IEEE Open Journal of
Circuits and Systems

Received 8 May 2022; revised 18 June 2022 and 5 July 2022; accepted 6 July 2022. Date of publication 12 July 2022; date of current version 19 July 2022.

Digital Object Identifier 10.1109/OJCAS.2022.3189550

Design Space Exploration of Single-Lane OFDM-Based Serial Links for High-Speed Wireline Communications

GAIN KIM (Member, IEEE)

Department of Electrical Engineering and Computer Science, Daegu Gyeongbuk Institute of Science and Technology, Daegu 42988, South Korea

This article was recommended by Associate Editor Y. Dong.

CORRESPONDING AUTHOR: G. KIM (e-mail: gain.kim@dgist.ac.kr)

This work was supported by the DGIST Start-Up Fund Program of the Ministry of Science and ICT Under Grant 2022010322.

ABSTRACT The 4-level pulse-amplitude modulation (PAM-4) with an analog-digital converter (ADC)-based receiver (RX) has become the most commonly employed modulation for ultra-high-speed serial links with the data rate above 100 Gb/s. To support the data rate of 200 Gb/s, orthogonal frequency division multiplexing (OFDM) has been studied recently as one of the possible modulation schemes in the next-generation serial links. The OFDM can feature high bandwidth efficiency without increasing the equalization complexity, leading to a reduced maximum signal amplitude attenuation and lower required DAC/ADC conversion rates given sufficient DAC/ADC resolutions such that the BER is not primarily limited by the data converters' resolution. This paper presents system-level modeling results of OFDM-based wireline serial links, with a particular emphasis on the impacts of the fast Fourier transform (FFT) processor's tap count on the serial link performance. The relationship among the cyclic prefix (CP), FFT tap count, and the link bit-error-rate (BER) are thoroughly explained. The analysis explains that the power consumption of a partially-serial FFT processor improves with a larger kernel FFT size, and simulation results show that the BER performance improves with the FFT size where an optimal CP length exists given the FFT size.

INDEX TERMS Serial link, discrete multitone, DMT, orthogonal frequency division multiplexing, OFDM, wireline communications, wireline receiver.

I. INTRODUCTION

THE DEMANDS for ever-higher communication bandwidth for data centers, cloud platforms, and communication applications led the data rate increase of not only wireless communication systems but also wireline transceivers (TRX). While the compound annual growth rate (CAGR) for global fixed broadband traffic between 2018 and 2023 is expected to be 20 % [1], the data rate growths of various wireline TRXs have been achieving about 2× every four years. Meeting this trend, data rates of high-speed wireline transmitter (TX) and receiver (RX) have been increasing from about 10 Gb/s [2]–[4] to 112 Gb/s [5]–[8] in the past two decades.

While the dominating modulation scheme for serial links with data rate up to 56 Gb/s has been non-return-to-zero (NRZ) [9]–[12], 4-level pulse-amplitude modulation (PAM-4) started to be dominating from per-lane data rate of 56 Gb/s to reduce the maximum channel loss that the signal experiences at Nyquist [13]–[17]. Despite the increased circuit complexity and latency coming from the additionally required forward error correction (FEC), recently demonstrated 112 Gb/s TRXs [6]–[8] employ PAM-4 due to its 2× better bandwidth efficiency as compared to that of NRZ, with an analog-to-digital converter (ADC)-based RX front-end with heavy equalization on the digital signal processor (DSP). However, if PAM-4 is employed for next-generation 224 Gb/s links, at least 56 GHz of analog bandwidth with a Baudrate of 112 GBaud will be required, which will be very challenging not only for the analog front-end (AFE) design but also for the RX DSP design due to the long feed-forward

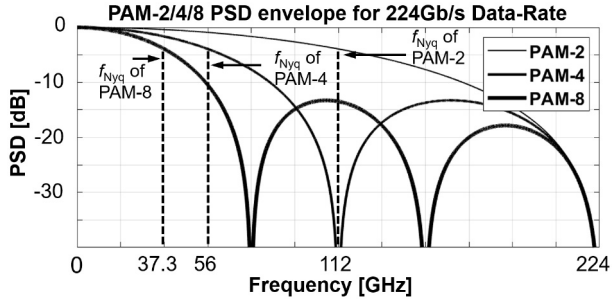

FIGURE 1. Power spectral density (PSD) envelope for randomly-generated PAM-2 (NRZ), PAM-4, and PAM-8 signal with ideal signal transition (i.e., without output bandwidth limitation and jitter) for 224 Gb/s data rate.

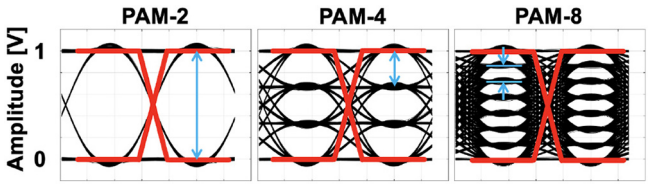

FIGURE 2. Eye diagrams of PAM-2/4/8 with normalized vertical full swing level. Peak-to-peak swings are illustrated in bold red lines and arrows in light blue color indicate the level-to-level amplitude differences in each eye diagram.

equalizer (FFE) taps and tight decision feedback equalizer (DFE) loop closure timing.

To shrink the occupied bandwidth and to reduce the Baudrate of signal with its data rate $>$ 200 Gb/s, looking for a pulse-amplitude modulation (PAM) with higher-order such as 8-level PAM (PAM-8) seems to be a natural way. As can be seen in Fig. 1, the Nyquist frequency (f_{Nyq}) of 224 Gb/s PAM-8 stream is 37.3 GHz while the f_{Nyq} of PAM-4 is 56 GHz. Such 33.3 % of bandwidth reduction relaxes the design complexity of the AFE and the Nyquist sampling ADC in terms of analog bandwidth and the operation speed as compared to those with PAM-4. On the other hand, from signal-to-noise ratio (SNR) perspective, PAM-8 is inherently more sensitive to crosstalk, random jitter of the sampling clock, and residual inter-symbol interference (ISI) on the noise margin as compared to PAM-4, due to the reduced PAM level-to-level distance given a fixed TX output dynamic range, as depicted in Fig. 2. Also, with a faster data rate, the DSP design complexity increases with a more extended channel pulse response that has to be equalized by an FFE with longer tap counts and with more stringent timing constraints for the loop closure of the DFE chain. Moreover, the design complexity of a loop-unrolled DFE, FEC latency, and complexity increase with the PAM order [18], making it not straightforward to move from PAM-4 to PAM-8.

To achieve an improved bandwidth efficiency with a relaxed RX complexity as compared to PAM-4, employing orthogonal frequency division multiplexing (OFDM) modulation to high-speed serial links has been studied. Prior-arts such as [19], [20] showed that 56-112 Gb/s optical data communication is feasible with discrete multi-tone (DMT) modulation (i.e., baseband OFDM with variable bit-loading for each OFDM sub-channel), and [21] demonstrated a 56 Gb/s RX data-path silicon for electrical links. System-level modeling of a DMT serial link for $>$ 100 Gb/s applications is provided in [22], and [23] compared DMT link and PAM-2/4/8 for ultra-high-speed links both with smooth channels and multi-drop-bus (MDB) interfaces. Timing recovery system with adaptive equalization for wireline DMT links is proposed in [24].

While energy- and area-efficient fast Fourier transform (FFT) processor is the key enabler of high-performance OFDM wireline RX data-path design, a thorough study on the impacts of the design parameters of FFT processors on the overall performance of OFDM-based high-speed serial links is missing. This paper presents area estimation of partially-serial FFT processors with a fixed compute throughput specifically for high-speed wireline applications, taking the required RX DSP clock frequency into account for the area-power metric. The effects of applying the DMT modulation to ultra-high-speed serial links are described, and the effects of different FFT tap size on the overall link performance are analyzed with simulations.

The rest of this paper is organized as follows. Section II briefly reviews the basics of the OFDM modulation and the wireline OFDM TRX architecture, and Section III analyzes the design complexity of the FFT processor with various design choices for the wireline OFDM RX data-path. Design considerations and optimization directions with a particular emphasis on the relationship between the FFT size and the link bit-error-rate (BER) performance for $>$ 200 Gb/s OFDM TRX are presented in Section IV. Simulation results are shown in Section V, and Section VI concludes the paper.

II. BACKGROUND

This Section provides the basics of the OFDM modulation and the overall architecture of the TRX that realizes an OFDM serial link.

A. OFDM BASICS

The principle of the OFDM is loading quadrature-amplitude modulated (QAM) symbols to multiple separate sub-channels in the frequency-domain. The left inset of Fig. 3 illustrates an 8-sub-channel OFDM where one 16-QAM symbol is loaded to each sub-channel equally spaced in the frequency-domain. By applying an 8-tap inverse discrete Fourier transform (IDFT) to the eight QAM symbols in the frequency-domain, the time-domain OFDM symbol can be computed, as illustrated in the right inset of Fig. 3. Note that the imaginary (Q) samples are all-zero in wireline applications by making the IDFT input Hermitian symmetry, while these samples are typically non-zero in wireless communications. More details on the signal generation will be shown later in this Section.

