

Timing Recovery and Adaptive Equalization for Discrete Multi-Tone Signalling in Wireline Applications

JEREMY COSSON-MARTIN¹ (Graduate Student Member, IEEE),
HOSSEIN SHAKIBA² (Senior Member, IEEE), AND ALI SHEIKHOLESLAMI¹ (Senior Member, IEEE)

¹The Edward S. Rogers Sr. Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON M5S 3G4, Canada

²Huawei Technologies Canada, Markham, ON L3R 5A4, Canada

This article was recommended by Associate Editor C.-M. Hsu.

CORRESPONDING AUTHOR: J. COSSON-MARTIN (e-mail: marti701@ece.utoronto.ca)

This work was supported in part by the Huawei Technologies Canada and in part by the Natural Sciences and Engineering Research Council (NSERC) of Canada.

ABSTRACT This paper proposes a discrete multi-tone timing-recovery system with adaptive equalization for ultra-high-speed wireline applications. It combines frequency-domain clock recovery with decision-directed equalization to improve receiver performance while eliminating the need for pilot carriers, thereby increasing spectral efficiency. Compared to a conventional pilot-carrier-based technique employing four pilot carriers and a 32-point FFT, this approach improves phase-error sensitivity by 3.6 times, tracking bandwidth by 1.7 times, increases the jitter tolerance slope by 20dB per decade at low frequency, and removes residual equalization error, resulting in an overall data-rate increase of 27%. The concept is validated at the system-level and gate-level through synthesis in an FPGA. A convergence analysis of both the adaptive equalizer and clock synchronization shows the system's ability to mitigate error propagation and remain synchronized in the presence of impairments. Finally, we highlight the system's ability to trade-off clock convergence versus phase error sensitivity. Either parameter can be adjusted by 15 times, optimizing the receiver over a broad range of signal conditions.

INDEX TERMS Adaptive equalization, clock and data recovery (CDR), decision-directed equalization, discrete multi-tone (DMT), orthogonal frequency division multiplexing (OFDM), SERDES, single-tap equalization, timing recovery, wireline.

I. INTRODUCTION

THE WIRELINE industry is approaching a turning point. Today's state-of-the-art 100Gb/s long-reach 4-PAM systems require complex equalization schemes to combat severe Inter-Symbol Interference (ISI) including feed-forward equalizers of 25+ taps [1], [2]. With an ever-increasing data-rate, transceiver complexity continues to worsen; it will soon lead to impractical power consumption [3]. As such, in the pursuit to reach channel capacity [4], alternative signal modulation techniques are being considered [5], [6], [7], including Discrete Multi-Tone (DMT) signalling [2], [8], [9].

DMT is a mature technology dominant in communication applications that must optimize spectral efficiency. This

includes low-rate baseband telecommunication systems such as ADSL and VDSL [9], [10] and long-haul optical applications such as single-mode and multi-mode fiber [11], [12]. Furthermore, DMT's bandpass counterpart, Orthogonal Frequency Division Multiplexing (OFDM), dominates the wireless industry, including LTE and 5G [13]. As demonstrated in [5], DMT improves control over equalization, leading to a higher data rate while consuming less power.

One barrier impeding the widespread adoption of DMT is that high-performance timing recovery and channel estimation techniques must reserve a large number of frequency bins [14], creating an overhead of 27% for a 32-point FFT system. This paper analyzes the problem and proposes a solution that avoids bin reservation while improving phase-error

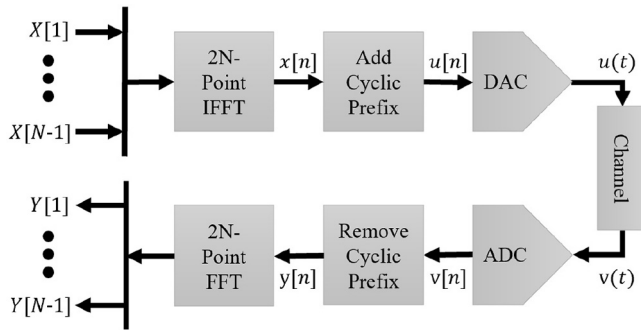


FIGURE 1. DMT top-level block diagram.

sensitivity, tracking bandwidth, jitter tolerance, and residual equalization error.

This paper is organized as follows: Section II gives a brief background, Section III investigates the overhead from bin reservation, Section IV describes the proposed receiver, Section V compares the solution to conventional approaches, highlights its frequency tracking capabilities, and discusses its implementation in a Cyclone V FPGA, Section VI analyzes the convergence of both the timing recovery and decision-directed equalizer, and Section VII concludes this paper.

II. BACKGROUND

This section provides a brief background of DMT signalling with a focus on equalization and timing recovery.

A. DISCRETE MULTI-TONE

As shown in Fig. 1, DMT splits a channel into $N-1$ orthogonal frequency bins and sends a tone in each where Quadrature Amplitude Modulation (QAM) adjusts the amplitude and phase to convey information.

At the input, bits are mapped to complex-valued symbols $X[1]$ to $X[N-1]$. Each is selected from a pre-defined constellation, encoding $\log_2(N_S)$ bits, where N_S is the constellation size. They are then fed to a $2N$ -point Inverse-Fast-Fourier Transform (IFFT) to create the tones $x[n]$. A Cyclic Prefix (CP), which is discussed later, is appended, and the signal $u(t)$ is sent through the channel. At the receiver, the CP is removed from $v(t)$, and the tones $y[n]$ are sent through an FFT to recover symbols $Y[1]$ to $Y[N-1]$. Finally, they are decoded and converted back to bits.

The Fourier transforms are modeled using (1) and (2) where $2N$ is the size of the transform.

$$x[n] = IFFT_{2N}\{X[k]\} = \frac{1}{2N} \sum_{k=0}^{2N-1} X[k]e^{j2\pi kn/2N} \quad (1)$$

$$Y[k] = FFT_{2N}\{y[n]\} = \sum_{n=0}^{2N-1} y[n]e^{-j2\pi kn/2N}. \quad (2)$$

The constellation size (bit-loading, N_S) and the separation between symbols (power-loading) is optimized using the Water Pouring Algorithm [15]. Power-loading is adjusted

to equate the error rate across all bins, whereas bit-loading is optimized to transmit as many error-free bits as possible using the least amount of signal power. In turn, this maximizes data rate given a target error rate and a set of channel impairments.

Wireless systems apply transforms containing thousands of bins [16]. However, due to low-latency requirements and complexity constraints, ultra-high-speed signalling is often bin limited. The receiver area increases proportionally to the FFT resolution. For DMT to be comparable in area and latency to a 4-PAM transceiver, $2N = 32$ is often selected [5]. Additionally, unlike wireless systems, wireline systems can only transmit real-valued signals. This necessitates a mirroring of the spectral content around the Nyquist frequency. Therefore, only the first N bins (bin 0 to bin $N-1$) can send unique information. The remaining N bins are reserved for sending the complex conjugate of the data [11]. Furthermore, DC (bin 0) is generally not used as it cannot support complex symbols without producing a complex output. As a result, the number of usable bins is limited to $N-1 = 15$ [5], [17].

