, "Silicon-organic hybrid (SOH) Mach-Zehnder modulators for 100 Gbit/s on-off keying," *Sci. Rep.*, vol. 8, no. 1, pp. 1–13, 2018, Art. no. 2598.
- [15] C. Kieninger *et al.*, "Silicon-organic hybrid (SOH) Mach-Zehnder modulators for 100 GBd PAM4 signaling with sub-1 dB phase-shifter loss," *Opt. Exp.*, vol. 28, no. 17, pp. 24693–24707, 2020.
- [16] S. S. Azadeh *et al.*, "Low V_{π} silicon photonics modulators with highly linear epitaxially grown phase shifters," *Opt. Exp.*, vol. 23, no. 18, pp. 23526–23550, 2015.
- [17] S. S. Azadeh, J. Nojić, A. Moscoso-Mártir, F. Merget, and J. Witzens, "Power-efficient lumped-element meandered silicon Mach-Zehnder modulators," *Silicon Photon. XV (Int. Soc. Opt. Photon.)*, vol. 11285, 2020, Art. no. 112850C.
- [18] G. Cong *et al.*, "Ultra-compact non-travelling-wave silicon Mach-Zehnder modulator," in *Proc. 45th Eur. Conf. Opt. Commun.*, 2019, pp. 1–3.
- [19] G. Cong, Y. Maegami, M. Ohno, and K. Yamada, "Ultra-compact non-travelling-wave silicon carrier-depletion Mach-Zehnder modulators towards high channel density integration," *IEEE J. Sel. Topics Quantum Electron.*, vol. 27, no. 3, Sep. 2021, Art. no. 8200311.
- [20] Y. Sobu, S. Tanaka, and Y. Tanaka, "High-speed-operation of all-silicon lumped-electrode modulator integrated with passive equalizer," *IEICE Trans. Electron.*, vol. 103, no. 11, pp. 619–626, 2020.
- [21] Y. Sobu, S. Tanaka, Y. Tanaka, Y. Akiyama, and T. Hoshida, "High-speed-operation of compact all-silicon segmented Mach-Zehnder modulator integrated with passive RC equalizer for optical DAC transmitter," in *Proc. Opt. Fiber Commun. Conf.*, Washington, D.C., USA: OSA Technical Digest (Optical Society of America), 2020, Paper M2B.8.
- [22] S. Zhou *et al.*, "Optimization of PAM-4 transmitters based on lumped silicon photonic MZMs for high-speed short-reach optical links," *Opt. Exp.*, vol. 25, no. 4, pp. 4312–4325, 2017.
- [23] N. Wolf *et al.*, "Electro-optical co-design to minimize power consumption of a 32 GBd optical IQ-Transmitter using InP MZ-Modulators," in *Proc. IEEE Compound Semicond. Integr. Circuit Symp.*, Oct. 2015, pp. 1–4.
- [24] B. Murray, C. Antony, G. Talli, and P. D. Townsend, "50-Gbps PAM4 and electrical duobinary modulation with a lumped silicon photonic MZM," in *Proc. 45th Eur. Conf. Opt. Commun.*, 2019, pp. 1–4.
- [25] J. O. Flower and S. C. Forge, "Developments in frequency-response determination using Schroeder-phased harmonic signals," *Radio Electron. Eng.*, vol. 51, no. 5, pp. 226–232, 1981.
- [26] D. Pérez-Galacho *et al.*, "Simplified modeling and optimization of silicon modulators based on free-carrier plasma dispersion effect," *Opt. Exp.*, vol. 24, no. 23, pp. 26332–26337, 2016.
- [27] H. Yu *et al.*, "Performance tradeoff between lateral and interdigitated doping patterns for high speed carrier-depletion based silicon modulators," *Opt. Exp.*, vol. 20, no. 12, pp. 12926–12938, 2012.
- [28] M. Félix Rosa *et al.*, "Design of a carrier-depletion Mach-Zehnder modulator in 250 nm silicon-on-insulator technology," *Adv. Radio Sci.*, vol. 15, pp. 269–281, 2017.
- [29] G. Ghione, *Semiconductor Devices for High-Speed Optoelectronics*, Cambridge, U.K.: Cambridge Univ. Press, 2009.
- [30] J. Witzens, "High-speed silicon photonics modulators," *Proc. IEEE*, vol. 106, no. 12, pp. 2158–2182, Nov. 2018.
- [31] H. Johnson and M. Graham, *High-Speed Signal Propagation: Advanced Black Magic*, Prentice Hall, 2003.
- [32] S. Caniggia and F. Maradei, *Signal Integrity and Radiated Emission of High-Speed Digital Systems*, Hoboken, NJ, USA: Wiley, 2008.
- [33] W. Li *et al.*, "100 Gbit/s co-designed optical receiver with hybrid integration," *Opt. Exp.*, vol. 29, no. 10, pp. 14304–14313, 2021.
- [34] K. Zhu, V. Saxena, and X. Wu, "Modeling and optimization of the bond-wire interface in a hybrid CMOS-photonic traveling-wave MZM transmitter," in *Proc. 29th IEEE Int. Syst.-Chip Conf.*, 2016, pp. 151–156.
- [35] I. García López *et al.*, "High speed BiCMOS linear driver core for segmented InP Mach-Zehnder modulators," *Analog Integr. Circuits Signal Process.*, vol. 87, no. 2, pp. 105–115, 2016.
- [36] D. Michelon, E. Bergeret, M. Egels, and A. Di Giacomo, "Wire-bonds used as matching inductor in RF energy harvesting applications," in *Proc. 10th Conf. Ph.D. Res. Microelectronics Electron.*, 2014, pp. 1–4.
- [37] L. Breyne *et al.*, "Electro-optic frequency response shaping in high speed Mach-Zehnder modulators," in *Proc. Opt. Fiber Commun. Conf. Exhib.*, 2020, pp. 1–3.
- [38] J. Hwang *et al.*, "A 32 Gb/s, 201 mW, MZM/EAM cascode push-pull CML driver in 65 nm CMOS," *IEEE Trans. Circuits Syst., II, Exp. Briefs*, vol. 65, no. 4, pp. 436–440, Apr. 2018.
- [39] A. Michard, J.-F. Carpentier, N. Michit, P. Le Maître, P. Bénabès, and P. M. Ferreira, "A sub-pJ/Bit, low-ER Mach-Zehnder-based transmitter for chip-to-chip optical interconnects," *IEEE J. Sel. Topics Quantum Electron.*, vol. 26, no. 2, Nov. 2020, Art. no. 8301910.
- [40] M. Webster *et al.*, "An efficient MOS-capacitor based silicon modulator and CMOS drivers for optical transmitters," in *11th Int. Conf. Group IV Photon.*, 2014, pp. 1–2.
- [41] C. Minkenberg, R. Krishnaswamy, A. Zilkie, and D. Nelson, "Co-packaged datacenter optics: Opportunities and challenges," *IET Optoelectron.*, vol. 15, no. 2, pp. 77–91, 2021.