Given a total usable bandwidth, the number of OFDM sub-channels defines the bandwidth of each sub-channel. As illustrated in Fig. 4, a larger number of sub-channels makes the sub-channel bandwidth narrower, and it lets each

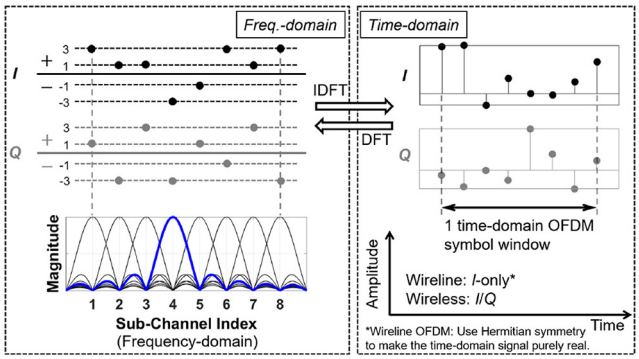

FIGURE 3. A 16-QAM (2 bits in *I*, 2 bits in *Q*) symbol allocation to orthogonal sub-carriers with equal spacing in the frequency-domain (left inset), and the time-domain complex OFDM symbol where its samples are computed by the IDFT operation from the frequency-domain symbols (right inset).

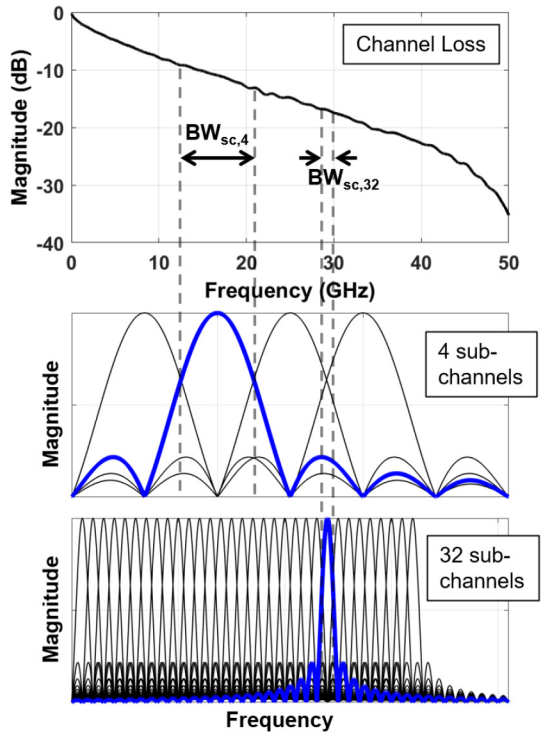

FIGURE 4. An example channel profile and the OFDM PSD plots with two different number of sub-channels (4 and 32).

sub-channel experiences more flat channel response for communicating over a lossy electrical link. In general, with narrower sub-channels, the OFDM signal is less affected by the channel delay spread, but increasing the number of sub-channels requires trade-offs in real-world implementation, which are discussed in later Sections.

B. CYCLIC PREFIX

A guard interval with samples between OFDM symbols, called cyclic prefix (CP), is one of the most important concepts in OFDM. Assume that there is a sequence of time-domain OFDM symbols, where each symbol is indicated as Symbol_k in Fig. 5(a), where $k \in \{1, 2, 3\}$. The time-domain

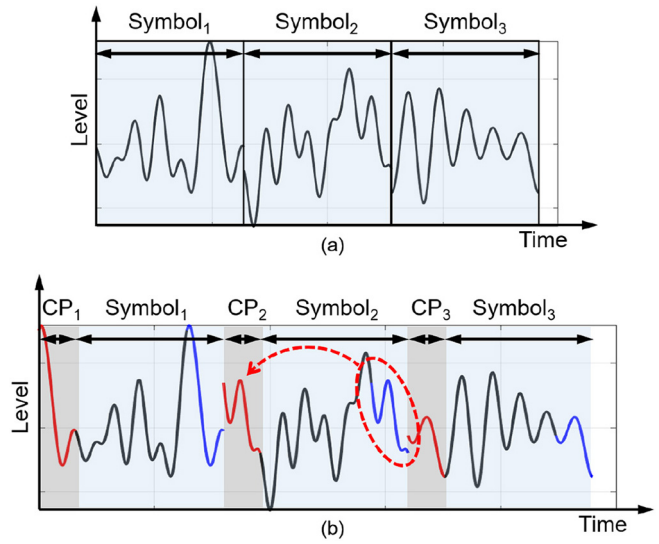

FIGURE 5. A sequence of time-domain OFDM symbols without the CP (a) and with the CP (b).

OFDM symbols will undergo a lossy channel with pre-and post-cursors. Then on the RX side, the current symbol (e.g., Symbol_2 in Fig. 5(a)) will be affected by the ISI coming from the neighboring symbols (e.g., Symbol_1 and Symbol_3 in Fig. 5(a)); hence the current symbol will experience SNR degradation.

To protect the OFDM symbol from ISI, a CP is inserted between the symbols, and these guard samples are a copy of a few last time-domain samples of the corresponding OFDM symbol, as shown in Fig. 5(b). By doing so, the OFDM symbol can be effectively protected from the ISI cursors caused by neighboring symbols. Acting as a guard interval, the CP also allows the sub-channels in an OFDM symbol orthogonal one from another, making it possible to use single-tap equalization for each sub-channel on the RX. More details on the CP for OFDM-based communication systems can be found in [25]. While the number of CP taps should be large enough to cover important ISI cursors, a larger number of CP taps does not guarantee better performance in high-speed wireline applications, and this will be studied more in detail in Section III.

C. OFDM TRX DATA-PATH

Fig. 6 shows the overall block diagram of the OFDM TRX data-path for wireline serial link. On the TX-chain, the FEC-encoded binary bitstream will be mapped to QAM symbols in the frequency-domain, and the bit-loading and/or power-loading (will be explained in Section IV) can be applied depending on the link configuration and the application. Then, the IDFT stage transforms the frequency-domain QAM symbols to the time-domain OFDM symbols, and the CP is inserted for the symbol protection from the ISI. The CP-inserted OFDM symbols will be serialized and transmitted toward the channel by a digital-to-analog converter (DAC). On the RX-chain, the reverse operations of the TX-chain are

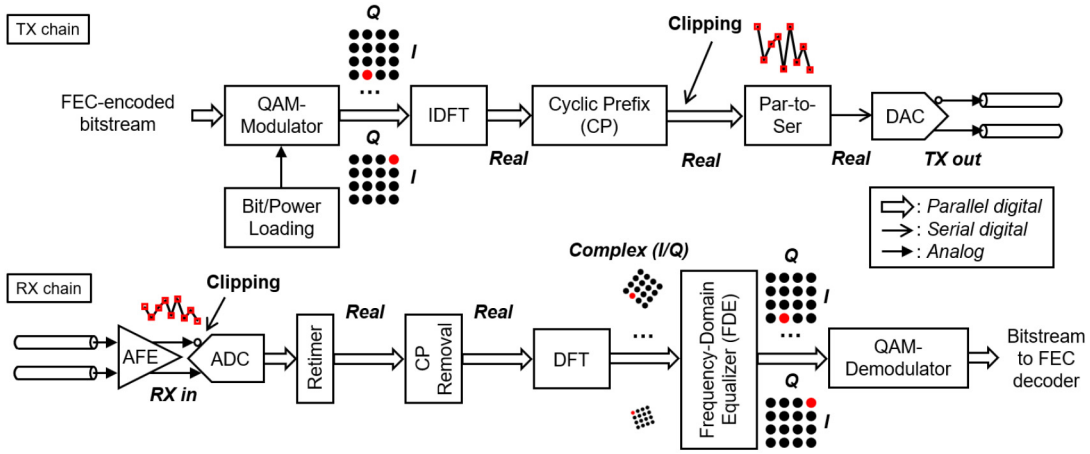

FIGURE 6. A simplified block diagram of the OFDM TRX data-path for a wireline serial link.

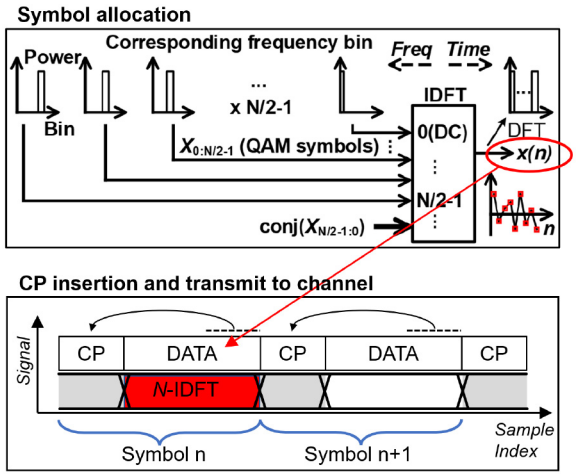

FIGURE 7. The symbol allocation to IDFT input bins in the frequency-domain with Hermitian symmetry to make sure that the IDFT output is purely real-valued (upper inset), and the CP insertion in the time-domain (bottom inset).