B. EQUALIZATION

Practical channels are band-limited, causing ISI. 4-PAM systems reduce ISI by amplifying the high-frequency content of the signal using equalizer filters. At higher data rates, ISI is more detrimental, requiring more sophisticated filters. Moreover, the optimization of equalizer coefficients is a non-trivial problem, often requiring an LMS or even a genetic search algorithm [1], [18]. These controllers are also becoming more complex, worsening the problem.

An advantage of DMT is its ability to eliminate the need for conventional equalization. Since each symbol is carried by a different tone and since tones are orthogonal from one another, the symbols will not interfere. Therefore, the system will experience negligible ISI. This permits single-tap frequency-domain equalization, to be discussed shortly. Moreover, unlike 4-PAM systems, each equalizer tap in a DMT system is optimized independently. This simplifies the control algorithm leading to more optimal spectral utilization with the potential to achieve a higher capacity [2], [8], [9], [19]. However, in DMT, three sources of impairment remain: Inter-Frame-Interference (IFI), attenuation, and phase delay.

DMT transmits data using frames. ISI causes energy from one frame to leak into the next, producing IFI. This removes orthogonality among bins and degrades performance. To mitigate this effect, a guard interval called a CP is added [20]. As shown in Fig. 2, this is created by appending the end portion of the frame to the front resulting in an overall length of $2N + N_{CP}$ Unit Internals (UI) where N_{CP} is the length of the CP. As demonstrated in [11], with a sufficiently long CP, IFI is removed, which ensures frames are independent of one another and guarantees orthogonality among bins. However, a longer CP increases the overhead, particularly for bin-limited systems. Although the required length depends on the sampling frequency and the channel characteristic, it




FIGURE 2. Guard interval removing Inter-Frame Interference (IFI).




FIGURE 3. Attenuation and phase distortion from a band-limited channel.

can be shortened by partially equalizing the signal before the ADC using continuous-time linear equalization.

As shown in Fig. 3, when a tone is sent through a band-limited channel, it will experience attenuation and phase delay. This manifests as constant scaling and rotation error at the output. Thus, single-tap equalizers are added after the FFT in the receiver to perform a complex multiplication per bin to correct the distortion.

Equalizing bins with a constant correction factor is not sufficient. In wireline applications, the channel is assumed short-term stationary. Although environmental factors such as temperature and aging affect the channel response, the change is assumed to be slow. However, when noise and jitter are applied, the link experiences rapid changes in magnitude and phase. As such, to the receiver, it does not appear stationary. This causes constellations to scale and rotate [21]. The solution is adaptive equalization, where we continuously measure the magnitude and phase response of the link to determine the required coefficients, a process called channel estimation.

Accurate channel estimation mitigates low-frequency disturbances. However, as illustrated in Fig. 4, poor channel estimation results in residual equalization error where the equalized constellations need further scaling and rotation [21], [22]. This worsens sensitivity to impairments as certain symbols land closer to decision boundaries. Two common approaches to channel estimation are: use of training sequence and use of pilot carriers.

1) USE OF TRAINING SEQUENCE

Periodically, the transmitter sends a training sequence composed of pre-defined symbols across each frequency bin. This uncovers the channel’s magnitude and phase response. As such, the receiver determines the required equalization by comparing each received symbol to its expected value,

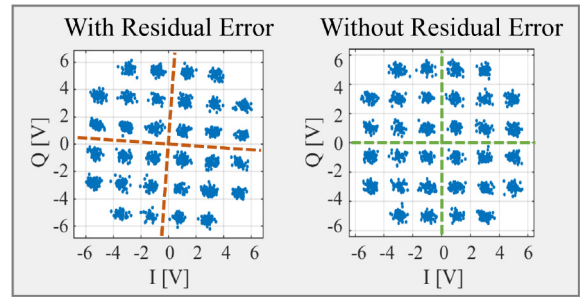


FIGURE 4. Comparison of 32-QAM constellation with and without residual equalization error.




FIGURE 5. Channel estimation using linear interpolation between pilot carriers.

calculating: $C[k] = E[k]/Y[k]$ where $C[k]$ is the required equalization, $E[k]$ is the expected symbol, and $Y[k]$ is the received symbol. A drawback of this approach is that channel information is only gathered when the sequence is sent. With noise and jitter being predominant impairments causing constellations to scale and rotate, the required equalization is continually changing. In order to track these changes, this method requires frequent training sequences; however, this incurs a large overhead.

2) USE OF PILOT CARRIERS

Similar to the previous method, this approach also transmits known symbols to uncover the required equalization. However, unlike the previous method, it permanently reserves a fraction of the bins for this purpose. These transmitted tones are called pilot carriers and are often spaced evenly across the bandwidth [21]. As such, an estimate of the link’s magnitude and phase response is gathered continuously; however, the receiver lacks information from unreserved bins. Assuming a relatively smooth channel response, this information is filled in by interpolating between adjacent pilots. However, this assumption is not always valid, especially for channels that experience significant time-varying discontinuities in their frequency response [21].

Fig. 5 depicts an example of magnitude interpolation with a discontinuity at bin 6. A similar figure can be created for phase interpolation. Noticeably, this approach experiences two drawbacks. First, interpolation does not approximate the channel perfectly; it will produce equalization error. And second, areas with severe discontinuities require additional pilots, lowering spectral efficiency.




FIGURE 6. Performing linear regression over rotated pilot carriers to estimate sampling phase error.

Nevertheless, pilot-carrier-based channel estimation is the predominant alternative in many standards, including the IEEE 802.16a DVB-T [14], [21] and IEEE 802.11a WLAN [23]. In Section IV, we propose a solution that overcomes the trade-off between channel estimation accuracy and spectral efficiency.

C. TIMING RECOVERY

Timing recovery deals with finding the clock frequency and phase that correctly positions the FFT window (coarse synchronization) and minimizes jitter (fine tracking) [19]. Three common approaches are: use of training sequence, use of cyclic prefix, and use of pilot carriers [23].

1) USE OF TRAINING SEQUENCE

This method periodically sends a training sequence. A cross-correlation of the expected sequence with the received signal creates a peak at the ideal window position [21]. Similar to training-sequence-based channel estimation, sampling phase error is only measured when the sequence is sent. As such, it is typically reserved for coarse synchronization as it is limited in tracking bandwidth. Enabling high-frequency jitter tracking requires more frequent training sequences, which incurs additional overhead.

2) USE OF CYCLIC PREFIX

This approach performs an auto-correlation with a lag of $2N$ to identify the repetition created by the CP and locate the beginning of each frame [21]. Assuming the CP is already present, this method does not incur additional overhead and gathers information continuously. However, it is susceptible to phase-error detection unreliability when severe ISI and low SNR are present [14], [21], [24]. Although reliability can be improved by time averaging, this lowers the jitter tracking bandwidth [20]. As such, this method is also typically reserved for coarse synchronization.

