Analog Coherent Detection for Energy Efficient Intra-Data Center Links at 200 Gbps Per Wavelength

Takako Hirokawa, *Student Member, IEEE, Student Member, OSA*, Sergio Pinna, *Member, IEEE*, Navid Hosseinzadeh, *Student Member, IEEE*, Aaron Maharry, Hector Andrade, Junqian Liu, Thomas Meissner, Stephen Misak, *Student Member, IEEE*, Ghazal Movaghar, Luis A. Valenzuela, *Student Member, IEEE*, Yujie Xia, Shireesh Bhat, Fabrizio Gambini, Jonathan Klamkin, *Senior Member, IEEE, Senior Member, OSA*, Adel A. M. Saleh, *Life Fellow, IEEE, Fellow, OSA*, Larry Coldren, *Life Fellow, IEEE, Fellow, OSA*, James F. Buckwalter, *Senior Member, IEEE*, and Clint L. Schow, *Fellow, IEEE, Fellow, OSA*

***Abstract*—As datacenters continue to scale in size, energy efficiency for short reach (<2 km) links is a major factor for networks that may connect hundreds of thousands of servers. We demonstrate that links based on analog coherent detection (ACD) offer a promising path to simultaneously achieving significantly larger link budgets and improved link energy efficiency. A complete analysis is presented that considers the power consumption of all the photonic and electronic components necessary to realize an ACD link architecture based on 50 Gbaud (GBd) quadrature phase-shift keying (QPSK) signaling combined with polarization multiplexing to achieve 200 Gb/s/λ. These links utilize receivers that incorporate an optical phase-locked loop (OPLL) to frequency- and phase-lock the local oscillator (LO) laser to the incoming signal. QPSK modulation offers compelling advantages both in achievable link budget and in energy efficiency. Indeed, low-complexity electronics based on limiting amplifiers can be used as opposed to the linear front-ends, A/D converters, and digital signal processing (DSP) required for higher-order QAM or PAM formats. Our analysis indicates that links with 13 dB of unallocated budget operating at error rates of $<10^{-12}$ can be achieved and is compatible with higher error rates that require forward error correction (FEC). We present a comparison of silicon and InP platforms and evaluate both traveling-wave and segmented modulator designs, providing an illustration of the wide design space before converging on the most promising architectures that maximize energy efficiency and minimize laser power. We establish the theoretical potential to achieve picojoule-per-bit energy efficiency targets.**

***Index Terms*—Coherent detection, data center, energy efficiency.**

I. Introduction

With ever-increasing demand for cloud services, evaluating interconnect technology benefits and tradeoffs anticipates future deployments of the data center through scaling baud rates, higher order modulation formats with more bits/symbol, polarization multiplexing, and adding additional wavelength division multiplexed (WDM) channels. Current data center links rely on intensity-modulated direct detection (IMDD) schemes due to their relative simplicity and correspondingly relatively low cost and power consumption. However, scaling IMDD links to 200 Gbps/lane will require a large jump in complexity and power consumption. A recent study showed the potential of a 100 GBd PAM-4 link to operate over a 400 m link distance [1]. However, heavy equalization was required, with 71 feedforward equalizer (FFE) taps and 15 decision feedback equalizer (DFE) taps, just to achieve a pre-FEC (Forward Error Correction) bit error ratio (BER) slightly below the soft decision (SD-FEC) limit of 2×10^{-2}. With such power-hungry equalization, the required received optical power was $> +7$ dBm, likely demanding an unfeasible output power from the transmitter (TX) source laser [1]. The limited prospects for scaling IMDD links to 200 Gbps/lane and beyond have driven substantial interest in developing a new generation of energy-efficient coherent links designed specifically for intra-datacenter applications [2]–[5].

A recent paper by authors from the Alibaba Group presents a detailed comparison of several variants of IMDD (PAM4, CAP16, DMT) against digital coherent (PDM-16QAM) for 400G links, backed up with experimental results, using metrics of minimizing laser and ASIC power consumption [4]. The authors conclude that coherent links have lower laser power requirements and comparable ASIC power dissipation and digital signal processing (DSP) complexity compared to the IMDD approaches. Recent work from Google provides a comparison up to 1.6 Tb/s, analyzing in detail multiple digital coherent (16, 32,

Manuscript received July 14, 2020; revised September 21, 2020; accepted October 5, 2020. Date of publication October 12, 2020; date of current version January 15, 2021. This work was supported in part by the Advanced Research Projects Agency-Energy under Grant ARPA-E, and in part by the U.S. Department of Energy under Award DE-AR0000848. *(Corresponding author: Takako Hirokawa.)*

Takako Hirokawa, Sergio Pinna, Navid Hosseinzadeh, Aaron Maharry, Hector Andrade, Junqian Liu, Thomas Meissner, Stephen Misak, Ghazal Movaghar, Luis A. Valenzuela, Yujie Xia, Fabrizio Gambini, Jonathan Klamkin, Adel A. M. Saleh, Larry Coldren, James F. Buckwalter, and Clint L. Schow are with the Electrical and Computer Engineering Department, University of California, Santa Barbara, CA 93106-9560 USA (e-mail: takako@ucsb.edu; pinna@ece.ucsb.edu; hosseinzadeh@ucsb.edu; amaharry@ucsb.edu; handrade@ucsb.edu; junqian@ucsb.edu; thomas_meissner@ucsb.edu; smisak@ucsb.edu; ghazalmovaghar@ucsb.edu; valenzuela@ucsb.edu; yujiexia@ucsb.edu; fgambini@ucsb.edu; klamkin@ece.ucsb.edu; adelsaleh@ece.ucsb.edu; coldren@ece.ucsb.edu; buckwalter@ece.ucsb.edu; schow@ece.ucsb.edu).

Shireesh Bhat was with the University of California, Santa Barbara, CA 93106-9560 USA. He is now with Juniper Networks, Sunnyvale, CA 94089 USA (e-mail: sbhat@ece.ucsb.edu).

Color versions of one or more of the figures in this article are available online at https://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JLT.2020.3029788

64QAM) and IMDD (PAM4, 6, 8) architectures [5]. The coherent links are projected to consume somewhat more power—on the order of 10-20%—but offer substantial advantages: greater tolerance to fiber impairments, higher spectral efficiency, and a large advantage in receiver (RX) sensitivity. For modulator drive swings less than $\sim 1V_{\pi}$, the gains in RX sensitivity are found to be mostly offset by large modulator losses and the PAM links are projected to achieve larger link budgets. The coherent links operate at 2X higher total bit rates, and with higher modulator drive voltages achieve 5-9 dB more link budget than the IMDD variants [5].

Digital coherent architectures commonly used in telecom interconnects are implemented with a free-running local oscillator (LO) which requires an RX chain consisting of a linear receiver front end followed by an analog-to-digital converter (ADC) to digitize incoming data. Doing so enables the DSP to perform functions such as carrier recovery, polarization demultiplexing and channel equalization to remove fiber propagation impairments such as chromatic dispersion (CD) and polarization mode dispersion (PMD). An alternative approach to coherent detection is analog coherent detection (ACD) which utilizes a highly integrated optical phase-locked loop (OPLL) to directly lock the frequency and phase of the LO laser to an incoming wavelength channel. Chip-scale integration enables low feedback loop delay and therefore high loop bandwidth, enabling the use of more easily integrated tunable LO lasers with MHz-scale linewidth [6]–[11]. Furthermore, the OPLL approach provides for the direct demodulation of complex signals at low uncorrected bit error rates (BER), with previous proof-of-concept demonstrations achieving BER $<10^{-12}$ for BPSK modulation up to 35 Gb/s [9]. Although latency may not be especially critical for our primary target application of intra-datacenter links where the use of FEC is ubiquitous, the potential to construct FEC-free coherent links offers a substantial advantage for highly latency-sensitive applications such as high-performance computing (HPC), Another key benefit of OPLL-based coherent detection in general, and offered by our OPLL ACD architecture is inherent wavelength selectivity. When the LO is locked to an incoming wavelength channel, other channels are rejected by the RX. For example, if the system channel spacing is 200 GHz, when the LO is locked to one of the wavelength channels in the incoming optical signal, the locked signal is down-converted to the baseband while the other wavelength channels are converted to 200 GHz or higher—far above the operating bandwidth of the receiver electronics. This wavelength selectivity can be exploited to reduce crosstalk requirements for future networks that incorporate photonic routing/switching and eases channel crosstalk requirements of on-chip wavelength multiplexing/demultiplexing components.

It is widely accepted that much of the complexity of traditional coherent DSP can be removed for datacenter applications [4], [5] where O-band operation of links up to 2km present negligible fiber impairments. Consequently, the biggest power savings offered by ACD arises not through the elimination of DSP, but through the removal of linear RX frontends and ADCs. QPSK as a modulation format uniquely takes full advantage of the direct demodulation capability enabled by ACD. At the output of the 90° hybrids in an ACD receiver, the I and Q channels have been separated and low-power electronics using limiting amplifiers can be used to make a binary decision, just like in the most power efficient non-return to zero (NRZ) on-off keying (OOK) links [12]. State-of-the-art ADCs have been developed with sufficient sampling rate and effective number of bits (ENOB) for 224 Gbps DP-16QAM coherent receivers with power consumption ranging between 235 mW [13] to 702 mW [14]. A dual-polarization I-Q receiver would require four such ADCs, resulting in a total ADC power consumption between 940–2808 mW or 4.2–12.5 pJ/bit based upon the efficiencies reported in [13], [14]. Our QPSK link architecture does not require these power-hungry components and full-link energy efficiencies of less than 5 pJ/bit are feasible. The substantial power savings advantage for QPSK does not straightforwardly scale to higher order QAM formats which require multiple decision thresholds for both I and Q channels, driving the need for A/D conversion.

