Data Center Links Beyond 100 Gbit/s per Wavelength

Jose Krause Perin, Anujit Shastri, and Joseph M. Kahn

E. L. Ginzton Laboratory, Department of Electrical Engineering, Stanford University, Stanford, CA, 94305 USA

ABSTRACT

Increased traffic demands within and between data centers now necessitate low-cost and low-power systems with per-wavelength bit rates beyond what can be easily achieved using conventional on/off keying. We review spectrally efficient links based on direct detection, Stokes vector detection, coherent detection and differentially coherent detection for data center applications. We show that limited spectral efficiency and power margin will inhibit scaling of direct detection-compatible formats beyond 100 Gbit/s. Stokes vector receivers can provide higher spectral efficiency without requiring a local oscillator laser, but require power-hungry analog-to-digital converters (ADCs) and digital signal processing (DSP). Similarly, existing DSP-based coherent systems designed for long-haul transmission may be excessively complex and power-hungry for short-reach data center links. We present low-power DSP-free coherent and differentially coherent alternatives that avoid high-speed ADCs and DSP and achieve similar performance to their DSP-based counterparts in intra-data center links and dispersion-compensated inter-data center links.

Keywords: Data center interconnect, direct detection, Stokes vector receiver, coherent detection, differentially coherent detection.

1. INTRODUCTION

Global data center internet protocol (IP) traffic is estimated to grow at a compound annual rate of 26.8% from 2015 to 2020, corresponding to a three-fold increase in five years.[1] Data center-to-data center IP traffic is expected to grow at an even faster rate of 31.9%.[1] This poses a significant challenge to continuously scaling the capacity of data center links while meeting strict constraints in cost and power consumption.

Scaling the capacity of data center links has long relied on using multiple wavelengths or multiple fibers to carry conventional on-off keying (OOK) signals. This strategy cannot scale much further, however, as 400 Gbit/s links, for instance, would require 16 lanes of 25 Gbit/s, resulting in prohibitively high cost, complexity, and power consumption. Recent research has focused on spectrally efficient modulation formats compatible with cost-effective intensity modulation and direct detection (IM-DD)[2–5] to minimize power consumption. These efforts led to the adoption of four-level pulse amplitude modulation (4-PAM) by the IEEE 802.3bs task force to enable 50 and 100 Gbit/s per wavelength. Nevertheless, 4-PAM systems already face tight optical signal-to-noise ratio (OSNR) and power margin constraints in amplified and unamplified systems, respectively. Moreover, next-generation interconnects will likely need to accommodate increased optical losses due to fiber plant, wavelength demultiplexing of more channels, and possibly optical switches. To alleviate some of these constraints, both mature and emerging technologies can help on a number of fronts. High-bandwidth, low-power modulators[6] will reduce intersymbol interference (ISI) and improve signal integrity. Segmented modulators[7] may simplify the transmitter-side electronics. Avalanche photodiodes (APD) and semiconductor optical amplifiers (SOA) can improve receiver sensitivity of 100 Gbit/s 4-PAM systems by 4.5 and 6 dB,[8] respectively. Improved laser frequency stability, either using athermal lasers[9] or frequency combs,[10] will enable dense wavelength-division multiplexing (DWDM) within the data center, possibly yielding a multi-fold increase in capacity.
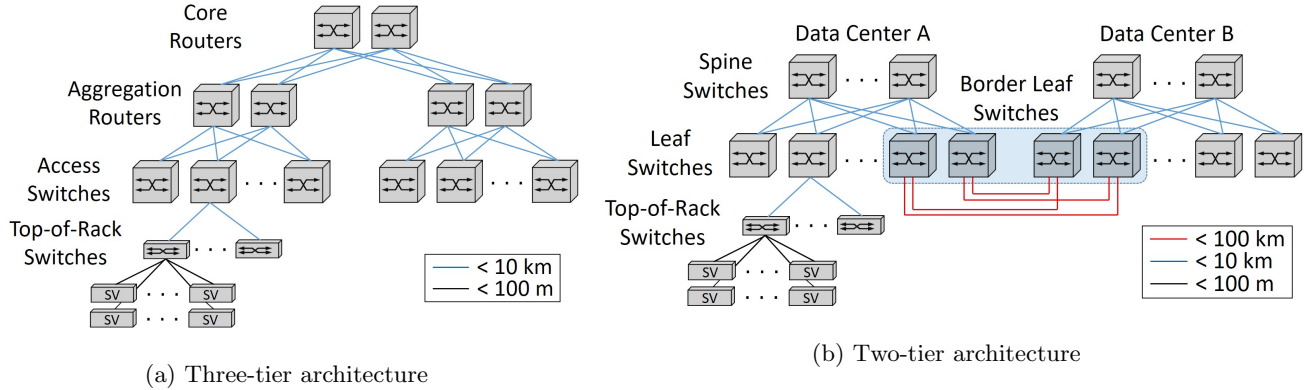
Figure 1: (a) A traditional three-tier data center architecture. (b) A newer two-tier data center architecture with every leaf switch connected to every spine switch. Intra-data center links are shown in blue and black, while inter-data center links are shown in red.

These technologies will extend the lifetime of 4-PAM, but they do not address the fundamental problem of such IM-DD systems, which is that they only exploit one degree of freedom of optical signals, namely, their intensity. Stokes vector detection has been proposed to enable up to three independent dimensions,[11] while avoiding a local oscillator (LO) laser and coherent detection. Nonetheless, Stokes vector receivers rely on power-hungry analog-to-digital converters (ADCs) and digital signal processing (DSP) and do not address the problem of high required OSNR in amplified links or poor receiver sensitivity in unamplified links. Coherent detection allows four degrees of freedom, namely two quadratures in two polarizations, and significantly improve receiver sensitivity due to the LO laser gain. Coherent receivers based on analog signal processing[12] are particularly promising architectures because of their low power consumption, as they avoid high-speed ADCs and DSP. DSP-based coherent receivers may also become attractive in the future, as demand for even higher spectral efficiency increases, and as those systems are optimized for low-power, short-reach applications by leveraging more advanced complementary metal-oxide semiconductor (CMOS) integrated circuit processes. The high spectral efficiency enabled by coherent detection, combined with its improved receiver sensitivity, will potentially blur distinctions between intra- and inter-data center links.

In this paper, we review and compare these different detection techniques and their enabling technologies. In Section 2, we start by reviewing data center networks and important characteristics of intra- and inter-data center links. In Section 3, we review recent research on modulators, in particular electro-optic Mach-Zehnder modulators (MZMs). In Section 4, we discuss optical fiber requirements. In Section 5, we discuss direct detection (DD)-compatible techniques including M-PAM and orthogonal frequency-division multiplexing (OFDM), also commonly referred to as discrete multitone (DMT). We present comparative results in terms of receiver sensitivity and required OSNR. In Section 6, we review Stokes vector receivers that allow utilization of more than one degree of freedom of the optical channel. In Section 7, we review digital and analog coherent receivers, as well as differentially coherent receivers. In Section 8, we compare the different modulation formats and detection techniques according to their overall complexity and DSP power consumption. In Section 9, we conclude the paper.

2. DATA CENTER NETWORKS

2.1 Network Architectures

Evolving traffic patterns due to virtualization and cloud computing have led to shifts from north-south traffic, i.e., traffic from outside data centers to servers, to east-west traffic, i.e., traffic from servers to other servers within the same data center or another one nearby.

Fig. 1a shows a traditional data center architecture consisting of three tiers. In this scheme, servers connect to access switches that then connect to two aggregation routers for redundancy. These aggregation routers are then connected to core routers with redundancy. While this is an efficient structure to manage north-south

Table 1: Impairments and constraints for intra- and inter-data center links.

Link Type	Reach (km)	Wavelength (nm)	Wavelength spacing	Main impairments	Amp.	Priorities
Intra-data center	≤ 10	1310	LAN WDM, CWDM	Pol. rotation	No	Power consumption, power margin, bit rate
Inter-data center	≤ 100	1550	DWDM	Pol. rotation, CD	Yes	Bit rate, power consumption
Long-haul	≤ 1000s	1550	DWDM	PMD, CD, Nonlinearities	Yes	Bit rate, reach

traffic, it is inefficient for east-west traffic. Traffic from one server to another in the same data center may travel up to the core layer and then back down, traversing two access switches, two aggregation routers, and a core router.

Hyperscale data centers have shifted to a flatter architecture consisting of two tiers,[13] as shown in Fig. 1b. In this configuration, servers are connected to leaf switches or to top-of-rack (TOR) switches that are connected to leaf switches, which in turn are connected to every spine switch, resulting in a multitude of paths. East-west traffic must now only travel to a spine switch before traveling back down to the desired leaf switch, resulting in low and predictable latency. Expanding the network is readily done by adding more leaf switches or spine switches, as needed. Fault tolerance is also improved, as a single spine switch failing will only result in a marginal decrease in performance. Achieving the full connectivity of the leaf-spine architecture does require more transceivers, as every leaf switch is connected to every spine switch.

Interconnection between nearby (< 100 km) data centers is achieved by interconnecting their border leaf switches, as illustrated in Fig. 1b. These inter-data center links have different constraints and impairments than the intra-data center links used within data centers.