3) USE OF PILOT CARRIER

This method uses pilot carriers to measure timing error. When the sampling phase deviates from its ideal position, constellations rotate [21]:

$$Y[k, \theta_{err}] = e^{-j2\pi k\theta_{err}/2N} \cdot Y[k] = e^{j\phi_{err}[k]} \cdot Y[k]. \quad (3)$$

θ_{err} is the phase error normalized to the sampling period. By monitoring the rotation error ϕ_{err} , a plot similar to Fig. 6

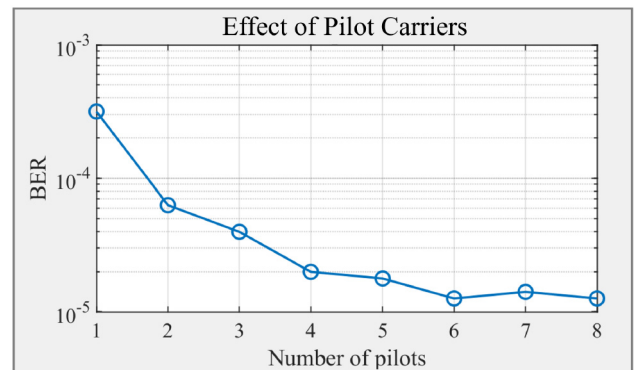


FIGURE 7. Performance impact from averaging over too few pilot carriers.

is generated. For a given drift in phase, higher-frequency pilots rotate more than lower ones. Taking a linear regression of these points and measuring the slope uncovers an accurate estimate of the sampling phase error.

This process has similarities to pilot-carrier-based channel estimation; however, the former monitors the required equalization on a bin-by-bin basis, the latter monitors the average group delay of all bins. Therefore the former applies interpolation, whereas the latter applies linear regression.

This approach is more accurate than the previous two alternatives, making it ideal for both coarse synchronization and fine jitter tracking [14]. The accuracy is analyzed in Section VI. However, the approach also has a drawback. There exists a trade-off between tracking accuracy and spectral efficiency. Additional pilots allow for more averaging in the frequency-domain, leading to more accurate estimates when subject to noise and jitter; however, reserving bins incurs additional overhead.

Pilot-carrier-based timing recovery is the predominant method for applications including IEEE 802.16a DVB-T and IEEE 802.11a WLAN as it achieves the best performance [21], [23]. However, the overhead associated with bin reservation is detrimental in bin-limited systems, affecting its practicality for the target application.

III. INVESTIGATING THE OVERHEAD

To investigate the effect of varying the number of pilot carriers, a 32-point FFT system-level model is developed and simulated. The model includes noise, crosstalk, quantization from a 7-bit 32GS/s ADC, attenuation from a 27dB channel at 16GHz [25], $0.01U_{rms}$ of random jitter, and $0.02U_{IP}$ of Dual-Dirac jitter (set to meet the CEI 56Gb/s jitter specification for channel compliance testing using reference transmitter and receiver [26]). Furthermore, the receiver can tolerate an additional $2U_{IPP}$ of sinusoidal jitter at 0.5MHz. Results are shown in Fig. 7.

As expected, decreasing the number of pilots worsens the Bit-Error-Rate (BER). With fewer than four pilots, the lack of averaging across bins impacts the reliability of the phase-error detection in the presence of impairments [14]. This sets the limit for proper operation. With this assumption, if each bin carries approximately the same number

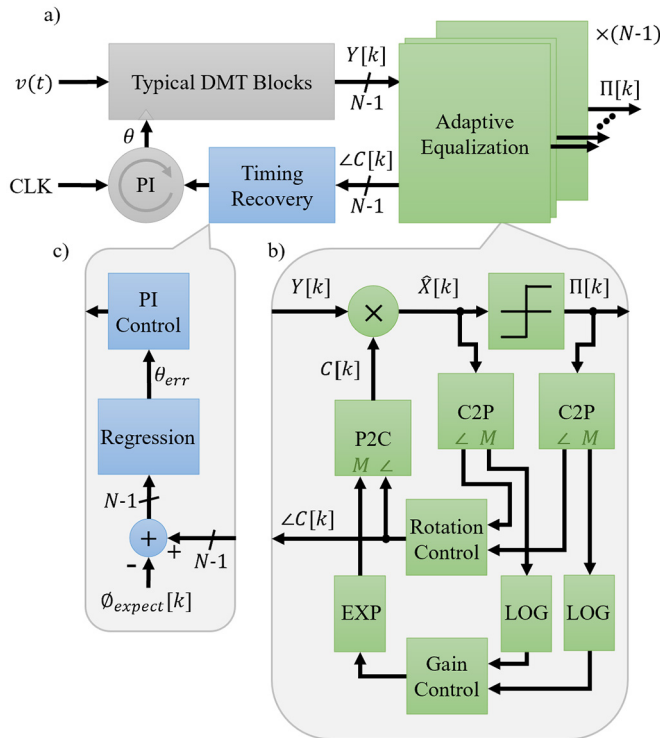


FIGURE 8. Proposed receiver block diagram: a) top-level, b) equalizer, c) timing recovery.

of bits, the system will experience a data-rate reduction of 27%. Considering this penalty, DMT becomes a less attractive alternative to 4-PAM signalling. Evidently, an alternative approach is needed with less overhead but similar tracking performance. Such a solution is introduced in the following section.

IV. PROPOSED DMT RECEIVER

A combination of adaptive decision-directed equalization and frequency-domain timing recovery is proposed, which does not require bin reservation yet outperforms the conventional pilot-carrier-based technique. The receiver is depicted in Fig. 8. In gray are conventional DMT blocks, in green is the proposed adaptive equalizer, and in blue is the proposed timing recovery.

At a high level, shown in Fig. 8a, the received signal $v(t)$ passes through typical DMT blocks, including an ADC, CP removal, and an FFT. Each frame sends its symbols $Y[k]$ to a set of $N - 1$ adaptive equalizers, one per bin, to correct scaling and rotation errors and produce final decisions $\Pi[k]$.

Each equalizer, shown in Fig. 8b, multiplies the input symbol $Y[k]$ by a complex coefficient $C[k]$ whose value is determined during the previous frame period, to produce the equalized symbol $\hat{X}[k]$. $\hat{X}[k]$ is then sliced to produce the final decision $\Pi[k]$, i.e., rounded to its nearest constellation point following (4) where S is the set of constellation points [19].

$$\Pi[k] = \arg \min_{s \in S} |\hat{X}[k] - s|. \quad (4)$$

The feedback path separates the magnitude and phase from symbols $\hat{X}[k]$ and $\Pi[k]$. This is accomplished using two cartesian-to-polar (C2P) converters. The two extracted magnitude values are sent to the gain controller and the two phase values to the rotation controller. Both controllers implement proportional-integral control. The outputs $|C[k]|$ and $\angle C[k]$ are converted back to cartesian form using a polar-to-cartesian (P2C) converter to produce the next frame's equalization value $C[k]$. In addition, the gain controller implements $\log_2(\cdot)$ (LOG) and $2^{(\cdot)}$ (EXP) converters; the reasoning will be discussed shortly.