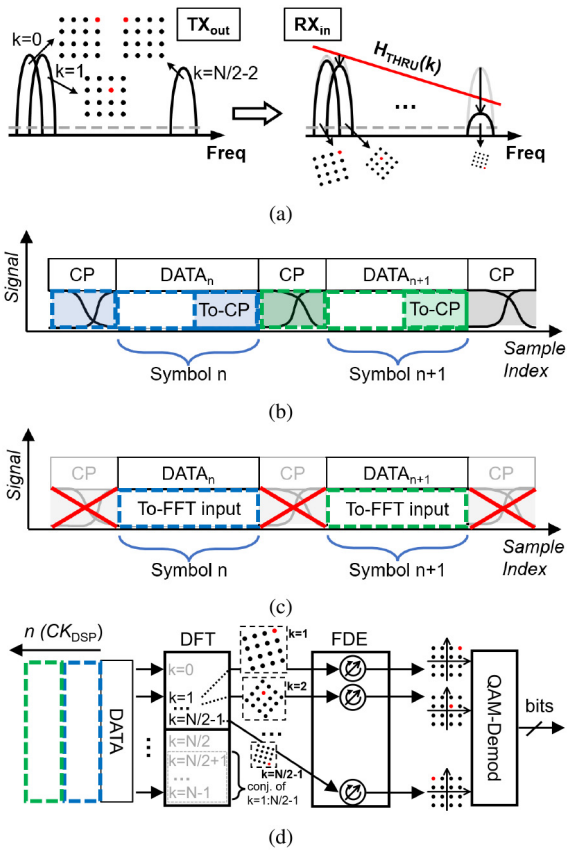

FIGURE 8. The effects of channel on the QAM symbols in OFDM sub-channels (a), the ISI-affected CPs in the time-domain (b), the CP-removal with proper symbol selection (c), and the RX data-path processing with CP-removed samples where DFT outputs corresponding to gray-texted indices are redundant (d).

performed overall, except for the frequency-domain equalization (FDE). As the QAM symbols delivered by different sub-channels experience different amplitude attenuation and angular rotation in the frequency-domain I/Q-plane, those received QAM symbols are rotated-back and amplified with proper complex-valued coefficients individually at each sub-channel data-path. This way of equalization is one of the main reasons why the OFDM is attractive for ultra-high-speed wireline serial links: the equalization can be performed by a single complex multiplication per sub-channel in DSP, substantially simplifying the equalization as compared to the time-domain PAM equalization.

Fig. 7 shows the QAM symbol allocation and the time-domain OFDM symbol generation on the TX DSP. Each arbitrary QAM symbol is allocated to the first half of the IDFT input bins, and the rest half of the bits are filled by the complex conjugate symmetry of the first half. This is to guarantee that the IDFT output is purely real-valued by having Hermitian symmetry at the input of the IDFT.

And each set of the IDFT output will be the data portion of the time-domain OFDM symbol, and the CP is inserted for protecting the OFDM data from the ISI.

On the RX data-path, the QAM symbols require equalization in the frequency-domain due to the frequency-selective loss and angular rotation caused by the channel, as shown in Fig. 8(a). The CP is affected by the ISI of the neighboring

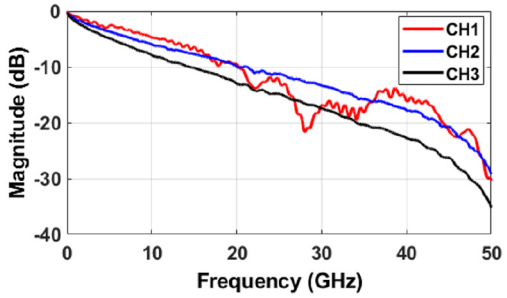

FIGURE 9. Measured frequency responses of point-to-point channels [28].

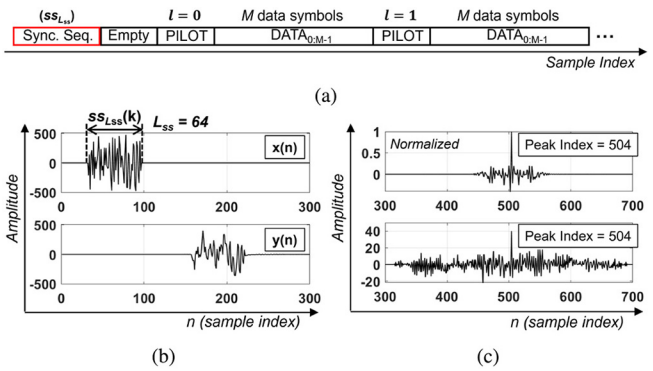

FIGURE 10. The transmit sequence blocks of an OFDM system (a), the transmitted and the received sequences in the time-domain for synchronization (b), and the correlation of $y(n)$ and $ss_{L_{SS}}$ (c) with the standard correlation expressed in (1) (upper inset) and with the computation-efficient method expressed in (2) (bottom inset).

symbols (Fig. 8(b)). The ISI-affected CPs are eliminated by simply ignoring these samples on the RX data-path (Fig. 8(c)), and the data symbols are sent to the DFT processor for the time-to-frequency-domain transform. Note that in Fig. 8(d), the upper half of the DFT output bins contain the same information as the bottom half of the DFT output since the TX IDFT input is Hermitian symmetry. Hence, only half of the DFT output should be connected to the FDE, reducing the complexity of the DFT processor.

D. TX-RX SAMPLE-WISE SYNCHRONIZATION

The TX-RX sample-wise synchronization is necessary for proper CP removal. For that, after the link power-up, a known synchronization sequence $ss_{L_{ss}}$ of length L_{ss} is transmitted through a channel as shown in Fig. 10(a), where *Empty* block, *PILOT* and *DATA* symbols are idle sequence, known symbols for channel estimation (CHEST), and data symbols, respectively. The RX continuously computes the correlation of the signal $y(n)$ coming for the channel (e.g., with CH3 of Fig. 9 for the simulation shown in Fig. 10(b)) with the $ss_{L_{ss}}$ to find the peak of the correlation which will return the synchronization position.

However, computing the correlation of two vectors expressed in (1) in real-time requires excessive DSP resources hence a sign-based correlation expressed in (2) can be a more efficient approach still enabling the TX-RX

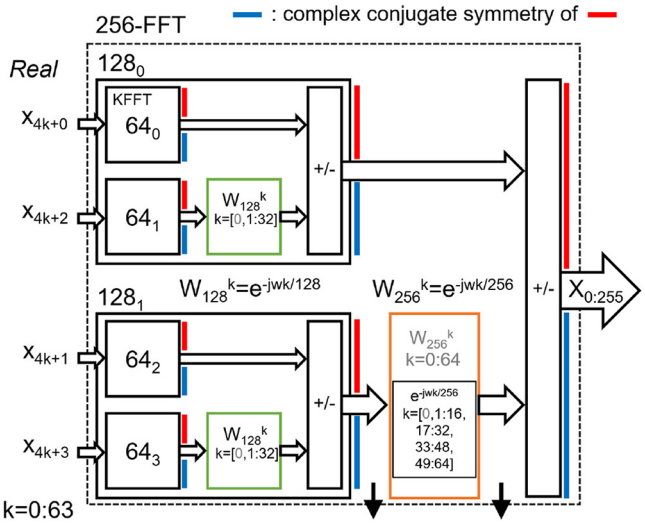

FIGURE 11. Block diagram of an example radix-2 fully-parallel 256-tap FFT processor based on 64-tap FFT kernels.

sample-wise synchronization.

$$z(n) = \sum_{k=0}^{L_{ss}-1} ss_{L_{ss}}(k) \cdot y(n+k) \tag{1}$$

$$z(n) = \sum_{k=0}^{L_{ss}-1} sign(ss_{L_{ss}}(k)) \cdot sign(y(n+k)) \tag{2}$$

While the correlation quality of (2) is worse than (1) as shown in Fig. 10(c), the detection reliability of (2) is robust with sufficiently large L_{ss} (e.g., ≥ 64) for wireline channels exhibiting ≥ 23 dB attenuation at 40 GHz.

III. FFT PROCESSOR ARCHITECTURE FOR WIRELINE OFDM AND ITS DESIGN COMPLEXITY

Among various building blocks of the OFDM RX data-path, the DFT processor and the FDE occupy the majority portion of the area and the power consumption. This section investigates the approximate area and the area-power product metrics of the DFT and the FDE. For area-efficient implementation, the FFT algorithm is widely used in DFT implementation. A parallel N-tap FFT, with $N \in \{4, 8, 16, 32, \ldots\}$, can be realized based on multiple smaller K-tap FFT kernels with a regular radix-based structure. An example of a fully-parallel 256-tap FFT implementation with real-valued inputs based on four 64-tap kernel FFTs (KFFT) is shown in Fig. 11, where the 64_0, 64_1, 64_2, 64_3 are four identical 64-tap KFFTs, green blocks are the first-stage twiddle factor multiplier array for twiddle coefficients $e^{-j\omega k/128}$ with $k \in \{0, 1, \ldots, 32\}$, and the orange block is the second-stage twiddle factor multiplier array for twiddle coefficients $e^{-j\omega k/256}$ with $k \in \{0, 1, \ldots, 64\}$.

The FFT can be designed to operate in partially-serial for area saving at the cost of higher DSP clock (CK_{DSP}) frequency as compared to a fully-parallel architecture given the same target throughput. As shown in Fig. 12, with a

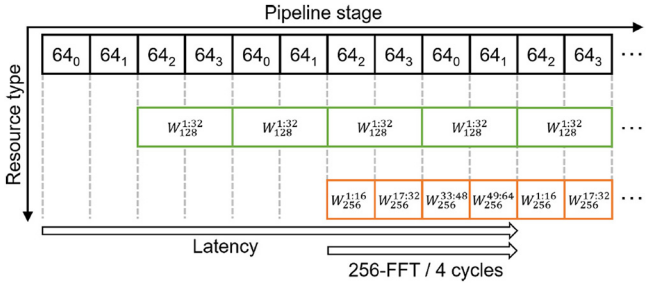

FIGURE 12. The partially-serial pipelined processing flow of a 256-tap FFT based on a 64-tap KFFT.

pipelined architecture, a single 64-tap KFFT can be shared as well as the first-stage and the second-stage twiddle factor multiplier arrays. For a partially-serial operation, a buffer memory should be present at the output of each sub-FFT stage to temporarily store the sub-FFT outputs produced during the pipelined process, as noted at the bottom in Fig. 11. The partially-serial pipelined architecture shown in Fig. 12 can be extended to larger N, such as $N = 512$ or $N = 1024$.