2.2 Intra- and Inter-Data Center Links

Table 1 summarizes the different constraints and impairments of intra- and inter-data center, in contrast with long-haul systems. In long-haul systems, the high cost and power consumption of complex designs are amortized, as a 3-dB improvement in receiver sensitivity may double the reach and nearly halve the number of required repeaters. Intra- and inter-data center links, however, have other design priorities such as cost, power consumption, and port density, and they face fewer propagation impairments, as polarization mode dispersion (PMD) and nonlinearities are typically negligible over these short propagation distances.

Fig. 2a shows an example system diagram for an intra-data center link. The transceivers in these links can use multiple wavelengths to achieve high bit rates, but they are typically multiplexed and demultiplexed within the module. Intra-data center links reach up to 10 km and typically operate near 1310 nm to minimize total chromatic dispersion (CD). In this small-CD regime, receiver-side electronic equalization is effective, as shown in the performance curves of Section 5. Moreover, intra-data center links are typically unamplified, resulting in low power margin. APDs and SOAs may improve the receiver sensitivity, as shown in[8] and discussed in Section 5. Current intra-data center links employ either coarse wavelength-division multiplexing (CWDM) with wavelength spacing of 20 nm, or LAN-WDM with wavelength spacing of 4.5 nm to avoid power-hungry laser temperature control. Dense WDM (DWDM) may become commercially viable by leveraging advances in athermal lasers and frequency combs.

Fig. 2b shows an example system diagram for an inter-data center link. Inter-data center links reach up to 100 km and operate near 1550 nm to leverage erbium-doped fiber amplifiers (EDFAs). CD is significant and must be compensated. As CD is a nonlinear operation in IM-DD systems, simple receiver-side electronic equalization is not effective. Alternatively, CD may be compensated optically by dispersion-shifted fibers (DCFs) or tunable fiber Bragg gratings (FBGs),[14] depicted in Fig. 2b by the block CD^{-1}. Though they are less flexible than electronic equalization, they are more power-efficient.
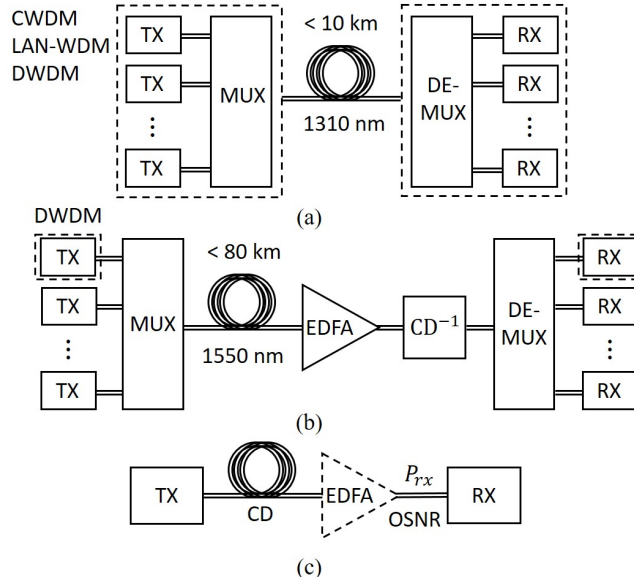
Figure 2: System-level diagrams of (a) intra-data center links, (b) inter-data center links, and (c) equivalent model for intra- and inter-data center links. The components enclosed in dashed lines in (a) and (b) may be encompassed by a single module. The fiber in (c) corresponds to total dispersion in intra-data center links, and residual dispersion after optical CD compensation in inter-data center links. The optical amplifier in (c) is the equivalent link amplifier, which is omitted when modeling intra-data center links.

Fig. 2c shows the equivalent model for intra or inter-data center links. In intra-data center links, the fiber represents the total link dispersion and the optical amplifier is omitted. The performance metric is the receiver sensitivity, which is the received power P_{rx} necessary to achieve a certain target BER, determined by the FEC threshold. For all scenarios, we consider the conventional hard-decision Reed-Solomon code RS(255, 239), which has a net coding gain of 5.6 dB at 10^{-12} BER, an input BER threshold of 1.8×10^{-4} to achieve 10^{-12} BER, and overhead of $\sim 7\%$. Note that the FEC choice is not critical for the performance comparison, since all schemes would benefit from a stronger FEC code.

In inter-data center links, the fiber in Fig. 2c corresponds to the residual CD after optical CD compensation. The EDFA corresponds to the equivalent optical amplifier, whose noise figure depends on the number of amplifiers in the link and on their individual noise figures. The performance metric is the OSNR necessary to achieve the target BER. The equivalent model of Fig. 2c allows treatment of intra and inter-data center links in the same framework. The difference lies only on whether the link is amplified or not. Hence, the same model could be used for intra-data center links amplified with SOAs.

We disregard fiber nonlinearities in the results presented in the next sections, since data center links are short and operate with relatively small power levels. Polarization effects are only studied for Stokes vector receivers and coherent receivers, discussed in Sections 6 and 7, respectively.

3. MODULATORS

Although present intra-data center transceivers typically use cost-effective directly modulated lasers (DMLs) or electro-absorption modulators (EAMs), future transceivers will likely shift to MZMs, which are already used in inter-data center transceivers, due to negligible chirp, high bandwidth and the ability to modulate both quadratures of the electric field. In this section, we restrict our focus to MZMs, but DMLs and EAMs are covered in detail in.[15,16] Table 2 shows properties of MZMs implemented in three different materials: lithium niobate ($LiNbO_3$), indium phosphide (InP), and silicon (Si). While $LiNbO_3$ modulators have long been used in long-haul systems, their size makes them unsuitable for the integration necessary for compact, low-power transceivers needed for intra- and inter-data center links. InP and Si, on the other hand, are small and are

Table 2: Comparison of electro-optic (EO) modulator material properties. Integration refers to integration with other components.

Material	Native laser	EO effect	$V_\pi L_\pi$ (V·cm)	Insertion loss	Thermal sensitivity	Integration	Comments
LiNbO$_3$	No	Inherent	10	Very low	Very low	No	Discrete components
InP	Yes	Inherent	0.5-1	Low	High	Yes	Small wafers (75 mm)
Si	No	Requires doping	0.2-2	Medium	Very low	Yes	Large wafers (300 mm)
SOH	No	Induced	0.05-0.2	Medium	High	Yes	Uncertain lifespan
Thin-film LiNbO$_3$	No	Inherent	2	Very low	Very low	Yes	

compatible with high degrees of integration.[17, 18] One of the primary drawbacks of Si compared to InP is the lack of a native laser. Si modulators are typically uncooled, owing to the small temperature sensitivity of the free-carrier plasma dispersion effect in Si. From a manufacturing standpoint, Si wafers are much larger and have higher yields than InP wafers. The ability to leverage the mature CMOS process is also a significant advantage for Si, as well as its being an inherently less expensive material. Manufacturing in Si can still be expensive, however, if not done at sufficient scale. High insertion loss is also a potential issue for many Si modulator implementations. Si and InP have their own advantages and disadvantages, but both are viable options for MZMs used in coherent intra- and inter-data center links.

Recent modulator research has focused on heterogeneous integration of other materials on Si as a way to leverage the various advantages of Si while overcoming its shortfalls, notably its potential bandwidth limitations at high symbol rates and its lack of a native laser, as noted above. While InP lasers are commonly used with Si modulators, they are external to the silicon photonic chip, resulting in high alignment costs or high insertion loss. Heterogeneous integration of the InP laser on Si allows for adiabatic coupling from the InP gain region to Si waveguides[19] and, in the future, may lead to lower insertion loss, better integration and cheaper manufacturing. Silicon-organic hybrid modulators are promising candidates for low-power coherent intra- and inter-data center links. By inducing strong electro-optic effects in organic materials placed in the cladding of the modulator waveguide, they offer short modulators with very low drive voltages,[6, 20] However, as the induced electro-optic effect can deteriorate within months, further research may be needed to ensure their viability. Integrating LiNbO$_3$ onto silicon can exploit its high-bandwidth electro-optic effect, resulting in wideband modulators[21] that still benefit from silicon's high ability of integration. Thin-film single-crystalline LiNbO$_3$ modulators have been demonstrated with small size, wide bandwidth, and low insertion loss, making them promising candidates for future data center links.[22]

Higher-order PAM and quadrature-amplitude modulation (QAM) signals may be generated using segmented MZMs.[23, 24] Segmented MZMs preclude the need for digital-to-analog converters (DACs) and linear drivers at the transmitter by allowing binary signaling and limiting drivers to control different segments of the MZM to generate the desired output signal. As segmented MZMs would require several extra segments to support pulse shaping and preemphasis, they are not as desirable for long-haul systems, but could result in significant power and complexity savings for data center links.

4. OPTICAL FIBER

The unique requirements of data center links may also motivate reevaluation of optical fiber CD characteristics. When power consumption is the primary concern, fibers with small CD or optical CD compensation should be preferred, since electronic compensation will inevitably be more power-hungry.

In intra-data center links, the limited power margin of DD links may not support, in the long term, increased optical losses due to the fiber plant, wavelength demultiplexing of more channels, and possibly optical switches. As a result, it may become convenient to operate intra-data center links near 1550 nm in order to leverage EDFAs. In this scenario, dispersion shifted fibers (DSF) with zero-dispersion wavelength near 1550 nm may

be preferred. Note that nonlinear fiber effects, which can be exacerbated by DSF, are negligible in intra-data center links, since they are short (up to a few km) and operate with relatively small power levels due to eye safety constraints. The DSF CD slope near 1550 nm should be small in order to maximize the number of WDM channels supported. Dispersion-flattened optical fibers with zero-dispersion wavelengths near both 1310 nm and 1550 nm bands would allow operability of intra-data center links in both bands.