The timing recovery, shown in Fig. 8c, extracts $N - 1$ rotation equalization values $\angle C[k]$ from the equalizers. These rotations are subtracted by $N - 1$ offset values, to be described later. The resultant rotation error $\phi_{err}[k]$ is sent to a linear regression block to estimate the sampling phase error θ_{err} . The error is input into a proportional-integral controller, which controls a phase interpolator and adjusts the sampling phase of the ADC.

We now provide more details on the proposed equalizer and timing recovery.

A. PROPOSED EQUALIZER

The proposed equalizer of Fig. 8b expands on work done by [19], [27], [28]. Instead of using pilot carriers, this approach performs decision-directed channel estimation. It determines the channel response by directly observing the scaling and rotation error of data-filled constellations. However, the proposed equalizer implements two features that differentiates it from previous works. First, it applies the CORDIC algorithm to separate magnitude error from rotation error. As mentioned, information from the rotation error is processed to produce accurate sampling phase error estimates. This combines equalization and timing recovery into a single system and achieves high performance without incurring any data rate overhead. Second, it implements the gain controller in the logarithmic domain. This removes amplitude-dependent stability allowing for a higher loop gain and thus a higher noise tracking bandwidth. Both features will be discussed shortly.

At startup, a training sequence is used to synchronize initial coefficients. At run-time, the equalized symbol $\hat{X}[k]$ is compared to the sliced symbol $\Pi[k]$. As illustrated by Fig. 9, $\hat{X}[k]$ is received at the location of the bullet. It is sliced to the nearest constellation point at the location of the **X**. A comparison of the two locations uncovers a ratio of magnitude error and a difference of phase error. Following (5) and (6), two control loops, one for gain and another for rotation, correct the error by adjusting $C[k]$. The control loops eventually drive the gain ratio to 1 and the phase difference to 0, ensuring symbols land in the center of their decision boundaries.

$$M_{err}[k] = \frac{|\Pi[k]|}{|\hat{X}[k]|} \rightarrow 1 \quad (5)$$

$$\phi_{err}[k] = \angle \Pi[k] - \angle \hat{X}[k] \rightarrow 0. \quad (6)$$

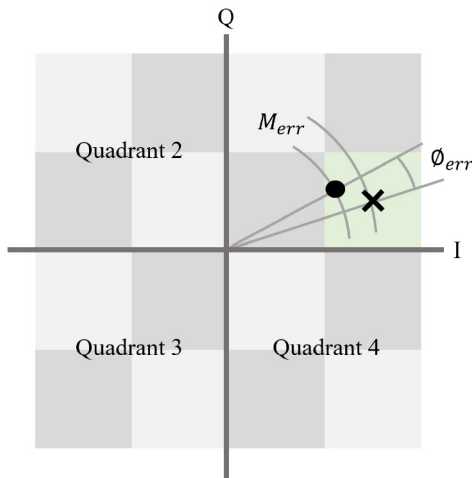


FIGURE 9. Decision-directed adaptive equalization comparing symbols before and after slicer.

Three techniques are adopted to minimize the complexity of the equalizer. First, the C2P and P2C converters implement the CORDIC algorithm [29], [30]. This uses a shift-and-add methodology to avoid multiplication. Second, the gain control loop attempts to drive the ratio of amplitudes towards unity on a symbol-by-symbol basis. This is different than a traditional AGC in which the goal is to drive the output level to a desired fixed value. In this work, we wish to control the ratio as opposed to the difference of the two amplitudes as the latter will create a nonlinear loop with stability concerns. This will require it to be severely over-damped and slow. To control the ratio, we have opted to work in the logarithmic domain as it results in less implementation complexity. As such, two LOG converters and an EXP converter are placed before and after the controller. The logarithmic and exponential converters also use a similar shift-and-add methodology [31]. And third, the C2P converter following the slicer and its corresponding LOG converter are implemented using lookup tables. This further reduces complexity since the number of possible input values is limited by the constellation size.

The proposed approach to channel estimation introduces two benefits. First, it does not need to reserve bins. As such, spectral efficiency is maximized. And second, it directly measures the equalization error of each bin. This error is driven to zero when integrated by the proportional-integral controllers, eliminating residual equalization error.

B. PROPOSED TIMING RECOVERY

All pilot-carrier-based timing recovery methods observe the rotation of received pilots. For this reason, those bins must send a known symbol and thus cannot contain random data. This paper's solution is to obtain timing information by monitoring the required rotation equalization of data-filled bins instead of the rotation itself.

At startup, a training sequence is used to synchronize equalization. Then the timing recovery enters synchronization mode, where it observes the lowest bin. By rotating the phase interpolator, it attempts to remove any required rotation equalization. Doing so achieves coarse frequency and phase synchronization.

At run-time, data is sent and the timing recovery enters normal operation mode where it observes the rotation equalization of all bins; this permits fine tracking. If phase drift is encountered, constellations will try to rotate. Referring to Fig. 6, we observe that for the same amount of sampling phase error, higher bins rotate more than lower ones. The adaptive equalizers compensate for this phenomenon by applying appropriate equalization, ensuring they stay square. Taking a linear regression of the rotation equalization produces an accurate estimate of the timing offset θ_{err} . This error is sent to a proportional-integral controller, which adjusts a phase interpolator and corrects the sampling phase of the ADC.

In the presence of frequency offset, the timing recovery observes a continuous shift in sampling phase error. As such, it repeatedly corrects this error by rotating the phase interpolator, in turn tracking the offset. The maximum frequency offset is analyzed in Section V.

The proposed timing recovery has the option to shift the target sampling phase of the receiver. Shifting the phase early allows the CP to correct both pre- and post-cursor ISI. This is done by adjusting the offset coefficients, changing the target rotation equalization to: $\phi_{expect}[k] = 2\pi kO/2N$ where O is the desired offset in UI.

This approach to timing recovery introduces two benefits. First, pilot carriers are no longer required. Therefore spectral efficiency is maximized. And second, unlike the conventional approach that estimates phase drift by averaging a finite number of bins, this solution uses information from every bin, providing a more accurate estimate and reducing sensitivity to impairments. As explained in Section VI, this approach overcomes the trade-off between timing-recovery accuracy and spectral efficiency, it maximizes both.

The combination of the proposed adaptive equalizer with pilot-carrier-less timing recovery lends to one additional improvement. Phase error is typically detected after the FFT in the Frequency Domain (FD) by measuring pilot rotation but is corrected before the FFT in the Time Domain (TD) by adjusting the ADC sampling phase. The loop latency is a limiting performance factor to timing recovery, as signals must convert from TD to FD and back to TD. The controller must wait for changes and thus cannot adjust quickly. Instead, the proposed solution partially corrects the effects of jitter directly in the FD by continuously updating equalization coefficients. This approach tracks low-frequency drift at the input sampler in the TD and high-frequency oscillations at the equalizer in the FD, bypassing much of the delay and minimizing constellation movement [27]. As will be shown next, this advantage improves the jitter tracking ability of the receiver.