For estimating the complexity of the FFT block in OFDM-based serial link applications, the CK_{DSP} at which the FFT runs has to be taken into account. It is worth noting that the half of the FFT output is redundant due to the Hermitian symmetry and is not used by the FDE. The Hermitian symmetry allows the FFT processor to be designed with half the number of arithmetic resources compared to the full FFT that receives complex input values [26]. In addition, some adders/subtractors and pipeline registers related to the unused output bins will be eliminated during the synthesis process, even if these bins are present in the block diagram, such as in Fig. 8(d). The number of complex multipliers M_{NFFT} in a partially-serial FFT can be written as

$$M_{\mathrm{NFFT}} = M_{\mathrm{KFFT}} + \frac{K}{2} + \frac{K}{4}\left(\log_2 \frac{N}{K} - 1\right), \tag{3}$$

where M_{KFFT} is the number of complex multipliers in a KFFT, K is the number of taps of the KFFT, and $N = 2^l \cdot K$ with $l \in \mathbb{Z}^+$ and $l \geq 2$. While the twiddle factor multipliers arrays except for the first-stage array are with non-constant complex multipliers, the KFFT block and the first-stage twiddle factor multipliers array are with constant multipliers, i.e., multipliers of which one input is a fixed coefficient at design-time. It is worth noting that the circuit area of a constant multiplier is half of a non-constant multiplier given the same bit-width.

Taking this into account, the number of multiplier cells required to implement a partially-serial N-tap FFT is plotted in Fig. 13 for three different KFFT tap sizes. In Fig. 13, the CK_{DSP} in each sub-figure is the required DSP clocking rate for achieving the FFT throughput for 224 Gb/s data rate, with all 64-QAM sub-channels, and 8 for the number of cyclic prefix (N_{CP}) taps. The blue lines show the sum of the number of non-constant complex multipliers in the twiddle factor multipliers arrays (>second-stage). The black lines indicate the sum of the number of constant multipliers in the KFFT

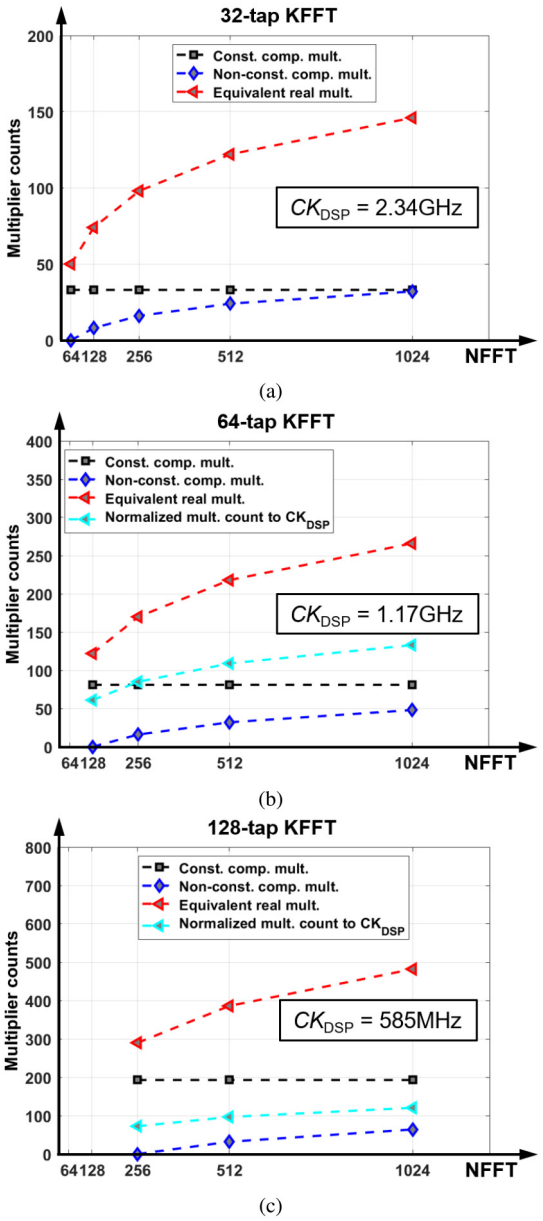

FIGURE 13. The number of multiplier cells in a partially-serial N-tap FFT with a 32-tap KFFT (a), a 64-tap KFFT (b), and with a 128-tap KFFT (c).

block and the first-stage twiddle factor multipliers array. The red lines show the equivalent number of real multipliers with the assumption that each complex multiplier includes three real multipliers and a constant multiplier has half the size of a non-constant multiplier. The light blue lines in Fig. 13(b) and Fig. 13(c) show the equivalent number of real multipliers normalized to the CK_{DSP} of Fig. 13(a), which indicate the power consumption metric of the N-tap FFT block. The number of required multipliers increases in a logarithmic way as a function of NFFT given a fixed target throughput in terms of the number of input/output samples. And from the light blue lines in Fig. 13, it can be seen that larger KFFT helps in decreasing the area-frequency product metric. Moreover,

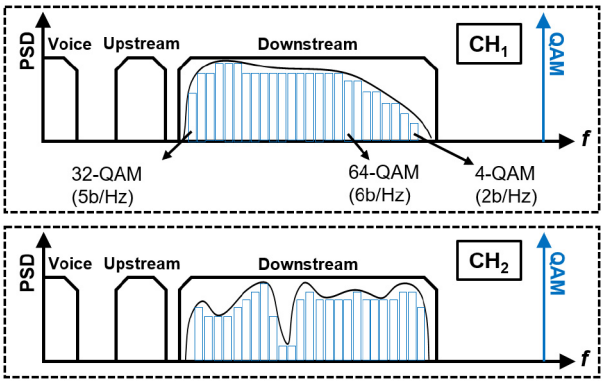

FIGURE 14. DMT modulation with bit-loading for adjustment to different channel SNR profiles (CH_1 and CH_2) in a DSL downstream example.

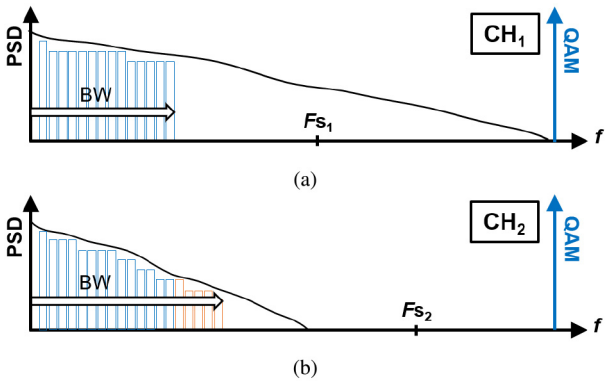

FIGURE 15. DMT bit-loading for two different channels with the same total data rate. The TRX with CH_2 (b) should run at a faster frequency with a larger analog bandwidth than with CH_1 (a).

with a larger KFFT, the DSP can operate at a lower clock rate; hence low-power standard cells with a smaller area can be used for the silicon implementation, which is not reflected in the graphs in Fig. 13. And as the FDE in a single-lane RX data-path consists of $\frac{K}{2}$ non-constant multipliers regardless of the NFFT tap sizes assuming that the FDE runs at the same clock rate CK_{DSP} as the FFT core, the only parameter that affects the power consumption of the FDE is the CK_{DSP}. In actual design, by allowing more parallelism (larger area) with lower clock frequency, lower-power standard cells can be used with lower VDD since the cells can have a longer propagation delay. Hence, lowering the clock frequency of the FDE helps in reducing the FDE's dynamic power consumption at the cost of increased area.

In short, the area and the power consumption of a partially-serial FFT processor increase in a logarithmic way as a function of NFFT. And for a given NFFT, the power consumption of FFT decreases with a larger KFFT. Under the assumption that the FDE runs at the same clock rate as the FFT, its power consumption also decreases with a larger sized KFFT.

IV. DESIGN CONSIDERATIONS FOR > 200 GB/S OFDM TRX DATA-PATH

A. BIT-LOADING VS HOMOGENEOUS QAM

An attractive property of OFDM is its ability to adapt to the SNR profile of the channel by adjusting the modulation order of QAM symbols and constellation amplitude loaded to each sub-channel. The OFDM with such an adaptation scheme is called discrete multitone (DMT). Bit-loading examples with DMT modulation for two different channels for digital subscriber line (DSL) applications are shown in Fig. 14, where the frequency responses of the channels exhibit largely different SNR profiles within a limited frequency band. With the DMT, communication with a reliable overall BER can be achieved by allocating proper number of bits to each sub-channel depending on the sub-channel SNR.

In ultra-high-speed wireline applications where the variable data rate is not acceptable, bit-loading should be applied under the constraint that the overall bandwidth efficiency is unchanged within a fixed frequency bandwidth as the channel condition varies. As shown in Fig. 15 where the data rates are the same in both Fig. 15(a) and Fig. 15(b), with a fixed data rate, a variable frequency bandwidth and sampling rate, the sub-channels at higher frequency band will exhibit relatively low spectral efficiency (compare Fig. 15(a) and Fig. 15(b)), resulting in the degradation of the average bandwidth efficiency.