Amplified inter-data center links would also benefit from small-CD fibers, as these links are limited by the amount of CD that can be compensated electronically or simply tolerated without compensation. As shown in Section 5, at bit rates of 100 Gbit/s, IM-DD systems can compensate for up to 80 ps/nm of dispersion using simple linear equalizers, while DSP-free coherent systems (Section 7) systems can tolerate up to 40 ps/nm without any compensation. As a result, these systems must employ some form of optical CD compensation, either using DCF or tunable FBG. Alternatively, non-zero DSF (NZ-DSF) with 1 or 2 ps/(nm×km) could eliminate the need of optical CD compensation for intermediate-reach inter-data center links, while allowing enough CD to minimize nonlinear effects.

5. DIRECT DETECTION

This section describes systems based on direct detection of the electric field. This includes M-PAM, OFDM, and single-sideband (SSB) modulation either of single-carrier or multicarrier formats. This section also covers the so-called Kramers-Kronig receiver, which reconstructs the phase of the received electric field from the received intensity waveform.

5.1 Pulse-Amplitude Modulation

In M-PAM, the information is encoded in M intensity levels. At the receiver, the optical signal is direct detected, filtered, and converted to the digital domain where adaptive equalization is performed. The equalizer may be a simple feedforward equalizer (FFE) or a decision-feedback equalizer (DFE). Alternatively, the receiver may perform maximum likelihood sequence detection (MLSD). Provided that CD is small, the IM-DD channel is accurately modeled as a linear channel. In this regime, an FFE exhibited only a 1-dB penalty with respect to the optimal and more complex MLSD.[2] For large CD, the fiber IM-DD channel is no longer approximately linear, and FFE or DFE are less effective.

The performance of an M-PAM system is determined by the noise variance at each intensity level. There are three scenarios of interest. The first consists of short-reach links in which the receiver uses a positive-intrinsic-negative (PIN) photodiode and thermal noise is dominant. In the next scenario, APD-based receivers have higher sensitivity, but shot noise becomes significant and will affect the noise variance at each level differently. Lastly, in amplified systems with either SOAs or EDFAs, the signal-amplified spontaneous emission (ASE) beat noise is dominant, resulting in different noise variances at the different intensity levels. Although the signal-ASE beat noise is not Gaussian, it can be accurately approximated as Gaussian, as systems with FEC operate at relatively high error rates. Analytical BER calculation for each of these scenarios is described in.[25]

When the noise is signal dependent, as in APD-based or optically amplified links, using equally spaced levels and decision thresholds is suboptimal. The intensity levels and the decision thresholds should be appropriately optimized to minimize the BER. While the exact optimization is intractable, nearly optimal performance is achieved by setting the intensity levels such that all error events are equally likely. This heuristics leads to an analytical equation for the intensity levels when the noise is Gaussian.[8] When the noise is not Gaussian, the intensity levels can be computed numerically using the Karhunen-Loève series expansion.[8]

Fig. 3a shows the receiver sensitivity of 4-PAM and other modulation formats at 100 Gbit/s vs. dispersion for unamplified systems based either on PIN photodiode or APD. Fig. 3b shows the required OSNR in amplified systems with either an SOA or an EDFA. The other modulation formats are discussed in the following subsections. The dispersion axis may be interpreted as total CD in intra-data center links, or residual CD after optical CD compensation in inter-data center links. These results were obtained following the analytical derivations presented in,[25] which are typically within 2 dB of the Monte Carlo simulations. Simulation parameters are given in [25, Table 3].

4-PAM outperforms all other candidates in all considered scenarios. The APD-based receiver has nearly 4 dB better sensitivity than the PIN-based receiver. The APD gain is optimized at each point following the procedure in,[8] and is approximately $G_{APD} \approx 12$. Level spacing optimization improves the receiver sensitivity by roughly 1 dB for an APD-based receiver (Fig. 3a), while in amplified systems (Fig. 3b), it results in $\sim$ 3-dB OSNR improvement.

As observed in Fig. 3, after roughly 50 ps/nm of dispersion, the penalty due to CD increases steeply. This penalty poses a limit in the reach of intra-data center links and restricts the maximum residual dispersion after optical CD compensation in inter-data center links. Several techniques have been proposed to extend the uncompensated reach. Perhaps the most effective is electronic pre-compensation,[26, 27] whereby the transmitted signal is filtered by the inverse of the fiber frequency response $H_{CD}^{-1}(f) = \exp(0.5j\beta_2(2\pi f)^2 L)$, where $\beta_2 = -(\lambda/2\pi c)D$, D is the dispersion parameter, and λ is the transmission wavelength. In theory, this pre-filtering can compensate any amount of CD, which must be known at the transmitter. The main drawback of this approach is its complexity. In addition to the filters having tens of taps for 100 km, the transmitter requires two DACs and drivers, as well as an in-phase and quadrature (I&Q) modulator.

Vestigial-sideband (VSB) and SSB modulation allows uncompensated transmission of 4-PAM over 80 km.[28] The signal resulting from the mixing of VSB or SSB signal with a strong carrier does not experience power fading. SSB modulation has generally better performance than VSB modulation, at the cost of more complex transmitter. In VSB, the intensity-modulated 4-PAM is generated as usual, and the negative sideband is suppressed by an optical filter. The transmitter laser and the optical filter must have fine wavelength stabilization in order to ensure filtering of the correct signal band. SSB modulation and signal-signal beat interference (SSBI) cancelation are discussed in Section 5.3 for OFDM, but the same considerations apply to single-carrier formats. A detailed comparison of SSB and VSB for OFDM is presented in.[29]

Line coding techniques such as duobinary 4-PAM[30] or Tomlinson-Harashima precoding[31] are less effective, since they do not avoid the power fading due to CD in IM-DD channels. Duobinary 4-PAM encoding and Tomlinson-Harashima precoding narrow the signal bandwidth, but even if the bandwidth is halved, the maximum tolerable dispersion is only on the order of 300 ps/nm.[25] In,[32] Tomlinson-Harashima precoding was used in conjunction with DFEs to compensate CD-induced power fading. Tomlinson-Harashima precoding mitigates the error propagation penalty in DFEs, which allows transmission of 56 Gbaud 4-PAM up to 50 km using soft-decision FEC and OSNR required of 38 dB.

5.2 Orthogonal Frequency-Division Multiplexing or Discrete Multitone

In OFDM, the information is encoded on narrowband and orthogonal subcarriers. In data center literature, OFDM is commonly referred to as DMT, which is terminology borrowed from wireline communications literature, where DMT is often used to describe an OFDM signal transmitted at baseband.

OFDM, in principle, offers higher spectral efficiency than 4-PAM, since the individual subcarriers can be modulated using higher-order QAM. Two variants of OFDM were originally proposed for intensity-modulated data center links: DC-biased OFDM (DC-OFDM) and asymmetrically clipped optical (ACO)-OFDM. These OFDM variants differ in how they meet the non-negativity constraint of the intensity-modulated optical channel, and they achieve different tradeoffs between power efficiency and spectral efficiency. In DC-OFDM, a relatively high DC bias is added to minimize clipping distortion. By contrast, in ACO-OFDM, the entire negative excursion of the signal is clipped, and clipping distortion is avoided by encoding information only on the odd subcarriers.[33]

Fig. 3a shows the performance of both DC- and ACO-OFDM in terms of receiver sensitivity vs. dispersion for unamplified systems, and Fig. 3b shows their performance for amplified systems in terms of OSNR required. 4-PAM outperforms both OFDM variants. DC-OFDM has a significant penalty due to the relatively high DC bias required to meet the non-negativity constraint of the intensity-modulated optical channel. Although ACO-OFDM has better performance, it requires prohibitively high DAC/ADC sampling rates due to its low spectral efficiency. In fact, the ACO-OFDM performance curves are not monotonic because, as dispersion increases, subcarriers near the first notch of the IM-DD channel frequency response achieve a poor SNR, so are not used.