In this paper we present a multi-wavelength analog coherent detection (ACD) architecture utilizing a chip-scale OPLL and based on 50 GBd polarization-multiplexed QPSK (PM-QPSK) for an aggregate data rate of 200 Gbps/λ. In addition to the link-level advantages in optical budget and power efficiency offered by QPSK-based ACD, we believe it will be advantageous to scale to bit rates of 800 Gb/s and beyond by using four or more WDM lanes, each carrying 200 Gb/s, as opposed to fewer lanes at higher per-λ bit rates. The large optical loss budget enabled by ACD further opens a wider space for network architecture designs offering greater flexibility and scalability through the insertion of optical wavelength-level routing and/or circuit switching devices in the data center network. Keeping the per-λ bandwidth granularity lower expands opportunities for network architectures with substantial power savings and enhanced operational flexibility as discussed in Section IV.

II. Coherent Link Energy Efficiency Model

In this section, we present an ACD link model that supports a quantitative exploration of the design space of modulator length, drive voltage, and TX source and LO laser powers. For ease of reference, we refer to the operating baud rates as 50 GBd, but all simulations are conducted at 56 GBd to allow for coding and FEC overhead. Furthermore, although our link architecture is capable of operating at uncorrected BER $<10^{-12}$, we assume a target BER of 1×10^{-5}, compatible with KR4-FEC (BER threshold $= 2.1 \times 10^{-5}$), and the KP4-FEC (BER threshold $= 2.2 \times 10^{-4}$) that is widely implemented in data center network switches [15].

The ACD link model consists of a quadrature phase-shift keying (QPSK) transmitter, a low-loss optical link ($<$ 2 km), and a homodyne coherent receiver. Fig. 1 illustrates a schematic of the dual-polarization QPSK (DP-QPSK) ACD link for the Si-based architectures. In Fig. 1, the transmitter (TX) laser light is split into two single-mode waveguides and modulated with IQ modulators. We consider two modulator architectures and two photonic integrated circuit (PIC) platforms. The first modulator is based on a traveling wave modulator (TW-MZM) design [16], while the second utilizes a segmented modulator (SEG-MZM) [17]. Both modulators have been demonstrated in Si and InP platforms [18]–[27]. We find that the choice of TX

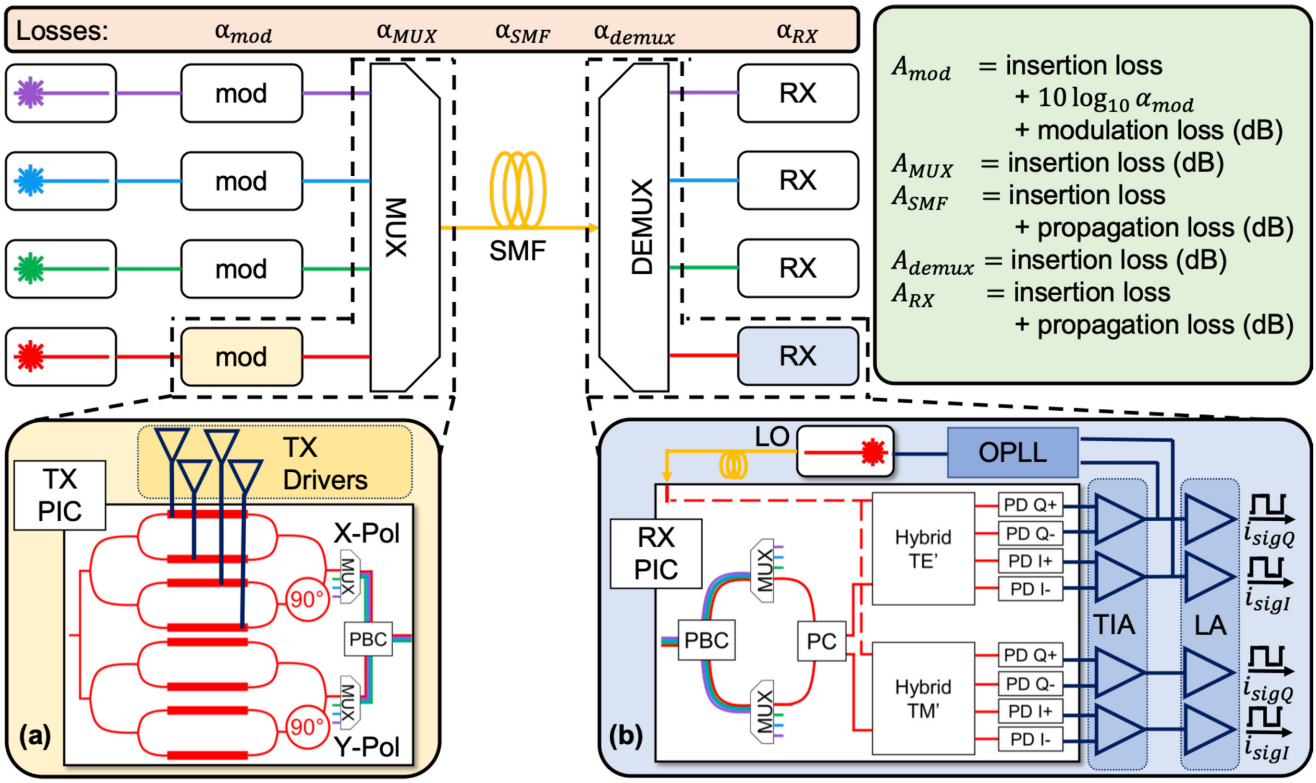

Fig. 1. Shows the link implementation for a QPSK link, where (a) shows the schematic for the QPSK transmitter (TX) considered in the model, while (b) shows the schematic for the receiver (RX), including the OPLL. Note that the MUX/DEMUX is included in the design for the TX and RX, respectively.

architecture and PIC technology both have a significant impact on the overall link performance and power budget. Previous comparisons of SEG-MZMs and TW-MZMs have been made for high-bandwidth radio over fiber (RoF) photonic systems [25]–[27]. RoF links based on SEG-MZM generally showed improvements in gain and noise figure over TW-MZM implementations at high frequency, but at the expense of higher power consumption.

Considering a differential driving signal, a phase modulated signal is realized when an MZM is biased at its null point, where the electric field transmission is 0, as shown in Fig. 2(a). At this bias point, the optical carrier undergoes a 180° phase shift when the input signal transitions from the logic 0 to 1 and vice versa even when the voltage swing is less than twice the full modulator half wave voltage (V_π) of the MZM. However, driving the modulator with a signal amplitude smaller than $2V_\pi$ leads to increased loss. Such loss can be estimated using the modulation factor (F_M), defined in linear units by:

$$F_M = \frac{1}{2}\left(1 - \cos\left(\pi \frac{V_{sig}}{2V_\pi}\right)\right) \tag{1}$$

where V_{sig} is the peak-to-peak drive voltage. The optical loss due to the modulation factor with respect to the drive voltage is shown in Fig. 2(b). A F_M of 1, which corresponds to a $2V_\pi$ drive voltage swing, leads to no modulation-induced loss, while F_M of 0.5 corresponds to a V_π voltage swing and 3 dB of induced optical loss. The modulation factor therefore presents a fundamental power consumption tradeoff: larger drive voltages reduce modulation loss at the expense of higher power dissipation for the electrical modulator driver circuits. Conversely, lower drive voltages reduce driver power but increase optical losses that need to be compensated by higher source and/or LO laser power levels. The modulation loss is therefore controlled by the drive voltage amplitude and is independent of MZM insertion loss. The MZM length is also a key parameter that trades off optical propagation losses against electrical power dissipation in the driver circuits. The relationship between MZM length and optical propagation loss is given by

$$P = P_{in}\, e^{-\alpha_{opt} L_{MZM}} \tag{2}$$

where P_{in} is the input power, L_{MZM} is the active length of the modulator, and α_{opt} is the loss per length of the active region. In InP, we measured this to be 4.34 dB/mm, while in Si we used a value of 1.5 dB/mm. Modulator V_π inversely depends on modulator length: longer modulators exhibit lower V_π, and require relatively lower drive voltages at the expense of higher optical propagation losses; shorter modulators have lower optical losses but higher V_π with correspondingly higher drive voltage requirements and accompanying driver power consumption.