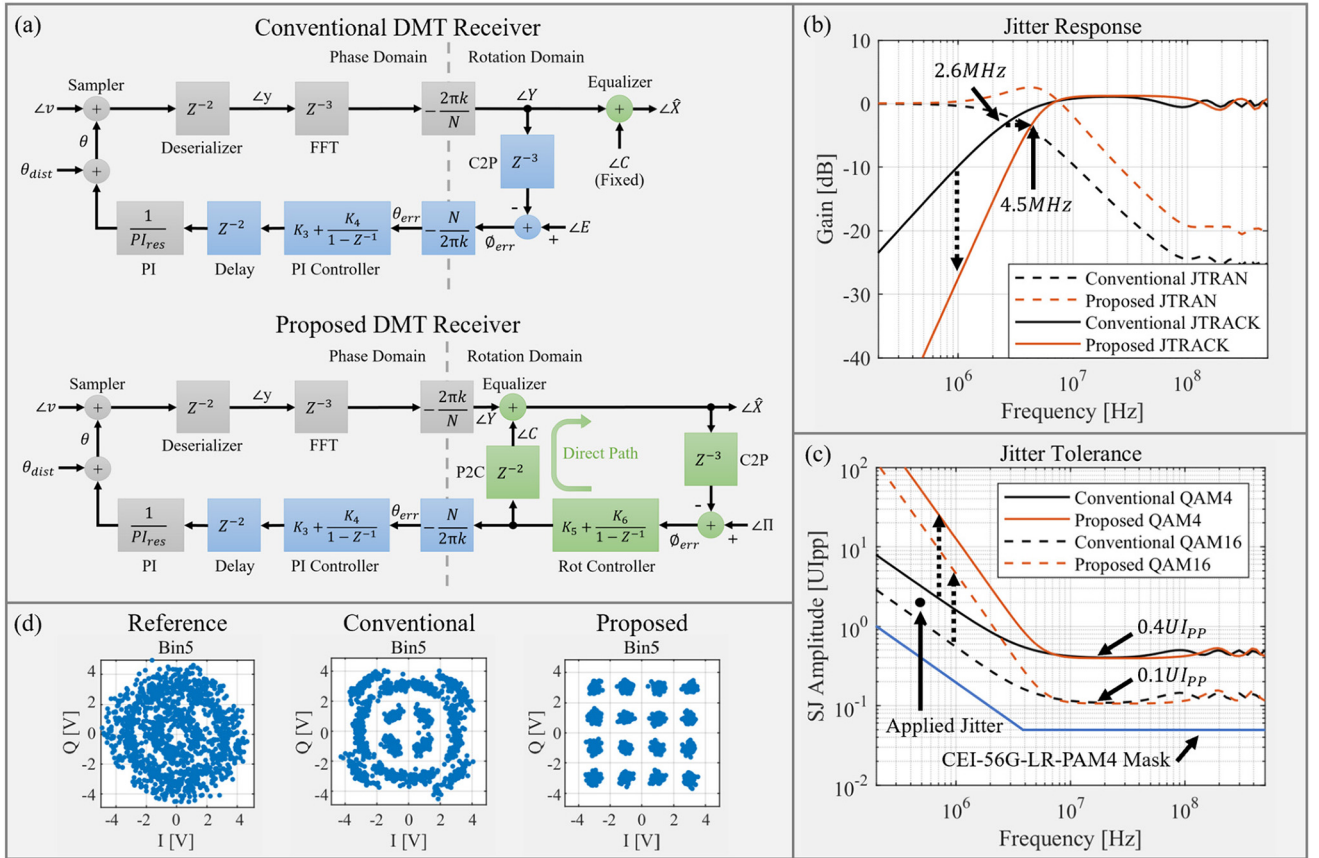

FIGURE 10. Performance comparison of typical versus proposed system.

TABLE 1. Complexity of proposed adaptive equalization and timing recovery.

Module	Block	Multipliers	Adders	Compares	LUTs
Equalizer (1 instance)	Complex Multiply	3	5	0	0
	Slicer	0	0	0	0
	C2P1	0	24	8	0
	C2P2	0	0	0	2
	LOG1	0	12	8	0
	LOG2	0	0	0	1
	EXP	0	12	8	0
	P2C	0	24	8	0
	Gain Control	2	3	0	0
Timing Recovery	Rotation Control	2	3	0	0
	Offset	0	1	0	0
	Regression	15	14	0	0
	PI Control	2	3	0	0

C. COMPLEXITY

The complexity for the proposed equalizer and timing recovery is shown in Table 1. The final implementation assumes 8-bit arithmetic and employs 8-stages for the C2P, P2C, LOG and EXP converters.

V. SYSTEM MODELING, SIMULATION, AND FPGA IMPLEMENTATION

A. SYSTEM MODELING

In this section, two linear models, shown in Fig. 10a, are created to enable a fair jitter performance comparison

between a conventional pilot-carrier-based timing-recovery system and the proposed solution. Both represent a 32-point FFT system with a 32GS/s ADC. For this comparison, the models are set in the phase and rotation domains corresponding to the time and frequency domains, respectively. As such, they can be simplified to only include gain and delay elements; these are the primary contributing factors to jitter performance. For example, the FFT is represented by a delay and a scaling factor. This assumes that every bin observes the same phase disturbance proportional to its bin number, and thus it is only necessary to model a single bin. The highest bin is modeled as it is the most sensitive. Delay elements are added to meet timing requirements for a synthesized 16nm FinFET process. Gain elements model the proportional-integral controllers and transition signals between either domain.

In the conventional model, sinusoidal jitter $\angle v$ is input to the system. A phase interpolator, modeled by the addition of θ , attempts to minimize this disturbance. The result is sent to the deserializer, the FFT, and transitioned to the rotation domain $\angle Y$. There, fixed rotation equalization $\angle C$ is added, and the signal $\angle \hat{X}$ is outputted from the system. In the feedback path, rotation $\angle Y$ is measured following the FFT and compared to an expected phase $\angle E$. This simulates measuring the rotation of a pilot carrier. The error ϕ_{err} is converted to the phase domain θ_{err} , sent through a proportional-integral

controller, scaled through the phase interpolator θ , and added to the input signal $\angle v$. Furthermore, disturbance θ_{dist} including $0.01UI_{rms}$ of random jitter and $0.02UI_p$ Dual-Dirac jitter is added to model typical PLL/VCO phase noise.

The proposed model is similar to the conventional; however, adaptive rotation equalization $\angle C$ is added at the output of the FFT before leaving the system. In the feedback path, the equalizer output $\angle \hat{X}$ is compared to the slicer $\angle \Pi$, whose value is assumed to be constant. The error ϕ_{err} is then sent through the first proportional-integral block, which models the rotation controller. The gain controller does not contribute to jitter tracking and is not included. At the output, one path directly controls the equalization coefficient $\angle C$. The other path is converted to the phase domain θ_{err} , where it enters the second proportional-integral controller and completes the loop by passing through the phase interpolator θ and adding to the input signal $\angle v$.