Similarly to 4-PAM, CD mitigation through linear equalization is only effective when CD is small. Bit loading and power allocation would allow OFDM variants to better exploit the power-faded optical channel resulting
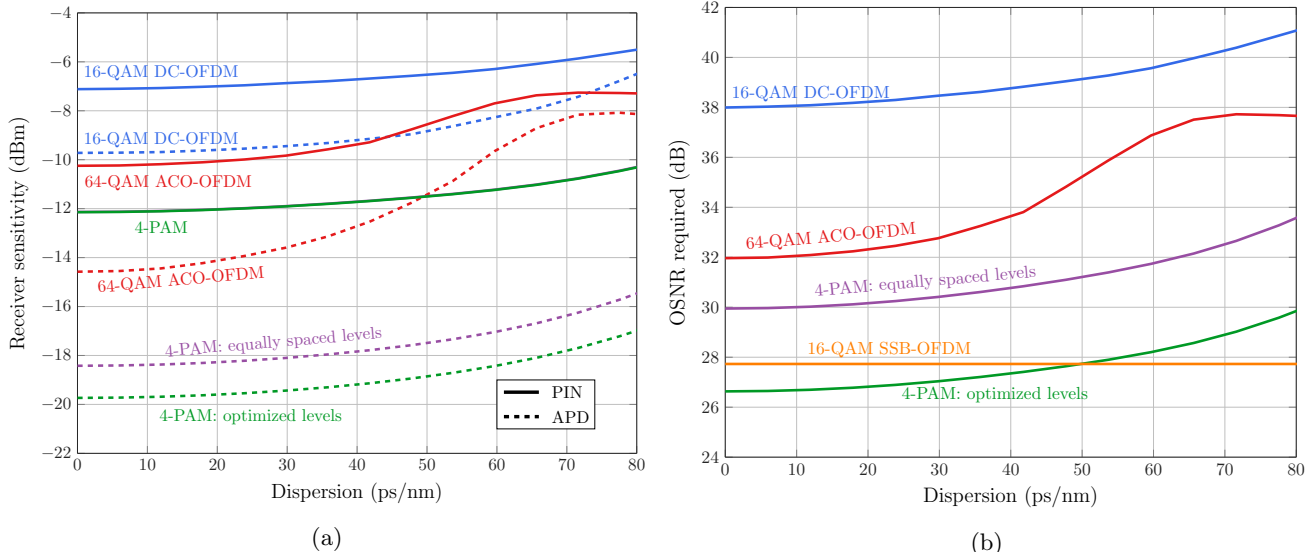
Figure 3: Comparison of performance of DD-compatible modulation schemes vs chromatic dispersion at 112 Gbit/s. Unamplified systems based on PIN photodiodes or APD are characterized in terms receiver sensitivity 3a, while amplified systems are characterized in terms of OSNR required 3b. The x-axis may be interpreted as total dispersion in intra-data center links or residual dispersion after optical CD compensation in inter-data center links (Fig. 2c).

from considerable CD, but such systems are unlikely to be practical, since DD also leads to intermodulation products that fall in the signal band. OFDM can also exploit pre-equalization techniques at the cost of double complexity at the transmitter by using two DACs, two drivers, and an I&Q modulator. More recently, however, SSB-OFDM has been proposed as a viable form of allowing uncompensated transmission for inter-data center links.

5.3 Single-Sideband OFDM

In SSB-OFDM, the subcarriers corresponding to the negative sideband are not modulated. The SSB-OFDM signal can still be directly detected, provided that a sufficiently strong unmodulated optical carrier is also transmitted. After DD, the mixing of the unmodulated carrier and the SSB-OFDM signal yields a real-valued double-sideband (DSB)-OFDM signal carrying the same information as the original SSB-OFDM signal. This DSB-OFDM signal does not experience the power fading characteristic of the IM-DD channel. In fact, the DSB-OFDM signal only experiences phase distortion, which can be effectively compensated by electronic equalization.

Fig. 3b shows the required OSNR for a SSB-OFDM assuming perfect SSBI cancellation. Since the detected DSB-OFDM does not experience power fading, the required OSNR does not vary with dispersion because. The ~ 28-dB OSNR required is similar to the OSNR required using Kramers-Kronig technique.[34]

Several SSBI cancellation techniques have been proposed with different efficacies and complexities. In,[35] SSBI cancellation is performed by using the received signal $y[k]$ to estimate the SSBI term by computing $|y[k] + j\mathcal{H}\{y[k]\}|^2$ and subtracting it from the received signal. A similar procedure is proposed in,[34] where the interference estimate is computed by linearization of the receiver. Due to noise, these techniques are most effective at high OSNR. Another technique is based on non-linear equalization based on truncated Volterra series.[36] The number of taps N_{taps} in a Volterra non-linear equalizer grows rapidly, and a simple time-domain implementation has complexity $\mathcal{O}(N_{taps}^2)$. In,[36] the Volterra nonlinear equalizer had 28 taps.

Another SSBI cancellation technique proposed in[34] is based on the so-called Kramers-Kronig (KK) receiver.[37,38] In contrast to previous techniques, the KK receiver reconstructs the phase of the electrical field from the detected intensity waveform. This reconstruction is only possible if the electric field signal is minimum phase, which is guaranteed by transmitting a sufficiently strong carrier.[37]

The KK phase retrieval technique is not restricted to SSB-OFDM signals. In fact, the KK phase retrieval technique was utilized to reconstruct a SSB 4-PAM signal,[38] and to reconstruct a M-QAM signal.[37] Note that for

Table 3: Experimental results of various Stokes vector receivers.

# of dims	Modulation	SE (bits/s/Hz)	Bit rate (Gbit/s)	Reach (km)	Wavelength (nm)	Ref.
2	DP-4-PAM	4	224	10	1310	[40]
2	16-QAM + Carrier	4	224	320	1550	[41]
3	DP-4-PAM + 4-PM	6	300	0	1550	[11]
4	DP-2-PAM + 8-PM + 8-DPM	8	320	10	1550	[39]

QAM, the information on the negative frequency sideband is not redundant. Hence, the transmitted signal must be frequency-shifted by $R_s/2$ with respect to the carrier, where R_s is the signal rate. Consequently, the spectral efficiency of KK M-QAM is halved: $0.5 \log_2 M$, which is the same spectrum efficiency achieved by $\sqrt{M}$-PAM modulation. Moreover, this is the same spectral efficiency achieved by carrierless amplitude and phase (CAP) modulation[2] without the SSB requirement and additional complexity of KK phase retrieval. However, CAP does not allow electronic CD compensation. For these reasons, the so-called KK receiver does not improve spectral efficiency or receiver sensitivity.

The KK phase retrieval does permit electronic CD compensation, but at arguably higher DSP complexity than the techniques described previously. The logarithm and square root computations require high-precision arithmetic as well as upsampling by a large factor in order to correctly represent $\ln \sqrt{y[k]}$ in the frequency domain. In,[37] an upsampling factor of three was recommended.

6. STOKES VECTOR DETECTION

The DD-compatible modulation formats presented in the previous section have the common drawback that they only exploit one dimension of the optical channel; i.e., optical intensity. Stokes vector detection was proposed to leverage more degrees of freedom of the optical channel while still using DD. In Stokes vector detection, the electric field components $|E_X|^2$, $|E_Y|^2$, $\mathrm{Re}\{E_X E_Y^*\}$, and $\mathrm{Im}\{E_X E_Y^*\}$ are detected, and subsequent DSP is realized to recover the transmitted information. Here, E_X and E_Y denote the electric field in the X and Y polarizations, respectively. This allows up to three independent degrees of freedom of the optical channel; i.e., intensity in X and Y polarizations, and inter-polarization phase. As shown by Morsy-Osman $et\ al$,[39] four-dimensional modulation may be achieved by also encoding information on the inter-polarization differential phase, but the receiver has to be significantly modified, increasing the power penalty due to power splitting. Table 3 summarizes recently published experimental results of Stokes vector receivers.

Although Stokes vector receivers achieve higher spectral efficiency by levering more than one degree of freedom of the optical channel, the required OSNR is similar to DD systems. For instance, a 112 Gbit/s DP-4-PAM + 4-PM achieves the target BER of 1.8×10^{-4} at OSNR of approximately 27 dB,[25] which is comparable to SSB-OFDM and 4-PAM with optimized level spacing (Fig. 3b). Moreover, the OSNR required varies with the received state of polarization. Fortunately, the OSNR required varies less than 1 dB for random states of polarization.[25]

CD compensation capability is also small, since the signal received by a Stokes vector receiver also experiences power fading characteristic of IM-DD channels. The exception for this is the M-QAM+Carrier format (Table 3), which, similarly to VSB and SSB modulation, relies on the mixing with a strong unmodulated carrier to produce a detected signal that does not experience power fading.

Stokes vector receivers do not require an LO laser, but the remaining receiver circuitry is comparable to that of a coherent receiver. Moreover, as discussed in Section 8, the receiver DSP power consumption is comparable to that of DP-16-QAM coherent receivers optimized for short-reach applications, but without the same receiver sensitivity or CD compensation capability.

7. COHERENT AND DIFFERENTIALLY COHERENT DETECTION

This section reviews DSP-based and DSP-free coherent and differentially coherent receiver architectures. We show that DSP-free coherent and differentially coherent receivers have substantially smaller power consumption

while achieving similar performance to their DSP-based counterparts. As a drawback, it is difficult to scale DSP-free coherent receivers to higher-order (M-ary, $M > 4$) QAM formats, even though the SNR in short-reach links would support these formats.

7.1 DSP-Based Coherent Receiver (DP-M-QAM)

Coherent detection based on high-speed DSP is a mature technology in long-haul systems, but it may be currently unsuitable for data center links, where cost and power consumption are paramount.

In a conventional dual-polarization DSP-based coherent receiver, the DSP chip performs functions such as polarization demultiplexing, PMD compensation, CD compensation, carrier recovery and clock recovery. Some implementations place the DSP chip on the line card itself with an analog interface to the pluggable transceivers, referred to as analog coherent optics (ACO). While this can increase transceiver port density, it essentially offloads the power consumption to elsewhere in the system.

The power consumption of the various operations performed by the receiver was extensively studied in.[42] The most power-hungry operations are CD equalization and polarization demultiplexing with PMD compensation, which together amount to roughly 55% of the receiver power consumption.[42] CD equalization is performed using nearly static frequency-domain equalizers with hundreds of taps. The subsequent 2×2 MIMO equalizer for polarization demultiplexing and PMD compensation is comprised of filters with typically less than 15 taps that are updated frequently to mitigate PMD and track changes in the received state of polarization.