We assume differential drive for both SEG-MZM and TW-MZM transmitters and incorporate polarization multiplexing to increase the link capacity by a factor of two. V_{sig} was swept from 0 to V_π. In the TW-MZMs, we assume that the microwave and optical velocities are matched such that the phase shift induced by the traveling wave electrodes is integrated along

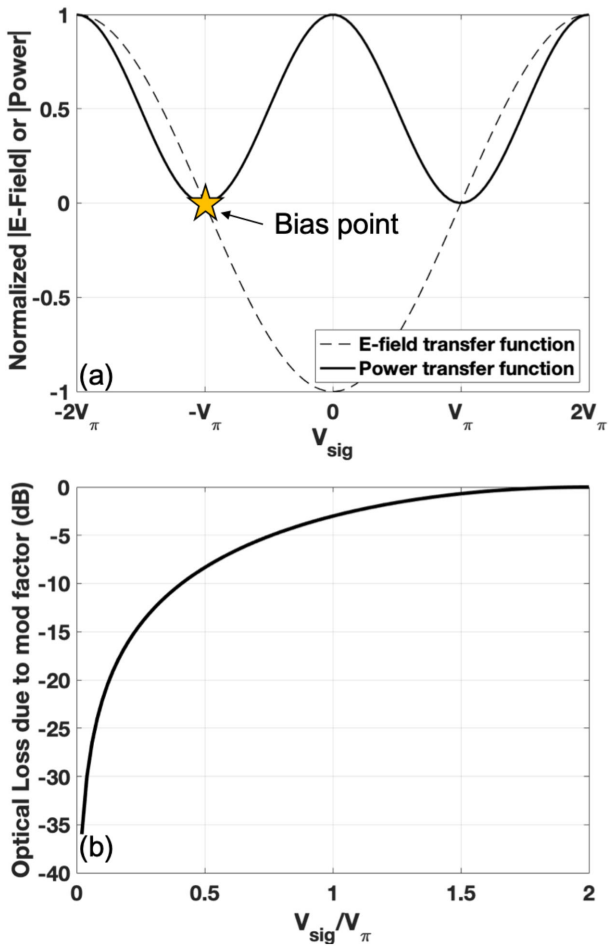

Fig. 2. (a) Electric field and associated output optical power of an MZM and the bias point of the MZM for QPSK modulation. (b) Optical loss in dB due to modulation factor plotted vs. the ratio of $\boldsymbol{V_{sig}/V_{\pi}}$.

the length of the MZM. If the optical and microwave velocities are not perfectly matched, there is a well-known degradation in bandwidth that would introduce an additional inter-symbol interference (ISI) power penalty. In the TW-MZMs, we account for electrode loss along the length of the modulator, which we estimated from simulations of traveling wave electrodes designed for operation above 50 GHz. The electrode loss was estimated to be 0.2 Np/mm and 0.4 Np/mm from simulations in Si and InP, respectively. For the SEG-MZMs, we assume that the driver accounts for the time delay between phase shifter sections and that the voltage delivered to each segment is the same for all segments.

For both driver power calculations, we consider only the power dissipation in the output stage. In both calculations, we assume a 45 nm CMOS technology and that η_{dr} is the efficiency of the driver. For the TW-MZM driver power consumption, we assume differential drive, current-mode logic operation, and that V_{sig} is the single-ended peak output voltage of the driver. The single-ended voltage swing delivered to the MZM is also dependent on the characteristic impedance of the transmission line (Z_0) which was set to 40 Ω for Si and 30 Ω for InP. 40 Ω was chosen for Si based on typical values found in the literature [28]–[31], while 30 Ω for InP was measured from initial designs. Changing the TW-MZM impedance to say, 50 Ω, would not change the bandwidth or the phase efficiency in the calculations, as this relationship is not captured by the model. The power consumption for the SEG-MZM is dependent on the length (l_{seg}) and capacitance (C_{seg}) of each segment, the number of segments being driven, and the baud rate (R_b). Like the TW-MZM driver calculation, V_{sig} is the single-ended peak output voltage of the driver. To calculate the number of segments, we first assumed that the length of each segment was 200 μm to ensure they would behave as lumped circuit elements under 50 GBd operation. The active length of the modulator was then divided by the segment length and rounded up to give an integer number of segments. C_{seg} was measured in initial test structures to be 0.27 fF/mm and 0.94 fF/mm for Si and InP, respectively. We define the driver efficiency as a ratio of the capacitance of each segment to the sum of the segment capacitance and the output capacitance of the driver. It is given by

$$\eta_{dr,SEG} = 1 - G_C \frac{V_{sig}}{t_r} \tag{3}$$

where G_C is the ratio of output capacitance to drain current for a given transistor, and t_r is the signal rise time. In other words, for a given drive voltage and rise time, the efficiency is set by the physical transistor parameters. We found that decreasing the process node to 22 nm, did not yield a significantly higher efficiency due to only slight changes between the processes in G_C. A 45 nm CMOS process, consistent with our calculations, and that will feature full monolithic integration with high-performance Si photonic devices is currently under development [32].

After propagating through up to 2 km of SMF, the WDM signal is coupled into the RX as shown in Fig. 1. While propagating through the fiber, the light undergoes random polarization rotation, necessitating polarization recovery in the RX PIC. In all of our analysis presented here, we assume a Si photonic implementation of the coherent RX PIC, as low-loss, on-chip polarization de-multiplexing can be much more readily realized compared to monolithic InP platforms. When light enters the RX, it first goes through a polarization splitter rotator (PSR) that separates incoming TE and TM components and rotates the TM component to TE for propagation through the on-chip waveguides that natively only support low-loss propagation of TE polarized light. After the PSR, we refer to the two propagating polarizations as TE' and TM', the latter of which has been rotated to TE. Both TE' and TM' contain a mix of the original transmitted polarizations, and must be further processed to recover and separate the original X- and Y-polarizations modulated at the TX.

The TE' and TM' signals next go through separate wavelength demultiplexers that separate individual wavelengths. For each wavelength, the TE' and TM' signals are processed by a polarization controller. Since polarization recovery is conducted for each wavelength, future networks incorporating wavelength routing or switching are readily supported—it doesn't matter if each wavelength entering the receiver has traversed a different path through the data center network. We have implemented a polarization controller using a three-stage device, described

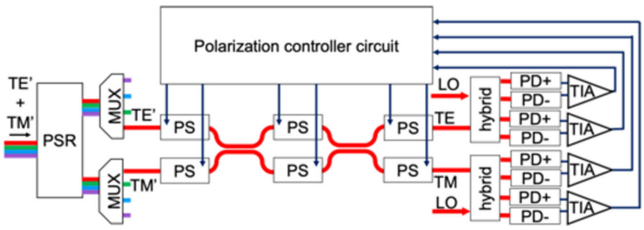

Fig. 3. Three-section polarization controller utilizing phase shifters (PS) after the polarization splitter rotator (PSR) with a polarization controller circuit. This scheme would be implemented for each wavelength.

in greater detail in [2], and shown schematically in Fig. 3. Six thermo-optic phase shifters can be configured to fully separate the original X- and Y-polarizations from the received TE' and TM' signals. After polarization recovery, the X and Y signals are sent to separate 90° hybrids. The polarization recovery scheme exploits a low frequency (few MHz) pilot tone impressed on the I-component of the X-polarization at the TX. In the RX, the low frequency pilot tone is separated from the information-bearing signal by means of a low-pass filter and fed to a low-speed microcontroller. A feedback loop tunes the six phase shifters in the polarization controller, minimizing the pilot tone for all the 90° hybrid output channels except the I-component of the X-polarization. The thermo-optic phase shifters have a response time on the order of tens of microseconds, sufficient to track polarization variations under nominal operating conditions, and can be controlled by a low-cost and ultra-low power microcontroller. We have previously demonstrated a multi-channel thermal phase shifter driver in [33]. The printed circuit board (PCB) implementation consumed a total of 400 mW in the unloaded configuration with an additional average 20 mW per thermal phase shifter connected to the driver. However, to be adapted for the polarization controller described here, the total channels would be reduced from 96 channels to six, with a total of three used at any time, thus significantly reducing the total power consumption of the circuit to be less than 100 mW.

The receiver includes an integrated LO laser for each wavelength channel that is split and then mixed with the incoming signal in separate 90° optical hybrids for the X and Y polarizations. Each optical hybrid produces four outputs, the I+/I- and Q+/Q- signal components, which are subsequently detected by high-speed photodiodes. The responsivity and input sensitivity of the photodiodes were set to 1 A/W and 50 μA. The detected photocurrents are converted to voltage signals by transimpedance amplifiers (TIAs) and then fully converted into digital signals via limiting amplifiers (LAs) that make hard binary decisions. A key advantage of our approach is being able to utilize TIA and LA circuits similar to those in proven NRZ designs that typically achieve the best energy efficiencies with lowest BER [12]. The TIA and LA outputs in both I and Q paths are tapped as inputs for the OPLL circuitry that keeps the LO frequency and phase-locked to the incoming signal. The OPLL is implemented as a Costas loop, providing phase and frequency detection, and is specifically designed for QPSK modulation [9]. The Costas loop architecture has been demonstrated for robust 40 Gbps BPSK operation (10^{-12} at 35 Gbps) across temperature variations of 2.6 °C [9], [10]. Furthermore, wide frequency pull-in range of +/− 30–40 GHz and phase-lock in less than 10 ns have been achieved [9], [10]. OPLLs also provide a high-level of wavelength selectivity through the rejection of all other incoming wavelength channels, reducing the sensitivity to optical crosstalk when scaling to higher numbers of WDM channels. [9]

TABLE I
CIRCUIT POWER CONSUMPTION

TW-MZM Driver	SEG-MZM Driver	TIA-LA-OB chain	OPLL
$\frac{V_{sig}^2}{8Z_0\eta_{dr,TW}}$	$\left\lceil\frac{L_{MZM}}{l_{seg}}\right\rceil\frac{C_{seg}V_{sig}^2 f_b}{\eta_{dr,SEG}}$	198 mW	53 mW