The timing recovery loop latency for both models is 10 clock periods. This delay is technology-specific and depends on the DSP clock frequency and arithmetic precision. The latency is split between 2 periods for the deserializer, 3 periods for the FFT, 3 periods to extract rotation error measurements from the received symbols, and 2 periods for the linear regression and for producing a sampling phase adjustment. Since both models require the same calculations, the phase tracking loop latency is identical, irrespective of the technology. However, introducing adaptive equalization adds a direct path in the proposed model whose loop has half the latency (5 clock periods). This reduction is achieved by bypassing the delay from the deserializer, the FFT, and from producing phase adjustments. Instead, it introduces only 2 clock periods to produce equalization adjustments. As will be shown next, this change improves the timing recovery bandwidth. Furthermore, but not shown, the gain control loop delay is 8 clock periods, 3 more than the rotation control loop to account for the LOG and EXP conversion.

The input-to-output transfer functions of both models, $F_1(z)$ and $F_2(z)$, are described by (7) and (8) respectively where $2N = 32$ is the FFT size, $k = 15$ is the observed bin number, $PI_{res} = 64$ is the PI resolution, and the loop gain coefficients are set to $K_1 = 3.0$, $K_2 = 1.2$, $K_3 = 9.0$, $K_4 = 1.5$, $K_5 = 0.08$, and $K_6 = 0.04$.

These equations extract two important curves: the Jitter Transfer Function (JTRAN) and the Jitter Tracking Function (JTRACK) [32]. JTRAN is typically a low-pass function describing the amount of jitter transferred from the input to the recovered clock. In contrast, JTRACK is typically a high-pass function describing the amount of jitter remaining in the recovered data. Eqn. (9) describes either curve.

$$\begin{aligned}
 F_1(z) &= \frac{A_1}{1 - A_1 B_1} \\
 A_1(z) &= -\frac{2\pi k}{2N} \cdot Z^{-5} \\
 B_1(z) &= \frac{K_1 + K_2/(1 - Z^{-1})}{PI_{res}} \cdot \frac{2N}{2\pi k} \cdot Z^{-5} \quad (7)
 \end{aligned}$$

$$\begin{aligned}
 F_2(z) &= \frac{A_2 F'_2(z)}{1 - A_2 F'_2(z) B_2 C_2} \\
 F'_2(z) &= \frac{1}{1 - C_2 D_2} \\
 A_2(z) &= -\frac{2\pi k}{2N} \cdot Z^{-5} \\
 B_2(z) &= -\frac{K_3 + K_4/(1 - Z^{-1})}{PI_{res}} \cdot \frac{2N}{2\pi k} \cdot Z^{-2} \\
 C_2(z) &= -\left(K_5 + K_6/(1 - Z^{-1})\right) \cdot Z^{-3} \\
 D_2(z) &= Z^{-2} \quad (8)
 \end{aligned}$$

$$\begin{aligned}
 JTRAN_1(z) &= \frac{F_1(z)}{A_1(z)} \\
 JTRAN_2(z) &= \frac{F_2(z)}{A_2(z) F'_2(z)} \\
 JTRACK_1(z) &= 1 - JTRAN_1(z) \\
 JTRACK_2(z) &= 1 - JTRAN_2(z). \quad (9)
 \end{aligned}$$

B. SIMULATION RESULTS

Fig. 10b depicts simulation results from a jitter response analysis. Two improvements are visible. First, the 3dB tracking frequency of the JTRACK curve increases by 1.7 times, from 2.6MHz to 4.5MHz. The reason for this improvement is the addition of the direct path, which halves the loop latency from 10 clock periods to 5 and permits a higher bandwidth while maintaining the same phase margin. Second, the proposed method experiences a steeper JTRACK drop-off at low-frequency due to the addition of the second loop, which changes JTRACK from a first-order to a second-order system with twice the roll-off.

Fig. 10c displays a jitter tolerance analysis. Deterministic jitter is applied whose amplitude, initially small, is increased until the highest bin rotates beyond the decision boundary of its constellation and the system's BER exceeds 10^{-6} . The disturbance frequency is swept up to the DSP Nyquist frequency. The test is repeated, once using 4-QAM constellations and again using 16-QAM. As a reference, the CEI-56G-LR-PAM4 jitter tolerance mask is shown [26]. We observe that the proposed system outperforms at low-frequency, where its tolerance increases with twice the slope of the conventional system. Again, this is thanks to the addition of the second loop. Furthermore, one key advantage of multi-tone systems is the ability to adjust the jitter tolerance. As shown in Fig. 10c, when the constellation size is increased, the tolerance is decreased, and vice versa. Therefore, by adjusting the bit-loading, it is possible to meet any practical jitter tolerance requirement. A theoretical limit of $N/2 = 8UI_{pp}$ is attained when allocating 4-QAM to the first bin and leaving the remaining bins empty. This is further discussed in Section VI.

Fig. 10d shows results from a top-level simulation. Here, the model from Section III is reused. The applied $2UI_{pp}$ of sinusoidal jitter at 0.5MHz is marked by a bullet in Fig. 10c, showing it exceeds the tolerance of the conventional 16-QAM model but meets the tolerance of the proposed model.

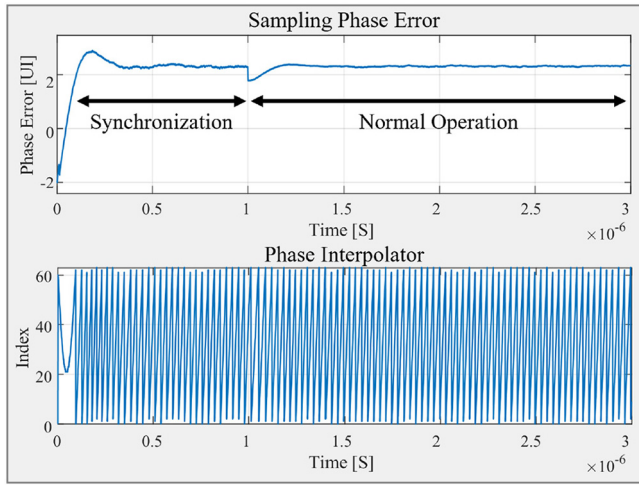


FIGURE 11. Timing recovery convergence when in the presence of a 0.1% frequency offset.

The left plot of Fig. 10d shows an output constellation when timing recovery is inactive. The center plot activates the timing recovery but leaves equalization fixed. Finally, the right plot implements the proposed system with adaptive equalization enabled. From left to right, the measured SNR is 12.4dB, 22.2dB, and 36.1dB. Thus, the right-most plot achieves the highest signal-to-noise ratio by more than 13dB; this validates the improvements found in the previous two analyses.