Eventually, the gain of moving from PAM-4 to OFDM may be questionable if the spectral efficiency of the DMT gets close to that of PAM-4. In addition, if the occupied bandwidth and the corresponding sampling rate vary depending on the channel condition, the entire TRX system including the AFE, data converters, and the DSP, should support a wide range of tuning and an arbitrary clock frequency. As shown in Fig. 15(b), when there are some sub-channels with a low number of bits per Hz (orange bars), these sub-channels will contribute to the increase of the spectral bandwidth, and a faster sampling rate (F_{s2}) as compared to F_{s1} of Fig. 15(a) will be needed accordingly. Then, all the related circuits will be operating at a higher frequency with higher VDD increasing the power consumption, and the DSP should be designed with higher-performance cells, which are not the best for a low-power design.

Thus, for ultra high-speed wireline applications with a fixed data rate, bit-loading for different channel profiles should be applied in such a way that the average bits/Hz does not change with the channel condition. For example, if less bits should be allocated to some low-SNR sub-channels compared to the reference bit-loading setting, more bits should be loaded to other sub-channels by the same amount together with a proper power loading as well. While simultaneous bit and power loading under the fixed-bandwidth constraint is likely to result in a better overall BER performance of the link, this work considers only power loading to simplify the problem while still allowing the adjustment of the signal to the channel profile. However, it is worth noting that although power loading can adjust the TX signal to point-to-point interfaces which have relatively smooth frequency response,

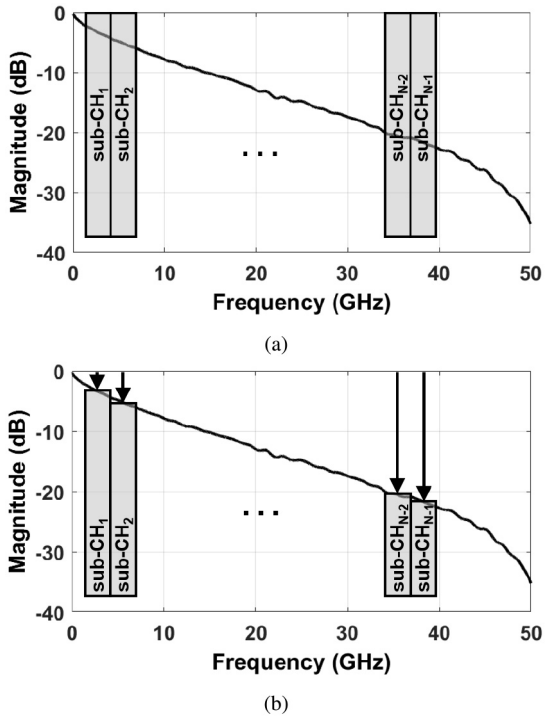

FIGURE 16. The illustration of sub-channel PSD at the TX output (a) and at the RX input (b) with a homogeneous QAM order and with the identical power for all sub-channels.

applying only power loading to communicate over a channel having significant notches in its frequency response (e.g., with multi-drop bus) would limit the advantages of using OFDM.

While moving to a higher-order QAM requires a tighter SNR margin, packing the data into a narrower spectral bandwidth brings several advantages, not limited to the improved bandwidth efficiency coming from the QAM-order itself. As the Baudrate reduces with the bandwidth reduction, the number of important ISI cursors of the time-domain pulse response will be reduced as well. Therefore, the required number of CP taps will be reduced, and this is translated into a reduced bandwidth overhead from the CP. Also, the relaxed data converters' sampling rate helps in increasing the overall SNR as the signal-to-noise and distortion ratio (SNDR) performance of ADC tends to be degraded with a faster sampling rate [27]. In addition, the power consumption of digital-intensive successive approximation register (SAR) ADCs and the modem DSPs will be reduced as well, where the power consumption tends to be increased exponentially with the operating frequency due to the requirement of higher supply voltages [21], [27]. However, it should be noted that packing the data into a narrower spectral bandwidth may not be always possible since it requires a larger FFT size, especially if the DSP is area-constrained.

B. POWER LOADING

Fig. 16(a) illustrates the PSD of the OFDM signal at the TX-side, with a homogeneous QAM order for every sub-channel having the same amount of power allocated to each

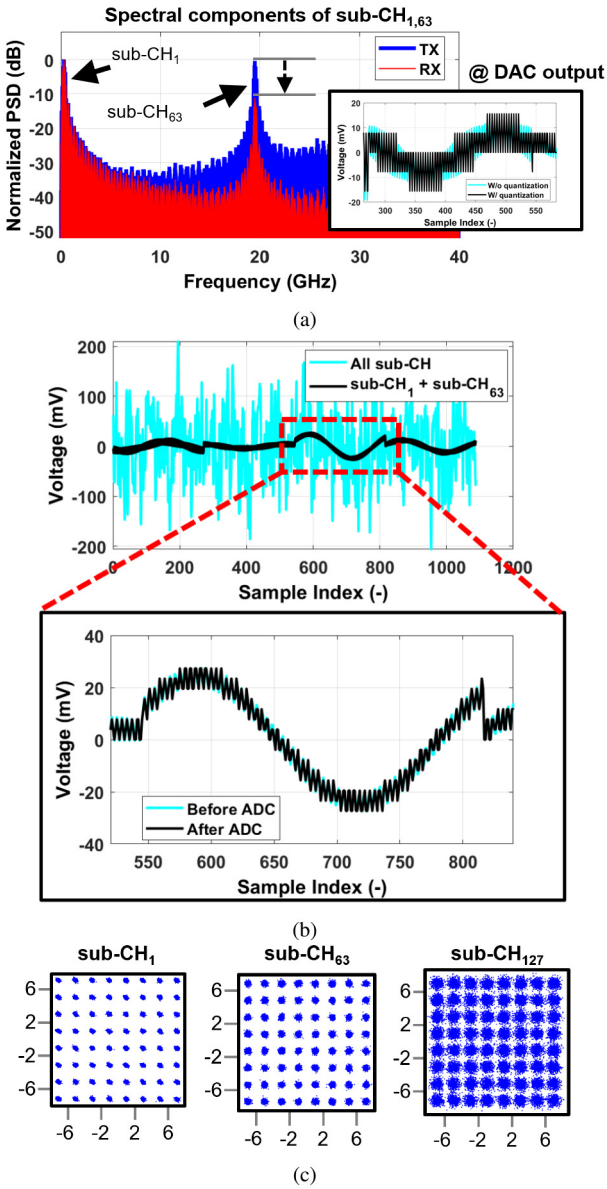

FIGURE 17. The extracted TX/RX PSD of the 1st and the 63rd OFDM sub-channels with identical transmitted sub-channel power and the time-domain DAC output signal (a), the received time-domain signal of the 63rd sub-channel contents superimposed on top of the 1st sub-channel signal (b), and the constellation diagrams of the equalized QAM symbols on the RX at the 1st, the 63rd, and the 127th sub-channels (c).

sub-channel. At the output of the TX DAC, the average PSD over the entire bandwidth will be flat with the same quantization SNR (QSNR) for every sub-channel, ignoring the sync roll-off effect at the output of the DAC. And on the RX-side, as shown in Fig. 16(b), the amplitude of low-frequency sub-channels are more or less preserved while the high-frequency contents are significantly attenuated due to the channel loss, resulting in an amplitude imbalance between the low- and the high-frequency contents. Therefore, the QSNR after the ADC will vary considerably depending on the sub-channel, and the BER of the high-frequency sub-channels will dominate the overall BER.

Fig. 17(a) shows an example of sending a 64-QAM OFDM with 127 sub-channels (256-tap FFT) without any power

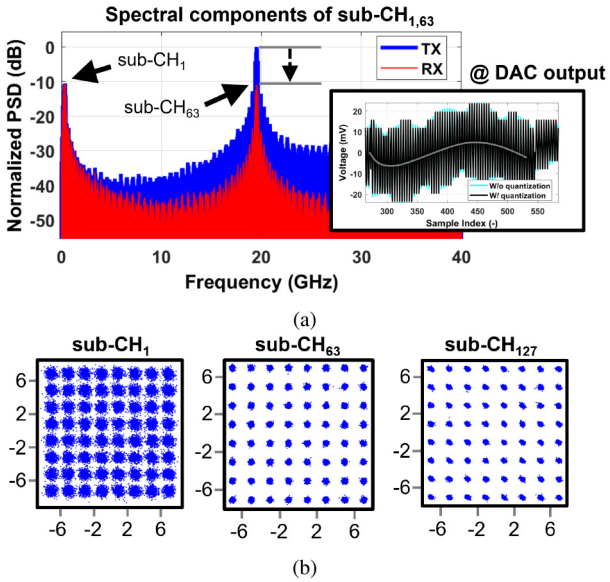

FIGURE 18. The extracted TX/RX PSD of the 1st and the 63rd OFDM sub-channels with channel-inverting power loading and the time-domain DAC output signal (a), and the constellation diagrams of the equalized QAM symbols on the RX at the 1st, the 63rd, and the 127th sub-channels (b).

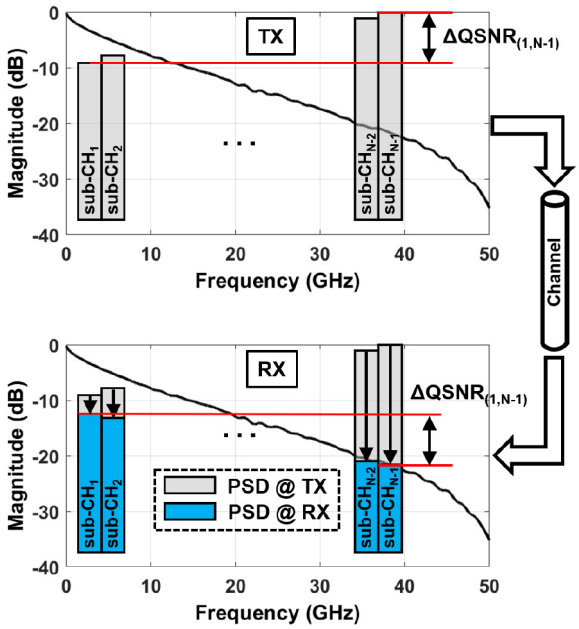

FIGURE 19. A balanced power loading for a QSNR-limited OFDM system.