The BER is determined by the Q-factor at the receiver, which is directly related to SNR. For QPSK, SNR = Q^2 [29]. All analysis here assumes a BER $=10^{-5}$, or Q $\approx$ 4.26, unless stated otherwise. This is sufficient to reach the KR4-FEC BER threshold of 2.1×10^{-5}. For homodyne detection, the SNR is

$$SNR = \frac{\langle I_{ac}^2 \rangle}{\sigma^2} = \frac{4R^2\alpha_{coh}P_{laser}P_{LO}}{2qR\Delta f\,(P_{LO} + \alpha_{coh}P_{laser} + I_d/R) + \sigma_T^2} \tag{4}$$

where R is the responsivity of the photodiode, $\alpha_{coh}P_{laser}$ and P_{LO} are the transmitter and LO laser powers at the photodiode, respectively, Δf is the bandwidth of the signal, I_d is the dark current in the photodiode, and σ_T^2 is the thermal noise power. In our model, F_M-induced losses are included in the total link attenuation, α_{coh}. [34]

The energy per bit calculation is given by

$$EPB = (P_{RX\ IC} + P_{OPLL} + P_{TX\ dr} + \eta P_{TX\ laser} + \eta P_{LO})/R_b \tag{5}$$

where $\eta P_{TX\ laser}$ and ηP_{LO} represent the wall plug efficiency of the TX and LO lasers, respectively; $P_{TX\ dr}$ represents the MZM driver, as described in the first two columns of Table I; $P_{RX\ IC}$ and P_{OPLL} represent the receiver chain—including TIA, LA, and output buffer (OB)—and OPLL including the polarization control loop, respectively; and R_b represents the total bit rate. The circuit power dissipation was extracted from transistor-level Spectre simulations in Cadence Virtuoso in the GlobalFoundries SiGe 8XP BiCMOS process for all circuits except segmented modulator driver that were calculated as indicated in Table I.

Other technology-dependent losses, such as waveguide passive attenuation, are also included in the model, and are listed in Table II. The laser efficiencies of the TX and LO lasers are both set to 20%. The TX laser power was swept from 0 to 30 dBm, while the LO laser power was swept from -10 to 20 dBm, though we do not expect to feasibly operate the lasers at powers over 15 dBm. In all the simulations that follow, unless otherwise noted, the unallocated link budget and symbol rate are set to 13 dB and 56 GBd, respectively.

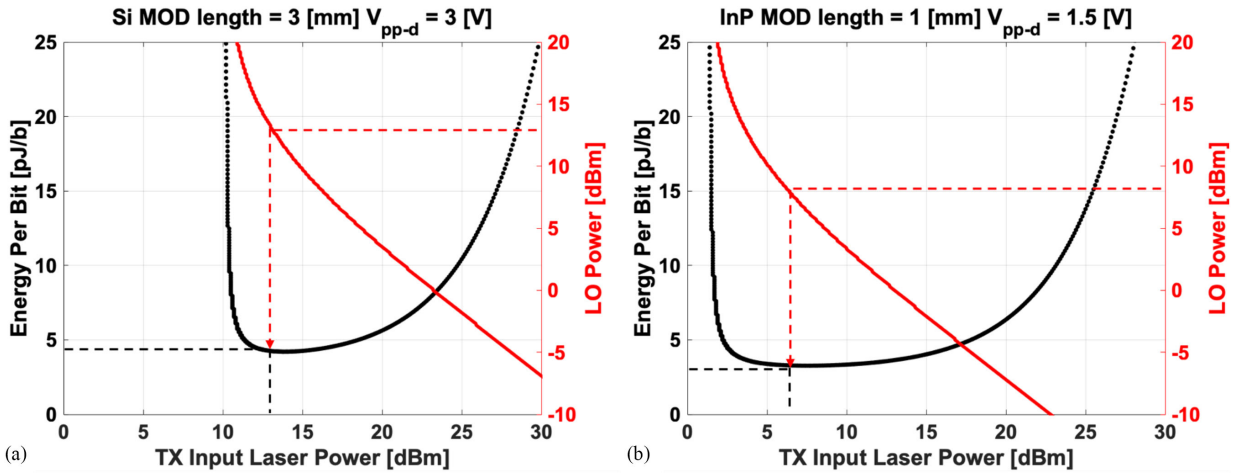

Fig. 4. Simulation results for (a) a 3-mm-long Si Tx/Si Rx TW-MZM and (b) a 1-mm-long InP Tx/Si Rx TW-MZM. The EPB curve (black) and LO power curve (red) correspond to BER $= 1 \times 10^{-5}$, below the KR4-FEC threshold of BER $= 2.1 \times 10^{-5}$.

TABLE II
INVARIANT LINK LOSS PARAMETERS

Parameter	InP (dB)	Si (dB)
FC	1.5	1.5
MUX	1.5	1.5
PSR	1	1
Splitter	3	3
Excess (TX)	2	2
PC	2	2
Hybrid	6	6
Excess (RX/LO)	1	1
TX Loss	9	10.5
RX Loss	13	13
LO Loss	10	11.5

Invariant loss parameters through the link are given in the table. The TX loss is equal to the sum of twice the FC loss, the loss of the MUX, PSR, and splitter, and excess loss for the Si TX. In the InP TX, the fiber coupling loss occurs once in the sum since it is assumed that there is an on-chip laser. The RX loss is equal to the sum of losses due to FC, MUX, PSR, polarization controller, hybrid, and excess losses. The loss in the LO path is the sum of FC, splitter, hybrid, and excess losses for the Si RX. For the InP implementation, the FC is neglected since it is assumed that there is an on-chip laser. FC = fiber coupling.

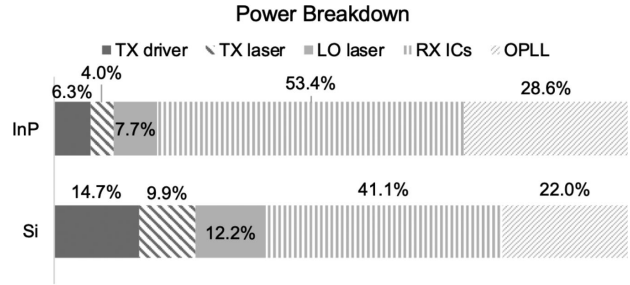

Fig. 5. The proportion of power taken up by each component in the link at the operating point indicated in Fig. 4.

III. RESULTS

Fig. 4 plots the LO power as a function of the TX laser power required to achieve a BER of 10^{-5}. Fig. 4(a) and (b) refer to Si photonic and InP TW-MZM, respectively. The plots also present the energy per bit (EPB) required to achieve BER $= 10^{-5}$ (Q $= 4.26$) as a function of the transmitted power. Fig. 4 therefore informs the available design and operating space for ACD links, indicating that sub-5 pJ/bit energy efficiency is possible for MZMs of practical active lengths in both Si and InP technologies. Fig. 5 shows the power contribution from the various components included in the link at the operating point indicated in Fig. 4. We found the expected trade-off between drive voltage and TX and LO laser powers; that is, one can reduce the drive voltage but must increase the laser powers to overcome the higher incurred losses. However, the reduction of drive voltage does not necessarily achieve the minimum EPB. In general, the EPB decreases with increasing MZM length because an increase in the modulation length leads to a reduction in the required drive voltage. This is true until the increasing electrical and optical losses due to the growing total length of active sections overcome any potential increase in modulation efficiency, thereby increasing the overall EPB, as shown in Fig. 6. Based on our modeling, this inflection does not occur until the length of the MZM grows to about 3-4 mm in Si and 1 mm in InP.

For Si MZMs, we chose a 3-mm active length as a reasonable tradeoff from a device-density and packaging perspective. From the literature, we assume $V_\pi L_\pi = 19$ V-mm [17], [35]–[39], which we have also confirmed through device testing. In estimating the TX loss, we included the optical waveguide loss due to undoped and doped sections as well as the splitters and couplers. Finding a point along the LO power curve in Fig. 4(a) that balances the TX and LO laser powers indicates around 13 dBm is required for both the LO and TX for the Si TW-MZM based link. For Si SEG-MZM links, balanced TX and LO laser powers were found to be ~12.5 dBm, and a similar EPB value of around 4 pJ/bit was projected for a voltage of 1.2 V_{pp-d}.

Transmitters incorporating InP MZMs offer improvements in both link efficiency and laser power requirements compared to Si MZMs. Since the InP platform shows higher modulation efficiency (lower $V_\pi L_\pi$) and higher passive waveguide loss

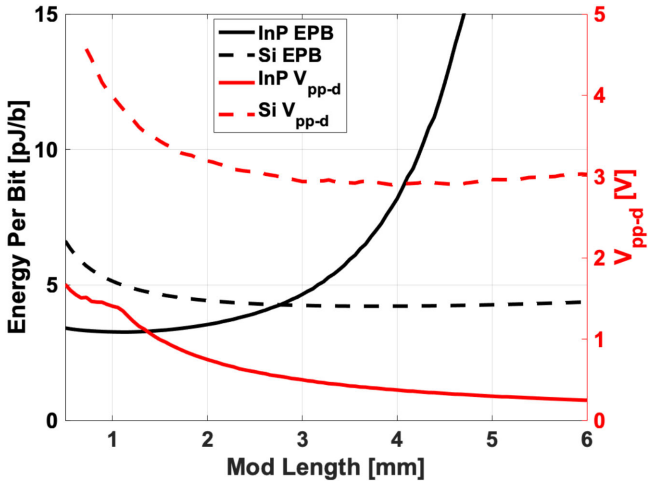

Fig. 6. Simulation results comparing the minimum EPB (black) and drive voltage (red) for InP and Si TW-MZMs. SEG-MZMs yield similar minimum EPB and lower drive voltage.