Recent works have focus on reducing number of operations necessary in CD and PMD equalization. For instance, Martins $et\ al$[43] have proposed a distributive finite-impulse response (FIR) equalizer that leverages the high multiplicity of the quantized FIR filter coefficients to drastically reduce the number of required operations. Compared to a conventional frequency-domain CD equalizer, their distributive FIR equalizer requires 99% fewer multiplications and 30% fewer additions.[43] As another example, in [12], we showed that the number of filters in the 2×2 MIMO equalizer can be halved provided that the skew between the two polarizations is much smaller than the sampling rate. This simplification only holds when the mean differential group delay (DGD) between the two polarizations is much smaller than the sampling rate, so that the two polarizations appear synchronized at the receiver. Assuming a sampling rate of 70 GS/s (oversampling ratio of 5/4 at 56 Gbaud), and PMD of 0.1 ps/$\sqrt{\text{km}}$, the small-DGD approximation holds up to ~ 200 km.

Fig. 4 shows the performance of various coherent and differentially coherent systems as a function of dispersion at 224 Gbit/s. These results were obtained following the analytical derivations presented in,[25] for simulation parameters given in [25, Table 5]. The curves in Fig. 4 for DSP-based receivers are flat across dispersion values, as CD is effectively compensated by electronic equalization. DSP-based coherent detection systems can use higher-order modulation, such as 16-QAM, to reduce the bandwidth required of electro-optic components. The other curves in Fig. 4 are discussed in the following subsections.

7.2 DSP-Free Coherent Receiver (DP-QPSK)

Coherent detection using analog signal processing was studied extensively in the 1980s and early 1990s,[44] but the advent of the EDFA and later DSP-based coherent detection diminished its popularity.

Fig. 5 shows the proposed implementation of a DSP-free coherent receiver.[12] Polarization demultiplexing is performed by optical phase shifters that are controlled by low-speed circuitry. Other receiver operations such as carrier recovery, timing recovery and detection are performed in the high-speed analog electronics stage. Timing recovery and detection may be realized using conventional clock and data recovery (CDR) techniques;[45] thus, we do not discuss them further herein.

The polarization controller, shown by the inset in Fig. 5, must recover the transmitted state of polarization by inverting the fiber polarization transfer matrix. Three cascaded phase shifter pairs can perform any arbitrary polarization rotation.[46] One way to control the amounts of differential phase shifts of each phase shifter is to transmit a marker tone in one quadrature of one polarization and minimize its presence in the other quadrature and polarization at the receiver. We have demonstrated compensation of up to 700 rad/s of polarization rotation in Monte Carlo simulations.[12] Faster tracking speeds can be obtained with larger changes of angles in the phase shifters, resulting in slightly higher average polarization compensation error, a tradeoff that must be optimized.
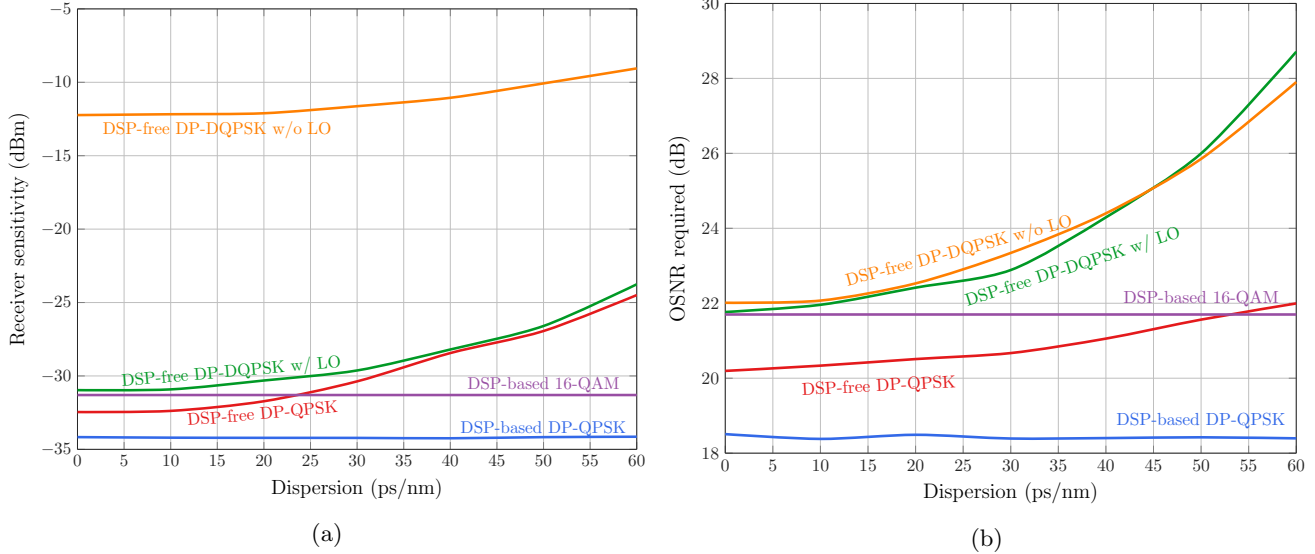
Figure 4: Comparison of performance of coherent detection schemes vs. dispersion at 224 Gbit/s. Unamplified systems are characterized in terms of (a) receiver sensitivity, while amplified systems are characterized in terms of (b) OSNR required. The x-axis may be interpreted as total dispersion in intra-data center links or residual dispersion after optical CD compensation in inter-data center links (Fig. 2c).
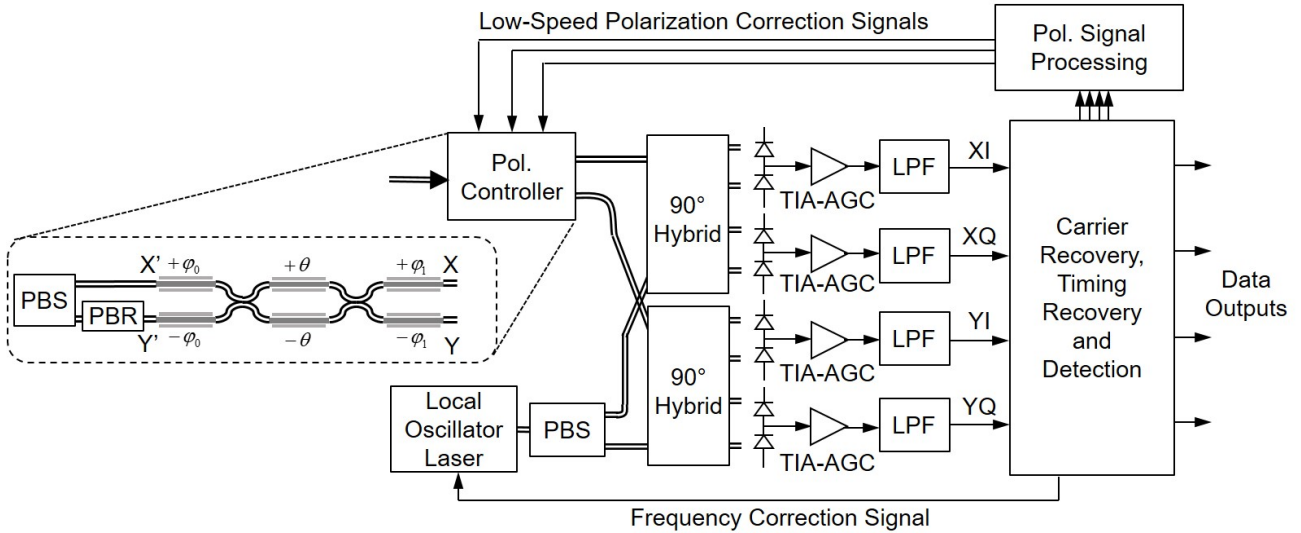

Figure 5: Block diagram of coherent receiver architectures based on analog signal processing. Acronyms: local oscillator (LO), polarization beam splitter (PBS), polarization beam rotator (PBR), transimpedance amplifier (TIA), automatic gain control (AGC), analog-to-digital converter (ADC), low-pass filter (LPF).

This method can be used for QPSK, as well as for 16-QAM and higher-order formats. Implementation of the phase shifters can be done using silica, LiNbO$_3$ or any other material that has low loss and allows integration of a sequence of phase shifters. The phase shifters and waveguides do not need to necessarily support two polarizations, as the demultiplexing is performed through phase shifts and coupling, not birefringence. Endless polarization control can be achieved by cascading more phase shifting sections or using a material such as LiNbO$_3$ for phase shifters and resetting them fast enough to allow interleaving and FEC to correct the burst errors. As polarization rotation through a fiber is a relatively slow (milliseconds), time-varying effect,[47] the polarization signal processing in Fig. 5 can be implemented with low-speed electronic microcontrollers.