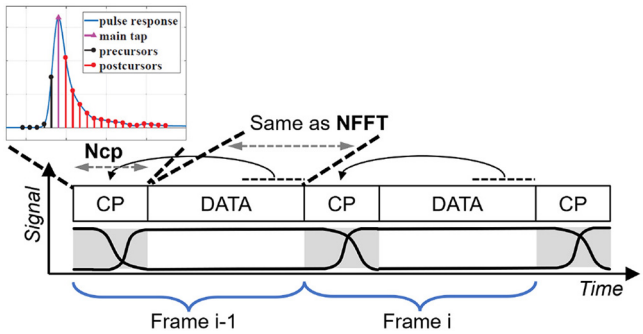

FIGURE 20. The time-domain OFDM samples with data symbols and with CPs of which number of taps is large enough to cover important ISI cursors.

loading on the TX-side. By extracting the spectral contents of the 1st and the 63rd sub-channels from the OFDM signal, it can be seen that the power of these two sub-channels are equal on the TX-side, while the high-frequency contents are attenuated by more than 10 dB on the RX-side. In the time-domain, at the DAC output, it can be seen that the amplitude of the low- and the high-frequency sinusoid basis are almost equal (right inset of Fig. 17(a)), but on the RX-side, it can be observed from Fig. 17(b) that the amplitude of the high-frequency basis is significantly more attenuated as compared to that of the low-frequency basis. And such amplitude attenuation translates into a QSNR degradation after the analog-to-digital conversion, where the resulting constellation diagrams are shown in Fig. 17(c) clearly showing the effects of the QSNR imbalance.

Fig. 18(a) shows the PSD of the 1st and the 63rd sub-channels extracted from the TX and the RX OFDM signal, with a complete inversion of the channel by the TX power loading. From the right inset of Fig. 18(a), it can be seen that the amplitude of the high-frequency component (63rd sub-channel basis) is much larger than that of the low-frequency component (the 1st sub-channel basis) at DAC output. While the 1st and the 63rd sub-channel power are similar on the RX-side, the QSNR penalty of the 1st sub-channel as compared to the 63rd sub-channel at the TX-side results in the imbalanced SNR over the sub-channels as can be seen from constellation diagrams shown in Fig. 18(b). Hence, a balanced power loading that allows all sub-channels to experience the equivalent QSNR degradation by the TX power-loading and the channel attenuation has to be applied in a QSNR-limited OFDM system, as illustrated in Fig. 19. It is worth noting that while this work considers consistent SNDR of DAC and ADC over the entire frequency bandwidth, in actual OFDM-based high-speed SerDes, the SNDR of DAC and ADC are frequency-dependent and the TX power loading should take this into account as well.

C. FFT SIZE, NUMBER OF CP TAPS, AND BER

In the time-domain OFDM signal, there are samples for data symbols (DATA in Fig. 20) that convey the actual information, and there are also CP samples that do not carry any data but just protect the data symbols from ISI. Given the Baudrate of the signal, N_{CP} should be large enough to cover important ISI cursors of the pulse response, and the number of samples in each data symbol is the same as NFFT, as shown in Fig. 20. The bandwidth of the total OFDM signal is expanded due to the overhead caused by the presence of the CP, and the overhead in percentage can be expressed as

$$O_{CP} = \frac{100 \times N_{CP}}{NFFT} \quad (4)$$

where O_{CP} is the bandwidth overhead coming from the CP. Given the N_{CP} determined by a pulse response, the O_{CP}

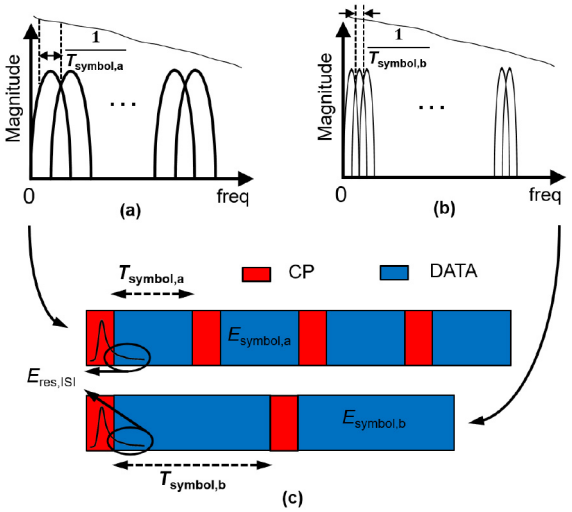

FIGURE 21. The OFDM sub-channel configurations for the same data rate with a wide sub-channel spacing (a), a narrower sub-channel spacing with more number of sub-channels (b), and the time-domain view with each configuration (c).

reduces as NFFT increases. Moreover, the residual ISI cursors not covered by the CP will have a reduced effect on the SNR degradation since the portion of the OFDM symbol energy will be increased over the residual ISI energy as

$$\frac{E_{\text{symbol,b}}}{E_{\text{res,ISI}}} > \frac{E_{\text{symbol,a}}}{E_{\text{res,ISI}}} \tag{5}$$

where $E_{\text{symbol,a}}$ and $E_{\text{symbol,b}}$ are the average OFDM data symbol energy with the OFDM sub-channel configurations shown in Fig. 21(a) and Fig. 21(b), respectively, and $E_{\text{res,ISI}}$ is the energy of the residual ISI cursors. The time-domain view of the CPs and the data symbols corresponding to the OFDM sub-channel configurations in Fig. 21(a) and Fig. 21(b) is shown in Fig. 21(c).

A larger NFFT does not only reduce the denominator of (4) but also reduces the required N_{CP}, further reducing the bandwidth overhead caused by the presence of the CP. An example illustrating the O_{CP} difference is shown in Fig. 22. Given a pulse response that requires $N_{\text{CP}} = 4$ with $NFFT = 16$ for a fixed target data rate, $O_{\text{CP}} = 25\,\%$ as illustrated in Fig. 22(a) where T_{16} is the duration required to send the data in an OFDM symbol including the CP. By doubling the NFFT (Fig. 22(b)), the O_{CP} is cut by half, and $T_{32} < 2 \times T_{16}$ meaning that the duration to send the same amount of data with the case of $NFFT = 16$ is shorter when $NFFT = 32$. For communicating at a fixed target data rate, the sample width in the time-domain can be adjusted such that $T'_{32} = 2 \times T_{16}$, resulting in a reduced Baudrate. The reduced Baudrate mitigates the ISI, and the N_{CP} can be lowered as shown in Fig. 22(c), and the lowered N_{CP} further reduces the O_{CP}.

Fig. 23(a) shows the OFDM signal bandwidth including the CP samples versus the FFT tap count, where the necessary CP length is obtained based on the ratio between the sum of the residual ISI cursors not covered by the CP and the main-cursor amplitude of the channel Baudrate pulse response. The N_{CP} corresponding to each NFFT is determined to satisfy

$$\frac{\sum ISI_{\text{Res}}}{A} < C_{\text{thres}} \tag{6}$$

where A is the main-cursor amplitude at Baudrate with a specific NFFT, and $C_{\text{thres}} = 0.25$ satisfies the BER requirement (BER$< 10^{-5}$) for any FFT size with NFFT ≥ 64 given the modeling environment in this work. Fig. 23(b) shows the N_{CP} obtained based on (6) versus NFFT, and the corresponding overhead for communicating at 224 Gb/s over channel 3 shown in Fig. 9.

V. SIMULATION RESULTS

This Section provides simulation results that characterize the BER performance of an OFDM serial link as a function of NFFT, N_{CP}, the DAC/ADC clipping ratio, and the DAC/ADC ENOB. With the time-domain simulation, a random binary stream of 48×10^6 bits is modulated to 64-QAM symbols, the data rate is fixed to 224 Gb/s, 7-bit resolution is considered for the DAC and the ADC, root-mean-square (RMS) random jitter (Rj_{RMS}) of 1 % unit interval (UI) is applied to data conversions, and a random noise RMS amplitude of 1.8 mV on the AFE is considered.

Simulation results shown in Fig. 24 are obtained by setting the DAC and the ADC clipping ratios to $R_{\text{clip,DAC}} = 3.41$ and $R_{\text{clip,ADC}} = 3.73$, respectively. The clipping ratios are defined as $\text{R}_{\text{clip,DAC}} = \beta_{\text{DAC}}/A_{\text{RMS,DAC}}$ and $\text{R}_{\text{clip,ADC}} = \beta_{\text{ADC}}/A_{\text{RMS,ADC}}$ for DAC and ADC, respectively, where β_{DAC} and β_{ADC} are saturation thresholds in absolute value on the DAC and the ADC, respectively, $A_{\text{RMS,DAC}}$ and $A_{\text{RMS,ADC}}$ are the root-mean-square values of the time-domain OFDM signal on the DAC output and on the ADC input, respectively. For DAC, saturation threshold is defined as the DAC analog differential output amplitude corresponding to the maximum DAC input digital code (in magnitude, signed number). And for the ADC, it is defined as the ADC input differential amplitude that converts to the maximum ADC digital code (in signed amplitude, centered at common-mode). It should be noted that over-clipping will cause excessive nonlinear distortion although it allows the overall signal power to be increased. On the other hand, under-clipping will make the overall SNR low despite the lowered clipping-induced nonlinear distortion. Hence, finding optimal clipping ratios both for the DAC and ADC that maximize the benefit coming from the overall SNR increase versus the penalty coming from the clipping-induced nonlinear distortion is important.