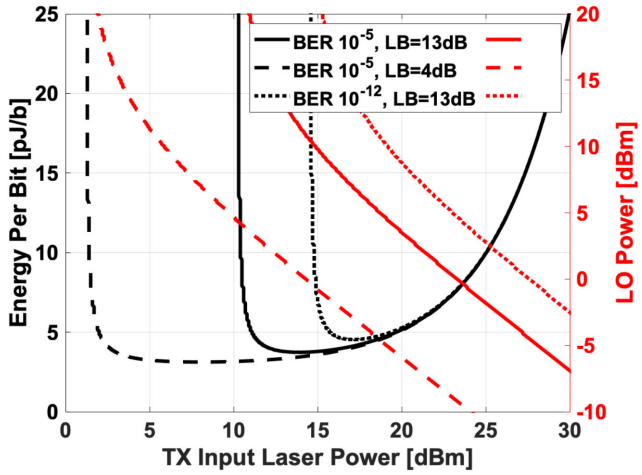

Fig. 7. Starting from the conditions used to generate the results for 3-mm-long Si TWMZMs shown in Fig. 4(a), LO power and EPB vs. TX power for BER 10^{-5} with a 4 dB link margin and BER 10^{-12} with a 13 dB link margin cases are shown. The drive voltage is set to 3 $\boldsymbol{V_{pp-d}}$. LB = link budget.

compared to the Si platform, we constrained the design to a 1-mm long device. We also found that a 1-mm long InP device is the optimal design point for achieving the minimum EPB, as can be seen in Fig. 6. For the InP simulations, we model $V_\pi L_\pi$ as 2 V-mm—an order of magnitude lower than Si. The InP TW-MZM design point is very realizable, requiring only 1.5 V_{pp-d} and roughly $+7$ dBm from both the input and LO lasers as can be seen in Fig. 4(b), while the InP SEG-MZM design point suggests that TX and LO laser powers of $+5$ dBm and drive voltage of 1 V_{pp-d} to be the ideal operating point for a similar EPB. InP device parameters are based on measured data of PICS fabricated in the UCSB Nanofabrication Facility.

Note that although InP TXs offer a more efficient solution than Si, a Si RX implementation is more favorable as on-chip polarization de-multiplexing is more readily implemented. In addition, an all Si implementation likely offers advantages in electronic and photonic integration, packaging, and cost. While it is possible to achieve low EPB, Si TXs have higher loss due to less efficient MZMs that degrade link efficiency and drive laser power requirements to challenging levels. Furthermore, both InP and Si SEG-MZMs will face significant challenges in integrating a large number of drivers alongside or flip-chipped on the modulator due to the large number of segments that make up the total active length required to achieve sufficient modulation. Thus, we focus on TW-MZMs for the rest of the analysis presented here. However, SEG-MZMs with integrated drivers could be a very compelling solution in monolithic processes that offer high-performance CMOS capable of 50 GBd operation [33].

In Fig. 6, we compare InP and Si TW-MZM EPB and drive voltage with respect to modulator length. The much lower $V_\pi L_\pi$ of InP contributes to a much lower drive voltage compared to the Si structure. However, a steep rise in EPB is projected for longer InP modulators. In this regime, propagation losses in the active sections dominate, driving up laser power requirements that overwhelm any power savings due to reduced drive voltage. This optimal point in EPB occurs at about 1 mm for InP due to much higher propagation losses in its active sections but is not a significant factor for modulators less than 6 mm long in Si. The 3-mm Si TW-MZM link can achieve energy efficiency approaching the InP MZMs, although still requiring higher power operating points for both TX and LO lasers. Combining this observation with the fact that Si offers a cost-effective platform for building large-scale, highly integrated PICs in 300 mm wafer manufacturing processes, we focus on Si TW-MZMs for the rest of our analysis in this paper. However, the observations and conclusions that follow are applicable to links incorporating InP MZMs.

To further explore the design space available for ACD link architectures, simulations were conducted for lower BER targets, reduced link budgets, and single polarization operation. Selected results are presented here for links with Si TW-MZM transmitters. Moving to a more aggressive BER target, 10^{-12}—often referred to as "error-free"—has a minor impact on the achievable EPB for fixed modulator lengths but drives both source and LO lasers powers considerably higher, to around $+16$ dBm, as seen in Fig. 7. Increasing the output drive voltage can also potentially result in decreased bandwidth of the driver circuits. Reducing the link budget from 13 to 4 dB has the opposite effect on link operation. Effects on the minimum achievable EPB are also minimal, but the required LO and input laser powers are reduced significantly to about $+6$ dBm. The EPB curve is also substantially flattened, indicating a wider range of choices for LO and TX laser power that achieve the optimal EPB.

Fig. 8 shows how EPB changes with voltage, assuming an operating point where the TX and LO laser powers are equal based upon the same link configurations analyzed in Fig. 7. Here it is evident that while the minimum EPB may occur at lower drive voltages, the TX and LO output powers may be unfeasible.

Finally, we investigated the case of a single polarization link, starting from the baseline Si TW-MZM case presented in Fig. 4(a). One may expect that the EPB will decrease with lower

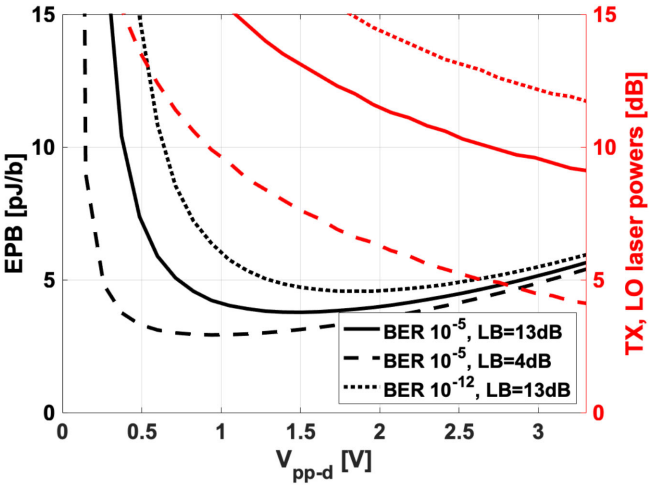

Fig. 8. Starting from the conditions used to generate the results for 3-mm-long Si TWMZMs shown in Fig. 4(a), and taking the operating point where the TX and LO laser powers are equal, this plot shows how the EPB changes for a given drive voltage and the required TX/LO laser powers to close the link. LB = link budget.

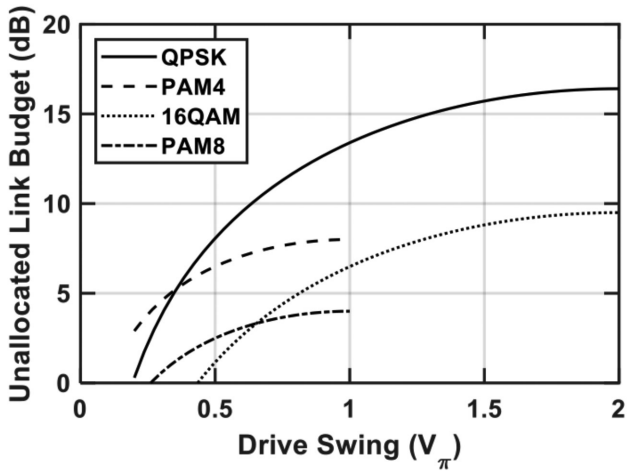

Fig. 9. Comparison of unallocated link margin in coherent and IMDD links, assuming MZM drive, and a target BER of 10^{-5}. The QPSK curve assumes analog coherent link performance as described in this work, while the other curves assume representative link performance projections for next-generation IMDD and digital coherent links [5].

complexity of the PIC due to the lack of polarization-specific components, but in fact, the change is not significant. This is because ICs take up a significant portion of the total link energy consumption at these optical powers (100s of mW compared to 10s of mW). Therefore, if we halve the number of ICs in the receiver, we reduce the IC energy consumption of the link almost by half for half the number of bits. On the other hand, by eliminating the need for sharing TX and LO lasers between two polarizations, the operating points of both lasers are reduced by roughly 3 dB to 8 dBm.

IV. Discussion

In the previous section we have shown how we can achieve sub-10 pJ/b, 200 Gbps per wavelength links. Increasing the data rate in a single lane by implementing a higher order modulation format is another path to higher aggregate link bandwidth but will decrease the unallocated link budget [5]. Utilizing the additional loss budget by inserting optical switches can decrease cost, latency, and power consumption of data centers. We will show in this section that increasing to higher order modulation rates decreases the unused link margin and may restrict connectivity between servers.