C. FREQUENCY OFFSET

Fig. 11 depicts the timing recovery convergence when the model from Section III is subject to its maximum frequency offset of 0.1% relative to the sampling frequency. The top plot shows the estimated sampling phase error, whereas the bottom shows the rotating phase interpolator. At startup, the timing recovery observes only the lowest bin to obtain coarse frequency and phase synchronization. At runtime, all bins are observed to obtain fine frequency and phase tracking. Following a slight correction, the timing recovery maintains a constant sampling phase error. This residual error can be removed by offsetting the receiver's ideal sampling phase. The timing recovery's convergence is primarily limited by the rotation speed of the phase interpolator.

D. FPGA IMPLEMENTATION

The proposed receiver is synthesized at gate level. A pipeline topology is used to meet timing requirements. Delay elements are placed between each pipeline stage to split the computation. The same delays have been modeled in Fig. 10a.

Following this, an Altera Cyclone V Field Programmable Gate Array (FPGA) is used to test the design in near real-time. Due to resource limitations, the FPGA could not fit the complete receiver. Instead, the design was tested in sections, including the deserializer, the FFT, and the adaptive equalizer. Timing recovery could not be implemented since




FIGURE 12. Adaptive equalizer test setup for FPGA implementation.

the chip does not contain a phase interpolator. Nevertheless, implementing portions of the design in an FPGA helps verify their functionality. As discussed next, it is possible to evaluate the likelihood of adaptive equalization instability in near real-time, helping characterize its robustness to impairments.

The test setup for the adaptive equalizer FPGA implementation is depicted in Fig. 12. A signal generator is implemented in Verilog to create symbols from an ideal constellation. Magnitude and phase distortion is applied by multiplying each symbol by a complex coefficient controlled by two potentiometers, one for magnitude and the other phase. The impaired symbols are fed to an instance of the adaptive equalizer running at 12.5MHz. At the output, symbols are compared before distortion is applied and after equalization. The symbol error rate is monitored to determine performance. Observing error-free operation signifies that the equalizer is able to track the disturbance and correct it. Assuming a 32GS/s ADC, the actual receiver is expected to operate with a 1GHz clock, 80× faster than the FPGA. However, since the relative loop delays of either implementation are the same, the stability behavior can still be characterized.

After synchronization, the potentiometers are adjusted to mimic a large disturbance. Results from this test observed a rapid reconvergence of the adaptive equalizer to the correct orientation without incurring error. As such, this test validates the equalizer's ability to reconverge following a disturbance. A more extensive stability analysis is conducted in the following section.

VI. CONVERGENCE ANALYSIS OF THE CONTROL LOOPS

A convergence analysis is performed in this section both for the adaptive equalizer and the overall timing recovery.

A. ADAPTIVE EQUALIZER STABILITY

The proposed adaptive decision-directed equalizer is susceptible to error propagation. At run-time, if the decision function produces a wrong symbol, coefficients will change in the wrong direction, making future errors more likely. This has the potential for a loss of synchronization. We now analyze the equalizer's susceptibility to positive feedback by analyzing its ability to re-converge in the presence of a step disturbance.

16-QAM and 128-QAM Decoder-Success-Rate (DSR) versus equalization error simulations are depicted in Fig. 13a and Fig. 13b respectively. This determines the maximum

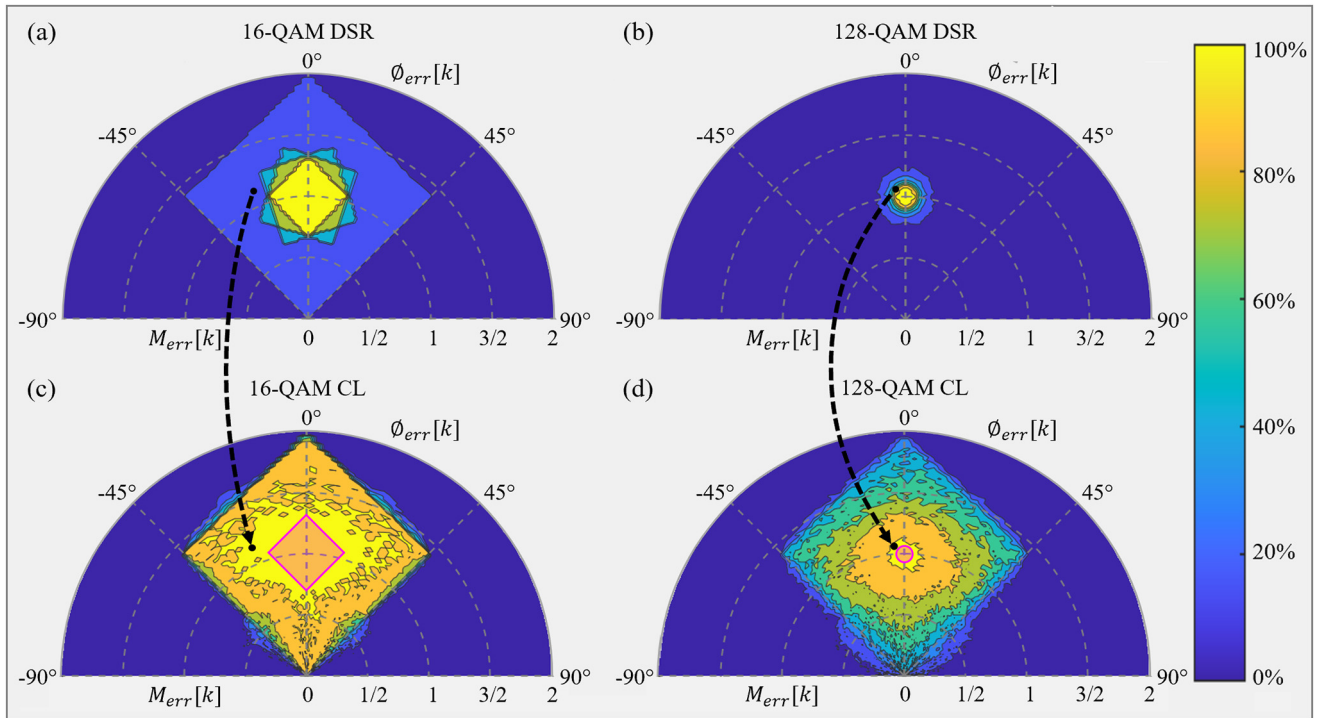


FIGURE 13. Decoder Success Rate (DSR) and Convergence Likelihood (CL) versus equalization error.

equalization error for the respective constellation. In each simulation, an ideal constellation is given rotation $\phi_{err}[k]$ and scaling $M_{err}[k]$ error. This causes some symbols to land outside their decision boundary. The percentage of correctly decoded symbols is recorded and shown as a contour plot. As expected, larger constellations have a narrower error-free region. 16-QAM can operate error-free within $\pm 15^\circ$ and $\pm 25\%$ scaling, whereas 128-QAM can only operate error-free within $\pm 5^\circ$ and $\pm 8\%$ scaling.