The CH3 shown in Fig. 9 is used as the channel for the following set of simulations. It can be seen from Fig. 24 that the BER performance improves with larger FFTs, and there is an optimal N_{CP} for each FFT size. From Fig. 24, it can be observed that the major ISI cursors should be covered by the CP, and a long CP can make the overall SNR even worse due to the increased signal bandwidth which makes the OFDM signal experience more severe channel loss

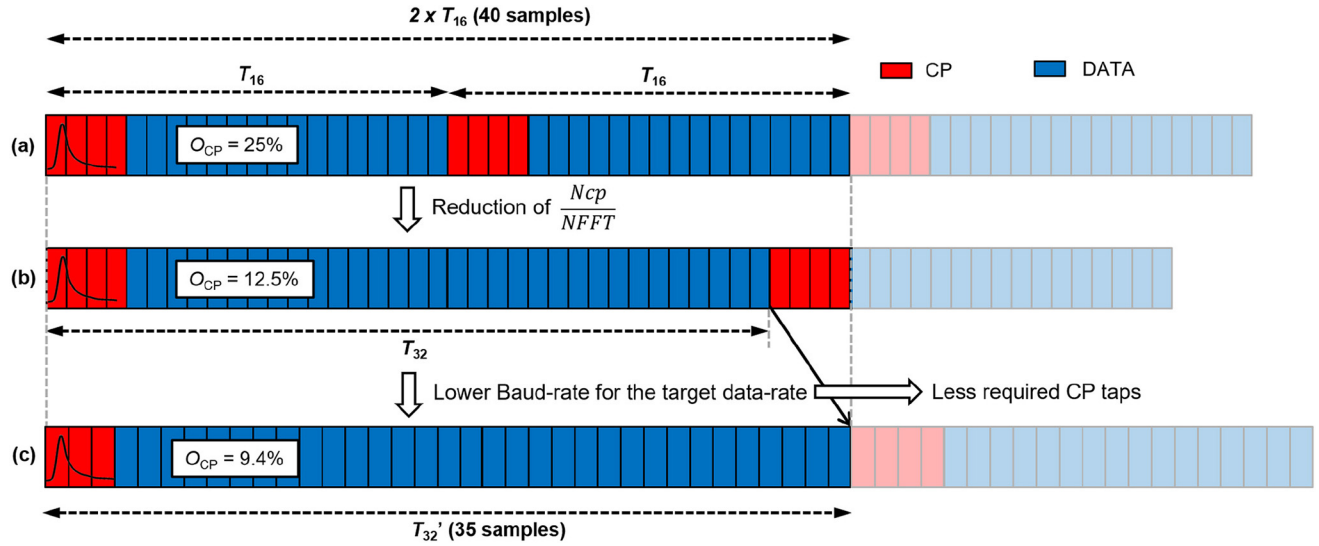

FIGURE 22. Illustration of the time-domain OFDM samples with $NFFT = 16$ and $N_{CP} = 4$ (a), with $NFFT = 32$ and $N_{CP} = 4$ (b), and with $NFFT = 32$ and $N_{CP} = 3$ (c) with the reduced Baudrate. The data rates are equal in (a) and in (c), and is the highest in (b).

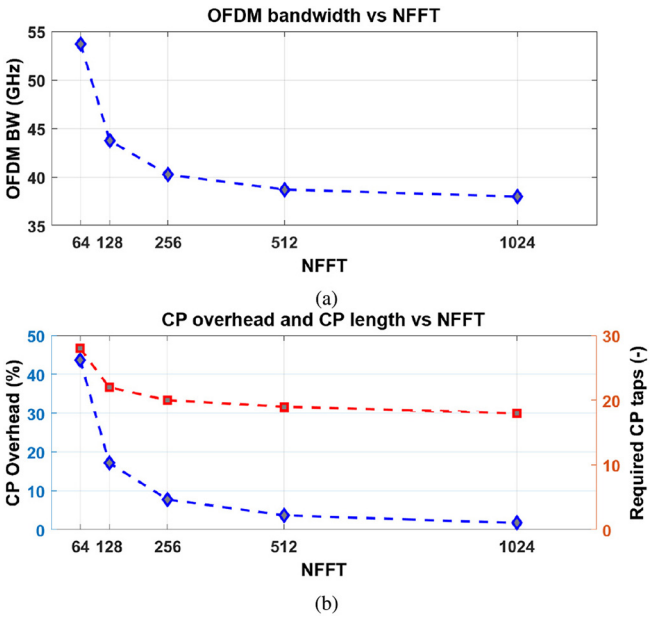

FIGURE 23. The OFDM signal bandwidth for different NFFT including the CP samples at 224 Gb/s data rate (a), and the required CP samples to meet the maximum BER target and the corresponding CP overhead in percentage (b).

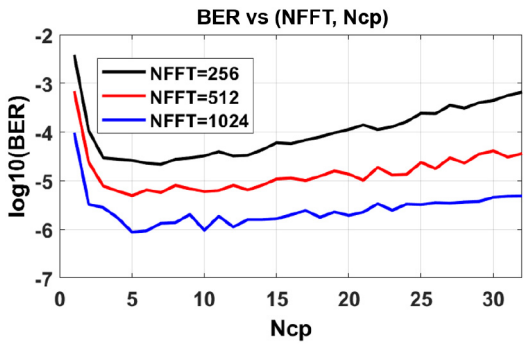

FIGURE 24. BER simulations as a function of N_{CP} for NFFT ∈ {256, 512, 1024}.

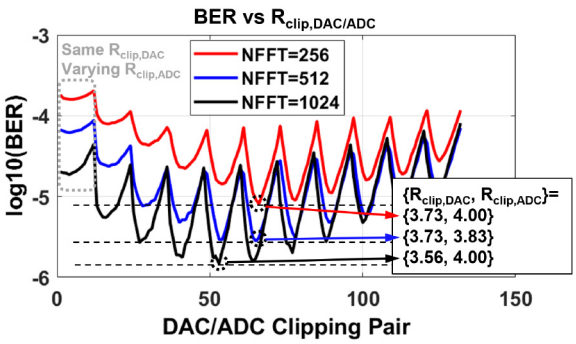

FIGURE 25. BER simulations as a function of the pair of the clipping ratios at the DAC and the ADC ($\{R_{clip,DAC}, R_{clip,ADC}\}$), where $R_{clip,DAC} \in [2.97, 4.70]$ and $R_{clip,ADC} \in [3.33, 5.40]$. The $R_{clip,DAC}$ and $R_{clip,ADC}$ ranges are divided into 11 and 12 points, respectively.

on average. With a 128-tap KFFT, the BER performance improves by more than an order of magnitude as NFFT enlarges from 256 to 1024, at the cost of 66% increased FFT arithmetic resources. Considering the overall RX DSP area where the FDE requires $3 \times K/2$ multipliers, the area overhead reduces to 40%. Furthermore, larger NFFT relaxes the Baudrate of the system (explained in Fig. 22), reducing the power consumption of the data converters.

Fig. 25 shows the BER simulation results performed by sweeping the DAC and the ADC clipping ratios for different NFFT settings, but with a fixed $N_{CP} = 8$ for CH3. It can be observed from Fig. 25 that the overall tendency that the higher FFT size results in better BER performance is maintained, but the best DAC and the ADC clipping ratio pair varies depending on the FFT size.

The BER simulation results by sweeping the effective number of bits (ENOB) of data converters are shown in

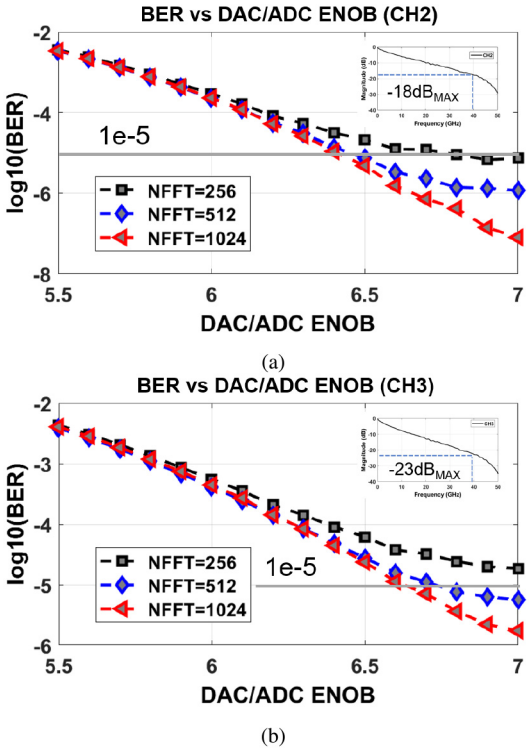

FIGURE 26. BER simulations as a function of DAC and ADC effective number of bits (ENOB) for NFFT $\in$ {256, 512, 1024} with CH2 (a) and CH3 (b) of which profiles are shown in Fig. 9.