A. Unallocated Link Budget

Optical switching is the subject of worldwide research, motivated by the promise of adding reconfigurability to data center networks and potentially improving overall data center energy efficiency [39]–[41]. The principle of adding a layer of arrayed waveguide grating routers (AWGRs) or optical switches layer to a data center to enhance scalability while reducing cost, power and latency, is described in [39] and [40], respectively. However, for optical switching to be practical, the links traversing the switches must either have enough budget to accommodate the losses of the switches, or the switches must be made transparent by incorporating optical gain. The latter approach, usually relying on semiconductor optical amplifiers, presents integration challenges in Si photonic platforms and also raises operational issues including added noise, gain uniformity across wavelengths, and crosstalk [42]. We believe the best approach to enable photonic switching is through expanding available link budgets to accommodate the insertion loss of switching or passive wavelength routing components. In order to assess the achievable link budgets offered by candidate link architectures, we follow the analysis approach in [5] to compare QPSK to both IMDD alternatives as well as 16QAM. The results are presented in Fig. 9 for an analysis conducted under a consistent set of assumptions for each link: the same laser powers, MZM modulators, and target BER of 10^{-5}. For drive swings above V_π, 16QAM can offer some improvement in budget compared to IMDD, but the advantages of QPSK are much more substantial. At full $2V_\pi$ drive levels, QPSK expands link budgets by 8 dB compared to PAM4 and 12 dB compared to PAM8. At a more practically realizable drive voltage of 0.6 V_π, QPSK offers increases of 2 dB and 6 dB compared to PAM4 and PAM8, respectively. In addition to enabling photonic switching in data center networks, the expanded link budgets offered by QPSK-based ACD can also potentially be used to improve transceiver yields owing to the reduced sensitivity to optical loss as well as reduce transceiver power consumption by lowering the operating points of the lasers and drivers.

B. Optical Switch-Based Networks

Here we illustrate the design flexibility offered by scaling bandwidth by adding additional 200 Gbps/λ WDM channels as opposed to increasing per-λ bit rates. A relevant design example is a data center network connecting 131,072, 50-Gb/s servers, using next-generation 51.2 Tb/s switches. A conventional design, using 800-Gb/s inter-switch links, would employ three levels of

electronic switches, with 51.2 Tb/s switches used in both the spine and aggregation layers, and a smaller 6.4 Tb/s ToR switch (supporting 64 servers per rack). If the 800 Gbps transceivers were realized using four independent, separable lanes operating at 200-Gb/s, the same number of servers could be supported using only two levels of 51.2 Tb/s switches interconnected by 200-Gb/s links. The drawback in this scenario is that four times as many fibers would be needed. However, the number of fibers can be reduced back to the original number of fibers as in the classic design by using four 200-Gb/s wavelengths per fiber, and inserting an optical wavelength-routing layer, consisting of 4x4 AWGRs, between the two electronic switching levels, described in detail in [39]. With only two levels of switches instead of three, this optimized design, taking full advantage of WDM parallelism, results in lower cost, latency, and power consumption. Thus, as described in the Introduction and as shown in Fig. 1, to realize an 800G link, we include four lanes multiplexed into a single fiber. Likewise, to scale to higher data rates, such as 1.6 Tb/s, eight lanes would be multiplexed into a single fiber, without changing the data format to take advantage of the inherent link budget advantage of QPSK as well as the efficiency advantages presented by being able to implement low-power electronics similar to the circuits used in NRZ link implementations that we have implemented in our ACD link.

The wavelength routing layer used to flatten the data center network in the above example could be based on all passive elements such as AWGRs, but there is also the possibility to add reconfigurability to the wavelength routing layer in the form of small (e.g., 4×4, 8×8, or 16×16) WDM photonic switches. Furthermore, the power consumed by optical switches are independent of the data rate of the signals they route. Optical switches do not need to perform power-hungry optical-electrical-optical conversions and instead perform the switching in the optical domain while being transparent to data format and data rate [43]. Consequently, as links move to higher data rates, their impact on the total energy per bit decreases.

Focusing on switches implemented in planar Si platforms, which offer a realizable path towards mass manufacturing, there have been several noteworthy recent demonstrations of Si photonic switches, including: a 4×4 switch with integrated gain [42], a 32×32 port polarization diverse switch [44], and a 240×240 port MEMS-based switch [45], among many others. We have demonstrated a wavelength-selective crossbar switch with multiple wavelength-selective elements at each cross-point [46]. Each switch offers a promising feature: the large port count from stitching of multiple die in [45], the path-independent loss in addition to the polarization diversity in [44], and the gain and custom ASIC integration in [42]. While these demonstrations differ in port count and switching time, they are non-blocking and exhibit losses less than the unallocated link budget assumed in the link model presented above. Having a large unallocated link budget eases the requirement of having ultra-low-loss photonic components in the switch design. This ensures that the insertion loss of the switch remains low so that the port count can be scaled to 32-64, offering flexibility in data center network architectures.

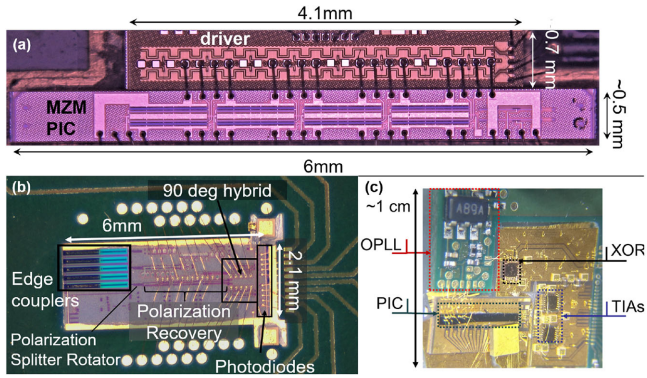

Fig. 10. Preliminary hardware for (a) a Si TX modulator and driver, (b) Si RX PIC, and (c) InP RX PIC packaged with an OPLL.

V. Conclusion

A comprehensive QPSK coherent link model has been presented and indicates that EPB under from 5-10 pJ/bit is possible with substantial improvements in optical loss budget. The simulation tool allows exploration of optical and electrical parameters that impact PIC design. Measured hardware will verify and refine the parameters used in the link analysis. Fig. 10 shows functional hardware that we have built and characterized to provide hardware-derived inputs to our modeling. Design and characterization results for the transmitter in Fig. 10(a) have been reported in [47], while receiver subsystems shown in Fig. 10(b) and (c) as well as other transmitters will be reported in forthcoming publications. Finally, we showed that QPSK links increase unallocated link budget. Analog coherent detection based on QPSK modulation has the potential to enable novel network designs incorporating wavelength routing and switching while simultaneously maximizing energy efficiency, facilitation future lower power data center network architectures that maximize overall data center efficiency.

Acknowledgment

The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

The authors would also like to thank GlobalFoundries for their foundry services. In particular, the authors would like to thank K. Giewont, A. Stricker, K. Nummy, D. Riggs, K. Dezfulian, and T. Letavic at GlobalFoundries for their support and assistance.

References

[1] X. Pang *et al.*, "200 Gbps/Lane IM/DD technologies for short reach optical interconnects," *J. Lightw. Technol.*, vol. 38, no. 2, pp. 492–503, Jan. 2020.

[2] J. K. Perin, A. Shastri, and J. M. Kahn, "Design of low-power DSP-free coherent receivers for data center links," *J. Lightw. Technol.*, vol. 35, no. 21, pp. 4650–4662, Nov. 2017.

[3] T. Hirokawa, S. Pinna, J. Klamkin, J. F. Buckwalter, and C. L. Schow, "Energy efficiency analysis of coherent links for datacenters," in *Proc. IEEE Opt. Interconnects Conf.*, 2019, pp. 1–2.