Carrier recovery is based on a phase-locked loop (PLL). The high-speed analog electronics stage is detailed
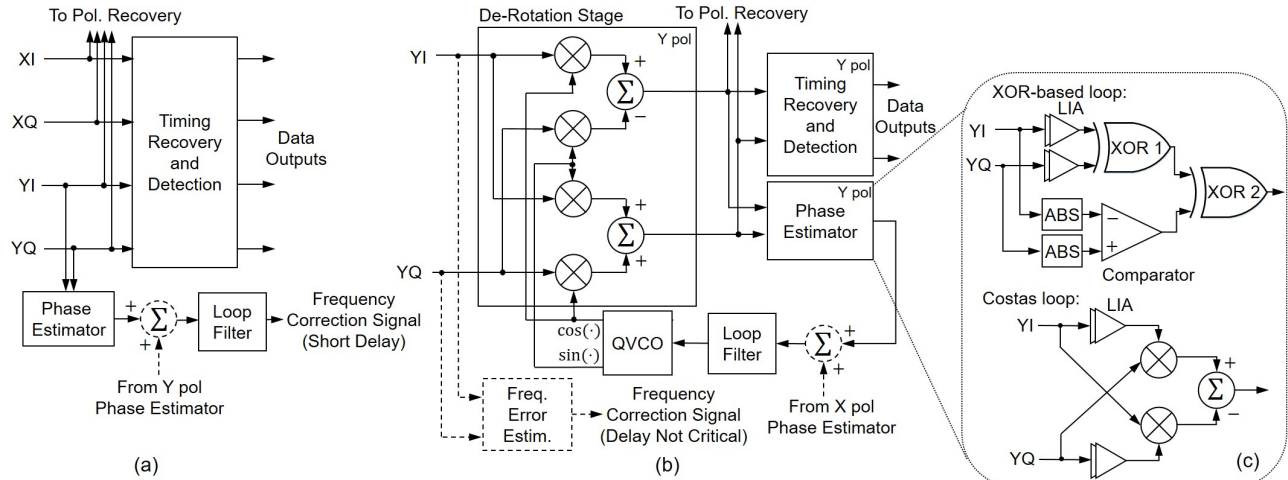
Figure 6: Block diagrams of carrier recovery for an analog coherent receiver based on analog (a) OPLL and (b) EPLL (shown for one polarization only). The phase estimator block is detailed in (c), where LIA denotes limiting amplifiers, and ABS denotes full-wave rectifiers.

in Fig. 6 for carrier recovery based on optical PLL (OPLL) and electrical PLL (EPLL). In an OPLL (Fig. 6a), the LO laser is frequency-modulated by the frequency correction signal generated by the CR stage. Hence, an OPLL requires a LO laser with wideband frequency modulation (FM) response and short propagation delay in the LO path to minimize the overall loop delay. Minimizing the loop delay is one of the main challenges in OPLL design, since the loop includes the LO laser, 90° hybrid, photodiodes, and all the subsequent electronics in carrier recovery, which may not be realized within the same chip. Notably, Park *et al* have demonstrated loop delays of only 120 ps for a highly integrated 40 Gbit/s binary PSK coherent receiver.[48] An EPLL (Fig. 6b) implementation eliminates requirements on LO laser FM response and on propagation delay at the cost of more complex analog electronics. Specifically, an EPLL requires a SSB mixer in each polarization to de-rotate the incoming signals (see Fig. 6b), since the transmitter and LO lasers are not phase locked. Additionally, the frequency offset between the transmitter and LO lasers must always be within the lock-in and hold-in ranges of the EPLL, which are practically limited by the voltage-controlled oscillator (VCO) frequency range (typically up to 10 GHz). This constraint can be satisfied by strict laser temperature control, whose cost and power consumption could be shared among several channels by using frequency combs for both the transmitter and LO. Alternatively, a frequency error estimation stage (Fig. 6b), based on a relatively simple frequency discriminator circuit,[49] may be used to keep the LO laser frequency sufficiently close to the transmitter laser.

For both OPLLs and EPLLs, the phase estimator for the loop filter shown in Fig. 6b can be implemented using analog multipliers, as exemplified by the Costas loop phase estimator. We propose a multiplier-free phase estimator (Fig. 6c), which relaxes signal integrity requirements. Our proposed XOR-based phase estimator operates by estimating the sign of the phase error instead of its amplitude. The receiver performance for either Costas or XOR-based loop depends on the phase error variance, which can be estimated using a small-signal approximation to linearize the PLL.[12]

Fig. 4 shows the performance curves for an analog coherent receiver. At small dispersion, the penalty with respect to the DSP-based receiver is due to imperfect receiver filtering. In our simulations, the LPF is a fifth-order Bessel filter with bandwidth of 39.2 GHz ($0.7R_s$ for 224 Gbit/s DP-QPSK), for which $\Delta f = 40.7$ GHz. Hence, the imperfect receiver filtering results in a 1.6 dB penalty compared to DSP-based receiver. As dispersion increases, the receiver sensitivity decreases or OSNR required increases sharply, since the receiver does not equalize CD. Nonetheless, the sensitivity would allow unamplified eye-safe systems near 1310 nm to achieve a reach up to 40 km. In fact, systems with 100 GHz wavelength spacing could support 49 channels with 5 dB of margin, and systems with 200 GHz wavelength spacing could support 25 channels with 8 dB of margin.

7.3 Differentially coherent detection (DP-DQPSK)

Differentially coherent detection is performed by computing the phase difference between two consecutive symbols. This precludes the need of an absolute phase reference, and hence carrier phase recovery is not necessary. Differential detection, however, has some disadvantages compared to coherent detection. First, for the same spectral efficiency, differential detection has an inherent SNR penalty, e.g., $\sim$ 2.4 dB for DQPSK compared to QPSK.[50] Second, differential detection restricts modulation to PSK formats.

Differential detection may be performed in the electrical domain (with LO) or in the optical domain (without a LO). In electronic differential detection, the phase difference between two consecutive symbols is computed using analog electronics.[25] DSP-based differential detection is less appealing because carrier recovery necessary for coherent detection only accounts for 4.4% of the receiver power consumption. In optical differential detection using delay interferometers, the receiver electronics must only perform timing recovery. This configuration does not employ an LO laser, which simplifies the receiver significantly. The delay caused by the delay interferometer is sensitive to the wavelength. As a result, the transmitter laser's frequency drifts can cause a penalty if not properly compensated by tuning the delay interferometer.[51] For DP-DQPSK, at 224 Gbit/s without delay interferometer tuning, a frequency drift of $\pm$800 MHz would incur a 2-dB penalty.

Fig. 4 shows the performance of homodyne DP-DQPSK for electronic (with LO) and optical (without a LO) differential detection. The curves of Fig. 4 assume no frequency error. In unamplified systems, the LO laser provides sufficient gain that electronic differential detection incurs only a small penalty relative to coherent detection. DQPSK without an LO has significantly poorer receiver sensitivity in unamplified systems, such as intra-data center links. However, the OSNR required in amplified systems remains approximately the same as that of a LO-based DQPSK receiver. This makes LO-free DQPSK an attractive option for amplified inter-data center links that have optical CD compensation, as they have the lowest receiver complexity among coherent and differentially coherent receivers.

8. COMPLEXITY COMPARISON

The previous sections compared the performance of the various modulation formats and detection techniques in terms of receiver sensitivity and OSNR required. This section focuses on the overall complexity and power consumption of these schemes.

Fig. 7 shows a coarse estimate of power consumption in 28-nm CMOS for various modulation schemes at 100 Gbit/s and 200 Gbit/s. The DSP-free receiver power consumption is estimated at 90-nm CMOS as detailed in.[12] The power consumption of DSP-based techniques is estimated using the power consumption models presented in.[42] First, the number of real additions and real multiplications is counted for all DSP operations (see[25] for details). Then, the power consumption is obtained by computing how much energy a given operation consumes. For instance, a real addition in 28-nm CMOS with 6-bit precision consumes 0.28 pJ, while a real multiplication with 6-bit precision consumes 1.66 pJ.[42] The power consumption estimates for DACs and ADCs assume that the power consumption scales linearly with resolution and sampling rate. The DAC figure of merit is 1.56 pJ/conv-step, while the ADC figure of merit is 2.5 pJ/conv-step.[42] Only OFDM formats are assumed to need high-resolution DACs, since single-carrier formats may avoid them assuming that pulse shaping and preemphasis are not performed. For all cases, the oversampling ratio assumed is $r_{os} = 5/4$, even though Stokes vector receivers and KK receivers have only been reported with $r_{os} = 2$.