Fig. 26, with a fixed $N_{\mathrm{CP}} = 8$ and with fixed clipping ratios $R_{\mathrm{clip,DAC}} = 3.41$ and $R_{\mathrm{clip,ADC}} = 3.73$ for all NFFT settings. The CH2 is less attenuative than CH3 (Fig. 9), and thus the overall SNR of the OFDM signal on the RX with CH2 is higher than with CH3 within the same bandwidth. With CH2, Fig. 26(a) shows that the link performance is limited by quantization noise for any simulated DAC/ADC ENOB when NFFT $=$ 1024. With NFFT $\in \{256, 512\}$, it can be observed that the BER saturates as the DAC/ADC ENOB improves, meaning that the link performance is limited by random noise. Noting that a large FFT size helps in reducing the overall occupied bandwidth of the OFDM signal thanks to the reduced CP overhead, it can be understood that the larger NFFT is, the narrower the occupied bandwidth is, then the lower the maximum and the average signal power loss is, and hence the higher the overall SNR of the OFDM signal is. From Fig. 26(b), the same tendency as in Fig. 26(a) can be observed with worse BER performances due to the reduced overall SNR within the same frequency bandwidth.

VI. CONCLUSION

The overall architecture of an OFDM wireline TRX is explained together with the data flow from the TX to the RX through electrical channels. The design of a partially-serial FFT architecture is presented, and the impacts of FFT design choices on the overall link performance are analyzed. Given a fixed compute throughput requirement, it is shown that the FFT area increases only logarithmically with the number of taps, and a larger size of kernel FFT helps in reducing the power consumption of the FFT processor for a fixed total number of FFT taps. Increasing the FFT tap counts reduces the required number of CP samples for communicating at the same target data rate, further relaxing the required Baudrate of the system. System-level modeling results show that the BER improves by more than an order of magnitude by increasing the FFT tap counts from 256 to 1024, at the cost of 66% of overhead in terms of the number of multipliers in the FFT. Taking the AFE and the DAC/ADC power and area into consideration, the BER improvement of more than a magnitude at the cost of 66% increase in the FFT area is a reasonable trade-off.

REFERENCES

[1] "Cisco annual Internet report (2018–2023) white paper," Cisco, San Jose, CA, USA, White Paper, 2020.
[2] S. Gondi, J. Lee, D. Takeuchi, and B. Razavi, "A 10Gb/s CMOS adaptive equalizer for backplane applications," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2005, pp. 328–330.
[3] M. Harwood *et al.*, "A 12.5Gb/s SerDes in 65nm CMOS using a baud-rate ADC with digital receiver equalization and clock recovery," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 436–438.
[4] A. Emami-Neyestanak *et al.*, "A 6.0-mW 10.0-Gb/s receiver with switched-capacitor summation DFE," *IEEE J. Solid-State Circuits*, vol. 42, no. 4, pp. 889–896, Apr. 2007.
[5] C. Menolfi *et al.*, "A 112Gb/s 2.6pJ/b 8-tap FFE PAM-4 SST TX in 14nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2018, pp. 104–106.
[6] J. Im *et al.*, "A 112-Gb/s PAM-4 long-reach wireline transceiver using a 36-way time-interleaved SAR ADC and inverter-based RX analog front-end in 7-nm FinFET," *IEEE J. Solid-State Circuits*, vol. 56, no. 1, pp. 7–18, Jan. 2021.
[7] P. Mishra *et al.*, "8.7 A 112Gb/s ADC-DSP-based PAM-4 transceiver for long-reach applications with >40dB channel loss in 7nm FinFET," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2021, pp. 138–140.
[8] M.-A. LaCroix *et al.*, "8.4 A 116Gb/s DSP-based wireline transceiver in 7nm CMOS achieving 6pJ/b at 45dB loss in PAM-4/Duo-PAM-4 and 52dB in PAM-2," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2021, pp. 132–134.
[9] J. F. Bulzacchelli *et al.*, "A 28-Gb/s 4-tap FFE/15-tap DFE serial link transceiver in 32-nm SOI CMOS technology," *IEEE J. Solid-State Circuits*, vol. 47, no. 12, pp. 3232–3248, Dec. 2012.
[10] T. Toifl *et al.*, "A 3.5pJ/bit 8-tap-feed-forward 8-tap-decision feedback digital equalizer for 16Gb/s I/Os," in *Proc. 40th Eur. Solid-State Circuits Conf. (ESSCIRC)*, 2014, pp. 455–458.
[11] T. O. Dickson, J. F. Bulzacchelli, and D. J. Friedman, "A 12-Gb/s 11-mW half-rate sampled 5-tap decision feedback equalizer with current-integrating summers in 45-nm SOI CMOS technology," *IEEE J. Solid-State Circuits*, vol. 44, no. 4, pp. 1298–1305, Apr. 2009.
[12] T. Toifl *et al.*, "A 2.6 mW/Gbps 12.5 Gbps RX with 8-tap switched-capacitor DFE in 32 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 47, no. 4, pp. 897–910, Apr. 2012.
[13] H. Yueksel *et al.*, "A 3.6pJ/b 56Gb/s 4-PAM receiver with 6-bit TI-SAR ADC and quarter-rate speculative 2-tap DFE in 32 nm CMOS," in *Proc. Eur. Solid-State Circuits Conf. (ESSCIRC)*, 2015, pp. 148–151.
[14] Y. Frans *et al.*, "A 56-Gb/s PAM4 wireline transceiver using a 32-way time-interleaved SAR ADC in 16-nm FinFET," *IEEE J. Solid-State Circuits*, vol. 52, no. 4, pp. 1101–1110, Apr. 2017.
[15] A. Roshan-Zamir *et al.*, "A 56-Gb/s PAM4 receiver with low-overhead techniques for threshold and edge-based DFE FIR- and IIR-tap adaptation in 65-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 54, no. 3, pp. 672–684, Mar. 2019.
[16] M. Pisati *et al.*, "6.3 A sub-250mW 1-to-56Gb/s continuous-range PAM-4 42.5dB IL ADC/DAC-based transceiver in 7nm FinFET," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2019, pp. 116–118.
[17] T. Ali *et al.*, "6.4 A 180mW 56Gb/s DSP-based transceiver for high density IOs in data center switches in 7nm FinFET technology," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2019, pp. 118–120.

[18] P. Sun, "100Gb/s single-lane SerDes discussion," IEEE, Piscataway, NJ, USA, IEEE document P802.3 New Ethernet Applications Ad Hoc, 2017.
[19] A. Dochhan *et al.*, "56 Gb/s DMT transmission with VCSELs in 1.5 um wavelength range over up to 12 km for DWDM intra-data center connects," in *Proc. ECOC*, Sep. 2016, pp. 391–393.
[20] N. Eiselt *et al.*, "Performance comparison of 112-Gb/s DMT, Nyquist PAM4, and partial-response PAM4 for future 5G Ethernet-based fronthaul architecture," *IEEE J. Lightw. Technol.*, vol. 36, no. 10, pp. 1807–1814, May 15, 2018.
[21] G. Kim *et al.*, "A 161-mW 56-Gb/s ADC-based discrete multitone wireline receiver data-path in 14-nm FinFET," *IEEE J. Solid-State Circuits*, vol. 55, no. 1, pp. 38–48, Jan. 2020.
[22] B. Vatankhahghadim, N. Wary, J. Bailey, and A. C. Carusone, "A study of discrete multitone modulation for wireline links beyond 100 Gb/s," *IEEE Open J. Circuits Syst.*, vol. 2, pp. 78–90, 2021.
[23] J. Salinas, J. Cosson-Martin, M. Laghaei, H. Shakiba, and A. Sheikholeslami, "Performance comparison of baseband signaling and discrete multi-tone for wireline communication," *IEEE Open J. Circuits Syst.*, vol. 2, pp. 65–77, 2021.
[24] J. Cosson-Martin, H. Shakiba, and A. Sheikholeslami, "Timing recovery and adaptive equalization for discrete multi-tone signalling in wireline applications," *IEEE Open J. Circuits Syst.*, vol. 2, pp. 856–868, 2021.
[25] J. Armstrong, "OFDM for optical communications," *IEEE J. Lightw. Technol.*, vol. 27, no. 3, pp. 189–204, Feb. 1, 2009.
[26] G. Kim, W. Kwon, T. Toifl, Y. Leblebici, and H. Bae, "Design considerations and performance trade-offs for 56Gb/s discrete multi-tone electrical link," in *Proc. IEEE 62nd Int. Midwest Symp. Circuits Syst. (MWSCAS)*, 2019, pp. 1147–1150.
[27] L. Kull *et al.*, "A 10-bit 20–40 GS/s ADC with 37 dB SNDR at 40 GHz input using first order sampling bandwidth calibration," in *Proc. IEEE Symp. VLSI Circuits*, Honolulu, HI, USA, 2018, pp. 275–276.
[28] "IEEE P802.3ck Task Force—Tools and Channels." 2020. [Online]. Available: http://www.ieee802.org/3/ck/public/tools/index.html

GAIN KIM (Member, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering from the Swiss Federal Institute of Technology in Lausanne (EPFL), Lausanne, Switzerland, in 2013 and 2015, respectively, and the Ph.D. degree in microsystems and microelectronics from EPFL, Lausanne, in 2018. From September 2016 to July 2018, he was with the High-Speed Interconnect Technology Group, IBM Research—Zurich, Rüschlikon, Switzerland, working on the design of ADC-based wireline receivers. From 2018 to 2020, he was a Postdoctoral Fellow with KAIST. From 2020 to 2022, he was with Samsung Research, Seoul, South Korea, as a Staff Engineer working on a baseband modem for 6G wireless communications. In 2022, he joined the faculty of the EECS Department, Daegu Gyeongbuk Institute of Science and Technology, Daegu, South Korea, where he is currently an Assistant Professor. He is the recipient of the 2018 IEEE Circuits and Systems Pre-Doctoral Scholarship Award.