[4] J. Cheng, C. Xie, Y. Chen, X. Chen, M. Tang, and S. Fu, "Comparison of coherent and IMDD transceivers for intra datacenter optical interconnects," in *Proc. Opt. Fiber Commun. Conf.*, 2019, Paper W1F.2.
[5] X. Zhou, R. Urata, and H. Liu, "Beyond 1 Tb/s intra-data center interconnect technology: IM-DD OR coherent?" *J. Lightw. Technol.*, vol. 38, no. 2, pp. 475–484, Jan. 2020.
[6] S. Ristic, A. Bhardwaj, M. J. Rodwell, L. A. Coldren, and L. A. Johansson, "An optical phase-locked loop photonic integrated circuit," *J. Lightw. Technol.*, vol. 28, no. 4, pp. 526–538, Feb. 2010.
[7] M. Lu, H. Park, E. Bloch, L. A. Johansson, M. J. Rodwell, and L. A. Coldren, "An integrated heterodyne optical phase-locked loop with record offset locking frequency," in *Proc. Opt. Fiber Commun. Conf.*, Mar. 2014, Paper Tu2H.4.
[8] M. Lu *et al.*, "A heterodyne optical phase-locked loop for multiple applications," in *Proc. Opt. Fiber Commun. Conf. Expo. Nat. Fiber Opt. Eng. Conf.*, 2013, pp. 1–3.
[9] M. J. W. Rodwell *et al.*, "Optical phase-locking and wavelength synthesis," in *Proc. IEEE Compound Semicond. Integr. Circuit Symp.*, 2014, pp. 1–4.
[10] M. Lu *et al.*, "An integrated 40 Gbit/s optical costas receiver," *J. Lightw. Technol.*, vol. 31, no. 13, pp. 2244–2253, Jul. 2013.
[11] A. Simsek *et al.*, "A chip-scale heterodyne optical phase-locked loop with low-power consumption," in *Proc. Opt. Fiber Commun. Conf. Exhib.*, 2017, pp. 1–3.
[12] L. Szilagyi, J. Pliva, R. Henker, D. Schoeniger, J. P. Turkiewicz, and F. Ellinger, "A 53-Gbit/s optical receiver frontend with 0.65 pJ/bit in 28-nm bulk-CMOS," *IEEE J. Solid-State Circuits*, vol. 54, no. 3, pp. 845–855, Mar. 2019.
[13] L. Kull *et al.*, "A 24-to-72 GS/s 8-b time-interleaved SAR ADC with 2.0-to-3.3 pJ/conversion and >30 dB SNDR at Nyquist in 14 nm CMOS FinFET," in *Proc. IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2018, pp. 358–360.
[14] K. Sun, G. Wang, Q. Zhang, S. Elahmadi, and P. Gui, "A 56-GS/s 8-bit time-interleaved ADC with ENOB and BW enhancement techniques in 28-nm CMOS," *IEEE J. Solid State Circuits*, vol. 54, no. 3, pp. 821–833, Mar. 2019.
[15] Y. Yue, Q. Wang, J. Yao, J. O'Neil, D. Pudvay, and J. Anderson, "400GbE technology demonstration using CFP8 pluggable modules," *Appl. Sci.*, vol. 8, no. 11, Oct. 2018, Art. no. 2055.
[16] M. Streshinsky *et al.*, "Low power 50 Gb/s silicon traveling wave Mach–Zehnder modulator near 1300 nm," *Opt. Express*, vol. 21, pp. 30350–30357, Dec. 2013.
[17] B. G. Lee *et al.*, "Driver-integrated 56-Gb/segmented electrode silicon Mach–Zehnder modulator using optical-domain equalization," in *Proc. Opt. Fiber Commun. Conf.*, 2017, Paper. Th1B.1.
[18] K. Goi *et al.*, "128-Gb/s DP-QPSK using low-loss monolithic silicon IQ modulator integrated with partial-rib polarization rotator," in *Proc. Opt. Fiber Commun. Conf.*, Mar. 2014, Paper W1I.2.
[19] P. Dong, L. Chen, C. Xie, L. L. Buhl, and Y.-K. Chen, "50-Gb/s silicon quadrature phase-shift keying modulator," *Opt. Express*, vol. 20, pp. 21181–21186, 2012.
[20] K. Goi *et al.*, "DQPSK/QPSK modulation at 40-60 Gb/s using low-loss nested silicon Mach–Zehnder modulator," in *Proc. Opt. Fiber Commun. Conf./Nat. Fiber Opt. Eng. Conf.*, 2013, Paper OW4J.4.
[21] N. Kono *et al.*, "Compact and low power DP-QPSK modulator module with InP-based modulator and driver ICs," in *Proc. Opt. Fiber Commun. Conf. Exp. Nat. Fiber Opt. Eng. Conf.*, 2013, pp. 1–3.
[22] E. Yamada *et al.*, "112-Gb/s InP DP-QPSK modulator integrated with a silica-PLC polarization multiplexing circuit," in *Proc. Nat. Fiber Opt. Eng. Conf.*, Mar. 2012, Paper PDP5A.9.
[23] N. Kikuchi, E. Yamada, Y. Shibata, and H. Ishii, "High-speed InP-based Mach–Zehnder modulator for advanced modulation formats," in *Proc. IEEE Compound Semicond. Integr. Circuit Symp.*, 2012, pp. 1–4.
[24] P. Evans *et al.*, "Multi-channel coherent PM-QPSK InP transmitter photonic integrated circuit (PIC) operating at 112 Gb/s per wavelength," in *Proc. Opt. Fiber Commun. Conf./Nat. Fiber Opt. Eng. Conf.*, 2011, Paper PDPC7.
[25] N. Hosseinzadeh, A. Jain, K. Ning, R. Helkey, and J. F. Buckwalter, "A 0.5–20 GHz RF silicon photonic receiver with 120 dB•$Hz^{2/3}$ SFDR using broadband distributed IM3 injection linearization," in *Proc. IEEE Radio Freq. Integr. Circ. Symp.*, 2019, pp. 99–102.
[26] N. Hosseinzadeh, A. Jain, K. Ning, R. Helkey, and J. F. Buckwalter, "A linear microwave electro-optic front end with sige distributed amplifiers and segmented silicon photonic Mach–Zehnder modulator," *IEEE Trans. Microw. Theory Technol.*, vol. 67, no. 12, pp. 5446–5458, Dec. 2019.
[27] N. Hosseinzadeh, A. Jain, K. Ning, R. Helkey, and J. F. Buckwalter, "A 1 to 20 GHz Silicon-Germanium low-noise distributed driver for RF Silicon Photonic Mach–Zehnder modulators," in *Proc. IEEE MTT-S Int. Microw. Symp.*, 2019, pp. 774–777.
[28] D. Patel *et al.*, "Design, analysis, and transmission system performance of a 41 GHz silicon photonic modulator," *Opt. Express*, vol. 23, no. 11, pp. 14263–14287, Jun. 2015.
[29] A. Samani *et al.*, "Silicon photonic Mach–Zehnder modulator architectures for on chip PAM-4 signal generation," *J. Lightw. Technol.*, vol. 37, no. 13, pp. 2989–2999, Jul. 2019.
[30] R. Ding *et al.*, "High-speed silicon modulator with slow-wave electrodes and fully independent differential drive," *J. Lightw. Technol.*, vol. 32, no. 12, pp. 2240–2247, Jun. 2014.
[31] 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, Feb. 2018.
[32] M. Rakowski *et al.*, "45nm CMOS — Silicon photonics monolithic technology (45CLO) for next-generation, low power and high speed optical interconnects," in *Proc. Opt. Fiber Commun. Conf. Exhib.*, 2020, pp. 1–3.
[33] T. Hirokawa *et al.*, "A wavelength-selective multiwavelength ring-assisted Mach–Zehnder interferometer switch," *J. Lightw. Technol.*, to be published, doi: 10.1109/JLT.2020.3011944.
[34] G. P. Agrawal, *Fiber-Optic Communication Systems*. Hoboken, NJ, USA: John Wiley & Sons, 4th ed., 2010.
[35] N. M. Fahrenkopf, C. McDonough, G. L. Leake, Z. Su, E. Timurdogan, and D. D. Coolbaugh, "The AIM photonics MPW: A highly accessible cutting edge technology for rapid prototyping of photonic integrated circuits," *IEEE J. Sel. Top. Quantum Electron.*, vol. 25, no. 5, pp. 1–6, Sep./Oct. 2019.
[36] H. Yu *et al.*, "Performance tradeoff between lateral and interdigitated doping patterns for high speed carrier-depletion based silicon modulators," *Opt. Express*, vol. 20, pp. 12926–12938, 2012.
[37] A. Rahim *et al.*, "Open-access silicon photonics platforms in Europe," *IEEE J. Sel. Top. Quantum Electron.*, vol. 25, no. 5, pp. 1–18, Sep./Oct. 2019.
[38] A. L. Giesecke *et al.*, "Ultra-efficient interleaved depletion modulators by using advanced fabrication technology," in *Proc. 42nd Eur. Conf. Opt. Commun.*, 2016, pp. 1–3.
[39] A. A. M. Saleh, "Scaling-out data centers using photonics technologies," in *Proc. Photon. Switching Conf.*, Jul. 2014, Paper JM4B.5.
[40] A. A. M. Saleh, A. S. P. Khope, J. E. Bowers, and R. C. Alferness, "Elastic WDM switching for scalable data center and HPC interconnect networks," in *Proc. 21st Opto Electron. Commun. Conf.*, 2016, pp. 1–3.
[41] G. Michelogiannakis *et al.*, "Bandwidth steering in HPC using Silicon nanophotonics," in *Proc. Int. Conf. High Perfor. Comput. Netw. Storage Anal.*, Nov. 2029, pp. 1–25.
[42] N. Dupuis *et al.*, "A 4 × 4 electrooptic silicon photonic switch fabric with net neutral insertion loss," *J. Lightw. Technol.*, vol. 38, no. 2, pp. 178–184, Jan. 2020.
[43] F. Testa and L. Pavesi, *Optical Switching In Next Generation Data Centers*. Springer, 2017.
[44] K. Suzuki *et al.*, "Nonduplicate polarization-diversity 32 × 32 Silicon photonics switch based on a SIN/Si double-layer platform," *J. Lightw. Technol.*, vol. 38, no. 2, pp. 226–232, Jan. 2020.
[45] K. Kwon, T. J. Seok, J. Henriksson, J. Luo, and M. C. Wu, "Large-scale silicon photonic switches," in *Proc. Photon. Electromagn. Res. Symp. - Spring*, 2019, pp. 268–273.
[46] T. Hirokawa *et al.*, "An all-optical wavelength-selective O-band chip-scale silicon photonic switch," in *Proc. Conf. Lasers Electro-Opt.*, 2020, Paper SF1L.2.
[47] N. Hosseinzadeh, K. Fang, L. Valenzuela, C. Schow, and J. Buckwalter, "A 50-Gb/s optical transmitter based on co-design of a 45-nm CMOS SOI distributed driver and 90-nm Silicon photonic Mach–Zehnder modulator," in *Proc. IEEE MTT-S Int. Microw. Symp. Dig.*, 2020.