Fig. 7(a and b) compare the power consumption of DD-compatible schemes at 100 Gbit/s for (a) a CD-compensated link where the residual CD is at most 80 ps/nm, and for (b) an 80-km uncompensated CD link. As expected, 4-PAM is more power efficient than the other formats. Compared to OFDM schemes, 4-PAM benefits from requiring lower sampling frequency, lower resolution, and performing time-domain equalization, which is more power efficient than frequency-domain equalization for short filters. However, in the high-uncompensated-CD regime, SSB modulation is the only viable choice. The SSBI cancellation in SSB-OFDM is assumed to be a Volterra nonlinear equalizer with 14 taps in (a) and 28 taps in (b). The power consumption of KK 4-PAM is excessively high due to the phase estimation using 3-times upsampling for computation of the Hilbert transform, as discussed in Section 5.3. Although not shown in Fig. 7, the power consumption of 4-PAM with MLSD in an uncompensated link would also be excessively high, since the complexity of the MLSD receiver
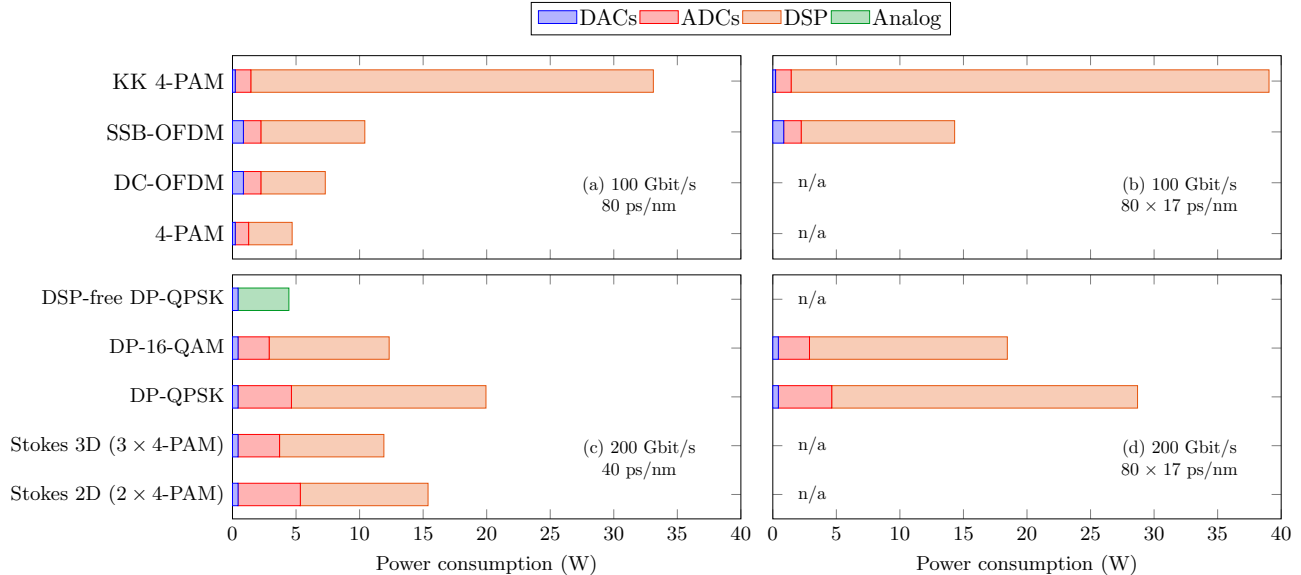
Figure 7: Coarse estimate of power consumption of high-speed DACs, ADCs, and DSP for various modulation schemes at 100 Gbit/s and 200 Gbit/s. DSP power consumption estimates are made for 28-nm CMOS using the models presented in.[42] DSP-free receiver power consumption is estimated for 90-nm CMOS.[12] The graphs to the left assume that CD is compensated optically and the residual CD is at most 80 ps/nm at 100 Gbit/s (a) and 40 ps/nm at 200 Gbit/s (c). The graphs to the right assume uncompensated transmission up to 80 km near 1550 nm. In this regime, most techniques cannot work due to the high uncompensated CD.

grows exponentially with the memory length of the Viterbi decoder. We do not include MLSD 4-PAM in the comparison of Fig. 7 due to the lack of models to translate branch metric computations into power consumption.

Fig. 7(c and d) compare schemes with higher degrees of freedom at 200 Gbit/s for (c) a CD-compensated link where the residual CD is at most 40 ps/nm, and for (d) an 80-km uncompensated link. DSP-free coherent is more power efficient as it avoids high-speed ADCs and DSP, which comes at the expense of small tolerance to CD. In the small residual CD regime (Fig. 7c), DSP-based coherent receivers have similar power consumption to that of Stokes vector receiver. The LO laser in coherent receivers provides improved receiver sensitivity, and it may account for up to 2.5 W of the total receiver power consumption.[42] In the high uncompensated CD regime (Fig. 7d), DSP-based coherent is the only viable option. The results of Fig. 7(c and d) also illustrate that it is more power efficient to operate with higher constellation sizes and more degrees of freedom in order to minimize the symbol rate.

9. CONCLUSION

Increase in traffic within data centers, as well as between data centers, will demand higher per-wavelength bit rates. DD-compatible formats can meet these needs in the short-term, but more degrees of freedom are needed to support higher per-wavelength bit rates. Stokes vector receivers allow more degrees of freedom, but rely on power-hungry ADCs and DSP. Coherent and differentially coherent detection methods enable up to four degrees of freedom while significantly improving receiver sensitivity. However, conventional DSP-based coherent receivers designed for long-haul transmission, which prioritizes performance, are suboptimal for data center applications, which prioritize cost and power consumption. By reducing receiver complexity and making system performance tradeoffs, the power consumption of coherent links can be made low enough for intra- and inter-data center applications. Following this philosophy, LO-based DSP-free coherent receivers seem particularly promising for intra-data and inter-data center links, whereas amplified inter-data center links could also support LO-free differentially coherent receivers.

REFERENCES

[1] Cisco, "Cisco Global Cloud Index: Forecast and Methodology, 2015–2020," *White Paper*, pp. 1 – 41, 2016.

[2] M. Sharif, J. K. Perin, and J. M. Kahn, "Modulation Schemes for Single-Laser 100 Gb/s Links: Single-Carrier," *J. Lightw. Technol.* **33**(20), pp. 4268–4277, 2015.

[3] J. K. Perin, M. Sharif, and J. M. Kahn, "Modulation Schemes for Single-Wavelength 100 Gbit/s Links: Multicarrier," *J. Lightw. Technol.* **33**(24), pp. 5122–5132, 2015.

[4] J. L. Wei, J. D. Ingham, D. G. Cunningham, R. V. Penty, and I. H. White, "Comparisons between 28 Gb/s NRZ, PAM, CAP and optical OFDM systems for datacommunication applications," *2012 Optical Interconnects Conference* **2**, pp. 3–4, may 2012.

[5] K. Zhong, X. Zhou, T. Gui, L. Tao, Y. Gao, W. Chen, J. Man, L. Zeng, A. P. T. Lau, and C. Lu, "Experimental study of PAM-4, CAP-16, and DMT for 100 Gb/s Short Reach Optical Transmission Systems," *Optics Express* **23**(2), pp. 1176–1189, 2015.

[6] M. Lauermann, R. Palmer, S. Koeber, P. C. Schindler, D. Korn, T. Wahlbrink, J. Bolten, M. Waldow, D. L. Elder, L. R. Dalton, J. Leuthold, W. Freude, and C. Koos, "Low-power silicon-organic hybrid (SOH) modulators for advanced modulation formats," *Optics Express* **22**(24), pp. 29927–29936, 2014.

[7] D. Patel, A. Samani, V. Veerasubramanian, S. Ghosh, and D. V. Plant, "Silicon Photonic Segmented Modulator-Based Electro-Optic DAC for 100 Gb/s PAM-4 Generation," *IEEE Photonics Technology Letters* **27**(23), pp. 2433–2436, 2015.

[8] J. Krause Perin, M. Sharif, and J. M. Kahn, "Sensitivity Improvement in 100 Gbit/s-per- Wavelength Links using Semiconductor Optical Amplifiers or Avalanche Photodiodes," *Journal of Lightwave Technology* **34**(33), pp. 5542–5553, 2016.

[9] S. H. Lee, A. Wonfor, R. V. Penty, I. H. White, G. Busico, R. Cush, and M. Wale, "Athermal colourless C-band optical transmitter for passive optical networks," *European Conference on Optical Communication, ECOC* **1-2**(August 2015), 2010.

[10] T. N. Huynh, R. Watts, V. Vujicic, M. D. G. Pascual, C. Calo, K. Merghem, V. Panapakkam, F. Lelarge, A. Martinez, B. E. Benkelfat, A. Ramdane, and L. P. Barry, "200-Gb/s Baudrate-Pilot-Aided QPSK/Direct Detection with Single-Section Quantum-Well Mode-Locked Laser," *IEEE Photonics Journal* **8**(2), 2016.

[11] M. Chagnon, M. Morsy-Osman, D. Patel, V. Veerasubramanian, A. Samani, and D. Plant, "Digital signal processing for dual-polarization intensity and interpolarization phase modulation formats using stokes detection," *Journal of Lightwave Technology* **34**(1), pp. 188–195, 2016.

[12] J. Krause Perin, A. Shastri, and J. Kahn, "Design of Low-Power DSP-Free Coherent Receivers for Data Center Links," *Journal of Lightwave Technology* **35**(21), pp. 4650–4662, 2017.

[13] R. Urata, H. Liu, X. Zhou, and A. Vahdat, "Datacenter Interconnect and Networking : from Evolution to Holistic Revolution," in *OFC*, **W3G.1**, 2017.

[14] M. Sumetsky and B. Eggleton, "Fiber bragg gratings for dispersion compensation in optical communication systems," *Journal of Optical and Fiber Communications Reports* **2**, pp. 256–278, Sep 2005.

[15] U. Troppenz, J. Kreissl, M. Möhrle, C. Bornholdt, W. Rehbein, B. Sartorius, I. Woods, and M. Schell, "40 Gbit/s Directly Modulated Lasers: Physics and Application," *Proceedings of SPIE* **7953**, pp. 79530F–79530F–10, 2011.

[16] N.-N. Feng, D. Feng, S. Liao, X. Wang, P. Dong, H. Liang, C.-C. Kung, W. Qian, J. Fong, R. Shafiiha, Y. Luo, J. Cunningham, A. V. Krishnamoorthy, and M. Asghari, "30GHz Ge electro-absorption modulator integrated with 3 μm silicon-on-insulator waveguide.," *Optics express* **19**(8), pp. 7062–7, 2011.

[17] C. R. Doerr, "Silicon photonic integration in telecommunications," *Frontiers in Physics* **3**(August), pp. 1–16, 2015.

[18] D. Thomson, A. Zilkie, J. E. Bowers, O. Alibart, V. D. Auria, M. D. Micheli, M. Smit, X. Leijtens, H. Ambrosius, and E. Bente, "An introduction to InP-based generic integration technology," *Semicond. Sci. Technol.* , pp. 1–41, 2014.

[19] G.-h. Duan, C. Jany, A. L. Liepvre, A. Accard, M. Lamponi, D. Make, P. Kaspar, G. Levaufre, N. Girard, J.-m. Fedeli, A. Descos, B. B. Bakir, S. Messaoudene, D. Bordel, S. Menezo, G. D. Valicourt, S. Keyvaninia, G. Roelkens, D. V. Thourhout, D. J. Thomson, F. Y. Gardes, and G. T. Reed, "Hybrid III − V on Silicon Lasers for Photonic Integrated Circuits on Silicon," *IEEE Journal of Selected Topics in Quantum Electronics* **20**(4), 2014.

[20] S. Koeber, R. Palmer, M. Lauermann, W. Heni, D. L. Elder, D. Korn, M. Woessner, L. Alloatti, S. Koenig, P. C. Schindler, H. Yu, W. Bogaerts, L. R. Dalton, W. Freude, J. Leuthold, and C. Koos, "Femtojoule electro-optic modulation using a silicon-organic hybrid device," *Light: Science & Applications* **4**(2), p. e255, 2015.

[21] A. Mercante, P. Yao, S. Shi, G. Schneider, J. Murakowski, and D. Prather, "110 GHz CMOS compatible thin film LiNbO3 modulator on silicon," *Optics Express* **24**(14), pp. 15590–15595, 2016.

[22] C. Wang, M. Zhang, B. Stern, M. Lipson, and M. Loncar, "Nanophotonic lithium niobate electro-optic modulators," *arXiv preprint arXiv:1701.06470* , 2017.

[23] A. Shastri, C. Muzio, M. Webster, G. Jeans, P. Metz, S. Sunder, and B. Chattin, "Ultra-Low-Power Single-Polarization QAM-16 Generation Without DAC Using a CMOS Photonics Based Segmented Modulator," *Journal of Lightwave Technology* **33**(6), pp. 1255–1260, 2015.

[24] A. Aimone, I. G. Lopez, S. Alreesh, P. Rito, T. Brast, V. Hohns, G. Fiolt, and M. Grunert, "DAC-free Ultra-Low-Power Dual-Polarization 64-QAM Transmission with InP IQ Segmented MZM Module," in *OFC*, p. Th5C.6, 2016.

[25] J. Krause Perin, A. Shastri, and J. M. Kahn, "Data center links beyond 100 Gbit/s per wavelength," *Optical Fiber Tech* , pp. 1–28, 2018.

[26] J. McNicol, M. O'Sullivan, K. Roberts, A. Comeau, D. McGhan, and L. Strawczynski, "Electrical domain compensation of optical dispersion," in *OFC*, **OThJ3**, 2005.

[27] Q. Zhang, N. Stojanovic, X. Changsong, C. Prodaniuc, and P. Laskowski, "Transmission of single lane 128 Gbit/s PAM-4 signals over an 80 km SSMF link, enabled by DDMZM aided dispersion pre-compensation," *Optics express* **24**(21), pp. 714–716, 2016.

[28] J. Lee, N. Kaneda, and Y. K. Chen, "112-gbit/s intensity-modulated direct-detect vestigial-sideband pam4 transmission over an 80-km ssmf link," in *ECOC 2016; 42nd European Conference on Optical Communication*, pp. 1–3, Sept 2016.

[29] D. Pilori, C. Fludger, and R. Gaudino, "Comparing DMT variants in medium-reach 100G optically amplified systems," *Journal of Lightwave Technology* **34**(14), pp. 3389–3399, 2016.

[30] H. Keangpo and J. Kahn, "Multilevel Optical Signals Optmized for Systems Having Signal-Dependent Noises, Finite Transmitter Extinction Ratio and Intersymbol Interference," *US Patent* **2**(12), 2004.

[31] K. Matsumoto, Y. Yoshida, A. Maruta, A. Kanno, N. Yamamoto, and K. i. Kitayama, "On the impact of tomlinson-harashima precoding in optical pam transmissions for intra-dcn communication," in *2017 Optical Fiber Communications Conference and Exhibition (OFC)*, pp. 1–3, March 2017.

[32] R. Rath, D. Clausen, S. Ohlendorf, S. Pachnicke, and W. Rosenkranz, "Tomlinson-Harashima Precoding for Dispersion Uncompensated PAM-4 Transmission with Direct-Detection," *Journal of Lightwave Technology* **35**(18), pp. 3909–3917, 2017.

[33] J. Armstrong and A. J. Lowery, "Power efficient optical OFDM," *Electronics Letters* **42**(6), 2006.

[34] Z. Li, S. Erkilinc, K. Shi, E. Sillekens, L. Galdino, B. Thomsen, P. Bayvel, and R. Killey, "SSBI Mitigation and Kramers-Kronig Scheme in Single-Sideband Direct-Detection Transmission with Receiver-based Electronic Dispersion Compensation," *Journal of Lightwave Technology* (to appear), 2017.

[35] S. Randel, D. Pilori, S. Chandrasekhar, G. Raybon, and P. Winzer, "100-Gb/s discrete-multitone transmission over 80-km SSMF using single-sideband modulation with novel interference-cancellation scheme," *European Conference on Optical Communication, ECOC* **2015-Novem**(4), 2015.

[36] L. Zhang, T. Zuo, Y. Mao, Q. Zhang, E. Zhou, G. N. Liu, and X. Xu, "Beyond 100-Gb/s Transmission Over 80-km SMF Using Direct-Detection SSB-DMT at C-Band," *Journal of Lightwave Technology* **34**(2), pp. 723–729, 2016.

[37] M. S. Antonio Mecozzi, Cristian Antonelli, "Kramers-Kronig coherent receiver," *Optica* **3**(11), pp. 1220—-1227, 2016.

[38] C. Antonelli, M. Shtaif, and A. Mecozzi, "Kramers-Kronig PAM transceiver," in *OFC*, 2017.

[39] M. Morsy-Osman, M. Chagnon, and D. V. Plant, "Four-dimensional modulation and stokes direct detection of polarization division multiplexed intensities, inter polarization phase and inter polarization differential phase," *Journal of Lightwave Technology* **34**(7), pp. 1585–1592, 2016.

[40] M. Morsy-Osman, M. Chagnon, M. Poulin, S. Lessard, and D. V. Plant, "224-Gb/s 10-km transmission of PDM PAM-4 at 1.3 um using a single intensity-modulated laser and a direct-detection MIMO DSP-based receiver," *Journal of Lightwave Technology* **33**(7), pp. 1417–1424, 2015.

[41] M. Y. S. Sowailem, T. M. Hoang, M. Chagnon, M. Morsy-Osman, M. Qiu, S. Paquet, C. Paquet, I. Woods, O. Liboiron-Ladouceur, and D. Plant, "100G and 200G single carrier transmission over 2880 and 320 km using an InP IQ modulator and Stokes vector receiver," *Optics Express* **24**(26), pp. 30485–30493, 2016.

[42] B. S. G. Pillai, B. Sedighi, K. Guan, N. P. Anthapadmanabhan, W. Shieh, K. J. Hinton, and R. S. Tucker, "End-to-end energy modeling and analysis of long-haul coherent transmission systems," *Journal of Lightwave Technology* **32**(18), pp. 3093–3111, 2014.

[43] C. Martins, F. Guiomar, S. Amado, R. Ferreira, S. Ziaie, A. Shahpari, A. Teixeira, and A. Pinto, "Distributive FIR-Based Chromatic Dispersion Equalization for Coherent Receivers," *Journal of Lightwave Technology* **PP**(99), pp. 5023–5032, 2016.

[44] J. Kahn, A. Gnauck, J. Veselka, and S. Korotky, "4-Gb/s PSK Homodyne Transmission System Using Phase-Locked Semiconductor Lasers," *IEEE Photonics Technology Letters* **2**(4), pp. 285–287, 1990.

[45] C. F. Liao and S. I. Liu, "40 Gb/s transimpedance-AGC amplifier and CDR circuit for broadband data receivers in 90 nm CMOS," *IEEE Journal of Solid-State Circuits* **43**(3), pp. 642–655, 2008.

[46] R. Noe, H. Heidrich, and D. Hoffmann, "Endless polarization control systems for coherent optics," *Journal of Lightwave Technology* **6**, pp. 1199–1208, Jul 1988.

[47] H. Bulow, W. Baumert, H. Schmuck, F. Mohr, T. Schulz, F. Kuppers, and W. Weiershausen, "Measurement of the maximum speed of PMD fluctuation in installed field fiber," in *OFC/IOOC*, pp. 83–85, 1999.

[48] H.-c. Park, M. Lu, E. Bloch, T. Reed, Z. Griffith, and L. Johansson, "40 Gbit/s Coherent Optical Receiver Using a Costas Loop," *Optics express* **20**(26), p. 7, 2012.

[49] H. R. Rideout, J. S. Seregelyi, S. Paquet, and J. Yao, "Discriminator-aided optical phase-lock loop incorporating a frequency down-conversion module," *IEEE Photonics Technology Letters* **18**(22), pp. 2344–2346, 2006.

[50] G. P. Agrawal, *Fiber-Optic Communication Systems*, Wiley, New York, 2002.

[51] H. Kim and P. J. Winzer, "Robustness to Laser Frequency Offset in Direct-Detection DPSK and DQPSK Systems," *Journal of Lightwave Technology* **21**(9), pp. 1887–1891, 2003.