Takako Hirokawa (Student Member, IEEE) received the B.S. degrees in applied math and engineering physics from the University of Colorado Boulder, Boulder, CO, USA, and the M.S. degree in electrical and computer engineering in 2016 from the University of California, Santa Barbara, CA, USA, where she is currently working toward the Ph.D. degree in electrical engineering.

Sergio Pinna (Member, IEEE) received the B.Sc. and M.Sc. degrees in telecommunications engineering from the University of Pisa, Pisa, Italy, and the Ph.D. degree in innovative technologies from the Scuola Superiore Sant'Anna, Pisa, in 2014. From 2010 to 2016, he was a Fellow with the Digital and Microwave photonics Group, National Photonic Networks Laboratory and Scuola Superiore Sant'Anna. In 2017, he joined the Integrated Photonics Laboratory, University of California Santa Barbara, as an Assistant Project Scientist.

Navid Hosseinzadeh (Student Member, IEEE) received the M.Sc. degree in electrical engineering from the Sharif University of Technology, Tehran, Iran, in 2014. He is currently working toward the Ph.D. degree in electrical and computer engineering with the Department of Electrical and Computer Engineering, University of California, Santa Barbara, Santa Barbara, CA, USA. His research interests include RF/millimeter-wave IC design in silicon and compound semiconductor technologies, silicon photonics integrated circuits, and quantum computers. In 2019, he was an R&D Intern with PsiQuantum Corp., Palo Alto, CA, USA. He serves as a Reviewer for the IEEE MICROWAVE AND WIRELESS COMPONENTS LETTERS, IEEE TRANSACTION ON CIRCUITS AND SYSTEMS, and IEEE TRANSACTION ON MICROWAVE THEORY AND TECHNIQUE.

Aaron Maharry received the B.S. degree in electrical and computer engineering from Ohio State University, Columbus, OH, USA, in 2017 and the M.S. degree in electrical and computer engineering in 2019 from the University of California, Santa Barbara, Santa Barbara, CA, USA, where he is currently working toward the Ph.D. degree in electrical and computer engineering.

Hector Andrade is currently working toward the Ph.D. degree in electrical and computer engineering with the University of California, Santa Barbara, Santa Barbara, CA, USA.

Junqian Liu is currently working toward the Ph.D. degree in electrical and computer engineering with the University of California, Santa Barbara, Santa Barbara, CA, USA.

Thomas Meissner is currently working toward the Ph.D. degree in electrical and computer engineering with the University of California, Santa Barbara, Santa Barbara, CA, USA.

Stephen Misak (Student Member, IEEE) received the B.S. degree in engineering physics with a minor in optical engineering from the Rose-Hulman Institute of Technology, Terre Haute, IN, USA, in 2017, and the M.S. degree in electrical and computer engineering in 2019 from the University of California, Santa Barbara, Santa Barbara, CA, USA, where he is currently working toward the Ph.D. degree in electrical and computer engineering.

Ghazal Movaghar is currently working toward the Ph.D. degree in electrical and computer engineering with the University of California, Santa Barbara, Santa Barbara, CA, USA.

Luis A. Valenzuela (Student Member, IEEE) received the B.S. degree in electrical engineering from Cal Poly Pomona, Pomona, CA, USA, in 2015 and the M.S. degree in electrical engineering in 2019 from the University of California, Santa Barbara, Santa Barbara, CA, USA, where he is currently working toward the Ph.D. degree in electrical engineering.

Yujie Xia is currently working toward the Ph.D. degree in electrical and computer engineering with the University of California, Santa Barbara, Santa Barbara, CA, USA.

Shireesh Bhat, biography not available at the time of publication.

Fabrizio Gambini, biography not available at the time of publication.

Jonathan Klamkin (Senior Member, IEEE) received the B.S. degree from Cornell University, Ithaca, NY, USA, and the M.S. and Ph.D. degrees from the University of California, Santa Barbara (UCSB), Santa Barbara, CA, USA. From 2008 to 2011, he was a member of the Technical Staff at the Electro-Optical Materials and Devices Group, MIT Lincoln Laboratory, Lexington, MA, USA. From 2011 to 2013, he was an Assistant Professor at the Institute of Communication, Information and Perception Technologies, Scuola Superiore Sant'Anna, Pisa, Italy. From 2013 to 2015, he was an Assistant Professor of Electrical and Computer Engineering (ECE) and Materials at Boston University, Boston, MA, USA. In 2015, he joined the ECE Department, University of California, Santa Barbara, where he is currently a Professor and the Director of the UCSB Nanotech. He has authored or coauthored nearly 200 journal and conference papers. He is a Managing Editor for the *Nanophotonics* journal, Vice Chair of the Microwave Theory and Techniques Society Subcommittee on Microwave Photonics, and Steering Committee Member of the JOURNAL OF LIGHTWAVE TECHNOLOGY. He was the Program Chair of the Integrated Photonics Research, Silicon and Nanophotonics Conference, in 2017, and the General Chair of the same conference, in 2018. He and his group members were the recipient of Best Paper Awards at the 2006 Conference on Optoelectronic and Microelectronic Materials and Devices, 2007 Microwave Photonics Conference, and 2017 and 2019 Asia Communications and Photonics Conference. He was the recipient of the NASA Early Career Faculty Award, the DARPA Young Faculty Award, and the DARPA Director's Fellowship.

Adel A. M. Saleh (Life Fellow, IEEE) is currently a Research Professor with the Department of Electrical and Computer Engineering and the Institute for Energy Efficiency, University of California, Santa Barbara (UCSB), Santa Barbara, CA, USA, since October 2011, conducting research on optical networking and photonic technology for chip-scale to global-scale networking applications, with emphasis on the interconnect networks of hyper-scale data centers and high-performance computers. From 2005 to 2011, he was a DARPA Program Manager, where he initiated several research programs on advanced optical networking. From 1999 to 2004, he held leadership positions in the optical networking industry, including Corvis, the first company to commercialize national-scale all-optical networks. From 1970 to 1999, he was with Bell Labs/AT&T Labs Research, conducting and leading research on optical and wireless networks. Between 1992 and 1999, he led the AT&T effort on several multi-million-dollar, cross-industry, DARPA-funded consortia that pioneered the vision of all-optical networking in backbone, regional, metro and access networks. He is a Fellow of the OSA.

Larry Coldren (Life Fellow, IEEE) received the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, USA, in 1972. After 13 years in the research area with Bell Laboratories, he joined the Department of Electrical and Computer Engineering, University of California, Santa Barbara (UCSB), Santa Barbara, CA, USA, in 1984 and was a co-founder of the Materials Department in 1986. From 2009 to 2011, he was acting Dean of the College of Engineering. He became Emeritus Professor in 2018 and was appointed a Research Professor that same year.

In 1991, he co-founded Optical Concepts, later acquired as Gore Photonics, to develop novel VCSEL technology and in 1998, he co-founded Agility Communications, later acquired by JDSU (now Lumentum), to develop widely-tunable integrated transmitters. At UCSB, he has worked on multiple-section widely-tunable lasers and efficient vertical-cavity surface-emitting lasers (VCSELs). His group has also developed a variety of high-performance InP-based photonic integrated circuits. He has authored or coauthored more than thousand journal and conference papers, eight book chapters, a widely-used textbook, and 63 issued patents, which have received more than 30 000 citations.

He is a Fellow of OSA, and the National Academy of Inventors as well as a member of the National Academy of Engineering. He was the Recipient of the 2004 John Tyndall Award, 2009 Aron Kressel Award, 2014 David Sarnoff Award, 2015 IPRM Award, and 2017 Nick Holonyak, Jr. Award.

James F. Buckwalter (Senior Member, IEEE) received the Ph.D. degree in electrical engineering from the California Institute of Technology, Pasadena, CA, USA, in 2006. From 1999 to 2000, he was a Research Scientist with Telcordia Technologies, Morristown, NJ, USA. In 2004, he was with the IBM T. J. Watson Research Center, Yorktown Heights, NY, USA. In 2006, he joined the Faculty of the University of California at San Diego, La Jolla, CA, USA, as an Assistant Professor and was promoted to an Associate Professor in 2012. He is currently a Professor of Electrical and Computer Engineering with the University of California, Santa Barbara, Santa Barbara, CA, USA. He was the recipient of the 2004 IBM Ph.D. Fellowship, 2007 Defense Advanced Research Projects Agency Young Faculty Award, 2011 NSF CAREER Award, and 2015 IEEE MTT-S Young Engineer Award.

Clint L. Schow (Fellow, IEEE) received the B.S., M.S., and Ph.D. degrees from the University of Texas at Austin, Austin, TX, USA.

After positions at IBM and Agility Communications, he spent more than a decade with the IBM T. J. Watson Research Center, Yorktown Heights, NY, USA, as a Research Staff Member and a Manager of the Optical Link and System Design group. In 2015, he joined the Faculty at the University of California, Santa Barbara, Santa Barbara, USA. He has led international R&D programs spanning chip-to-chip optical links, VCSEL and Si photonic transceivers, nanophotonic switches, and new system architectures enabled by high-bandwidth, low-latency photonic networks.

Dr. Schow has been an Invited Speaker and served on committees for numerous international conferences, including roles as a General Chair for the Optical Fiber Communications Conference (OFC), the Optical Interconnects Conference, and the Photonics in Switching Conference. He has authored or coauthored more than 200 journal and conference articles and 33 issued patents. He is a Fellow of the OSA.