16-QAM and 128-QAM Convergence Likelihood (CL) simulations are depicted in Fig. 13c and Fig. 13d respectively. These simulations have a similar setup to the FPGA implementation, shown in Fig. 12, and determine the effect of error propagation. Once again, random symbols from an ideal constellation are generated. Then a step disturbance adds rotation $\phi_{err}[k]$ and scaling $M_{err}[k]$ error to the symbols as they are input to an instance of the equalizer. The disturbance affects the decisions made by the equalizer, causing a potential for error-propagation and convergence to an incorrect orientation. Following a sufficiently long settling time, the success rate of proper convergence is recorded and averaged over one hundred simulations per data point, creating a statistical result.

As expected, if the disturbance remains within the pink shaded region where the decoder performs error-free, the equalizer has a 100% chance of proper convergence. Moreover, the equalizer can still converge correctly if the disturbance goes beyond the decoder's error-free region. For example, as shown by the arrows, the system has perfect convergence even with a DSR of 25% and 60% for a 16-QAM and 128-QAM constellation, respectively.

As mentioned in Section II, DMT systems employ channel optimization algorithms to determine the constellation size for each bin [2], [8]. Constellations begin small and grow incrementally until the channel capacity is reached [15]. The Bit-Error-Rate (BER) of the system dictates this optimization. Assuming Gray mapped constellations, the symbol error rate (SER) (the inverse of DSR) is related to the BER, proportional to the number of encoded bits within the constellation. As such, for a 128-QAM constellation targeting a BER of 10^{-6} , the algorithm will ensure less than seven symbols in a million land incorrectly. Since the system will remain converged up to an SER of 40%, it will also remain converged while targeting practical SER such as 7×10^{-6} . Thus, the concern for loss of synchronization is low, assuming the system operates with a typical SER.

B. CLOCK SYNCHRONIZATION CONVERGENCE

The proposed timing recovery operates by measuring the rotation of constellations. At run-time, if a constellation experiences a large rotation creating errors in the decision function, the timing recovery unit will produce incorrect sampling phase error estimates, and the clock could lose synchronization. We now analyze the timing recovery's susceptibility to jitter by analyzing its ability to re-converge in the presence of a step disturbance.

For this analysis, the same 32-point FFT system from previous sections is simulated where the feedback loops controlling both the equalization and the sampling phase are disconnected. Then phase error is manually applied from $-8UI$ to $+8UI$ producing phase-detection curves shown

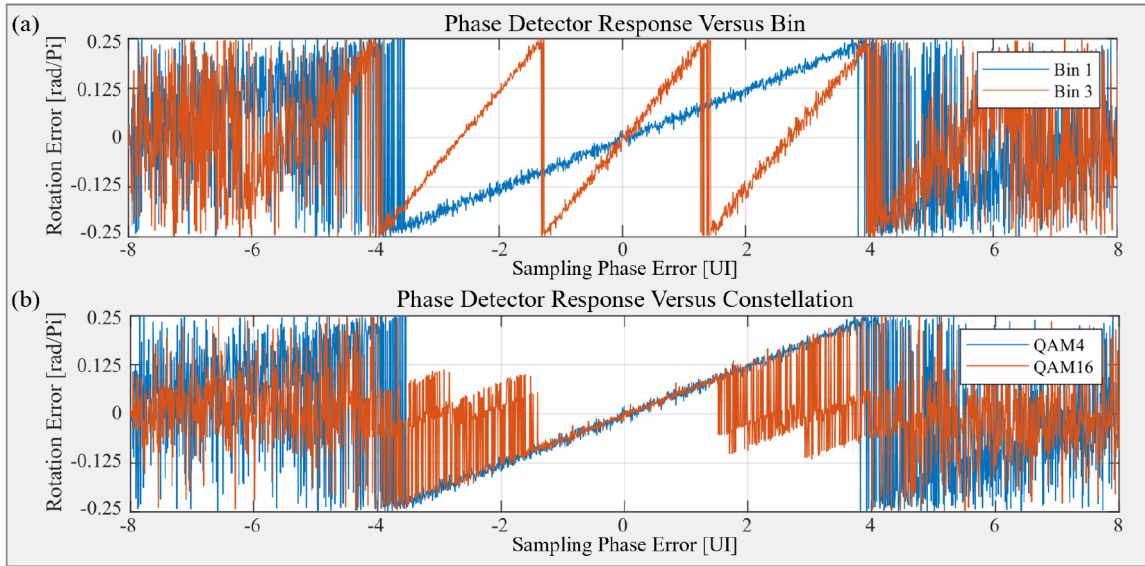


FIGURE 14. Sampling phase error detection curves.

in Fig. 14. Fig. 14a compares the phase-detection for different frequency bins while using 4-QAM constellations whereas Fig. 14b compares the phase-detection for different constellation sizes while observing the lowest frequency bin.

From Fig. 14a, the convergence region is limited to the linear portion of the detection curve prior to hitting a discontinuity. With bins 1 and 3 both containing a 4-QAM constellation, this region is limited to $\pm 4UI$ and $\pm 4/3UI$, respectively. As such, lower frequency bins experience a larger region of convergence but measure less rotation error for the same sampling phase error, producing a more coarse estimate.

This convergence region $Conv[k]$ in UI is described by (10) and the measurement sensitivity $\theta_{LSB}[k]$ in UI defining the smallest measurable change in phase is described by (11). We assume rotation is measured using R -bit arithmetic. As explained in Section IV, a more reliable phase error estimate is achieved by averaging over multiple bins. This improved sensitivity θ_{AVG} is described by (12) where P is the set of reserved bins, and N_P is the number of reserved bins.

$$Conv[k] = \pm \frac{2N/8}{k} \quad (10)$$

$$\theta_{LSB}[k] = \frac{2N}{k * 2^R} \quad (11)$$

$$\theta_{AVG} = \frac{1}{N_P} \sqrt{\sum_{k \in P} \left(\frac{\theta_{LSB}[k]}{k} \right)^2}. \quad (12)$$

Assuming a 32-point FFT system with 8-bit arithmetic, a conventional pilot-carrier-based approach limited to 4 pilot carriers in bins 1, 5, 10, and 15 achieves a sensitivity of $3.1 * 10^{-2}UI$. In contrast, the proposed system, after averaging across all 15 bins, achieves a sensitivity of $8.7 * 10^{-3}UI$. This is an improvement of 3.6 times compared to the alternative.

Fig. 14b confirms that larger constellations experience their discontinuity sooner than smaller ones and thus have a reduced convergence region. Symbols further from the origin travel

more distance for the same amount of rotation. As such, they reach their decision boundary sooner. This effect can be modeled with a convergence penalty set by results in Fig. 13.

The above two analyses consider the system at run-time. Thus, the convergence is limited by the decision boundary of constellations. However, as demonstrated in Fig. 11, at startup, when a training sequence is sent and when only the lowest bin is observed, the convergence region is extended to $\pm N = \pm 16UI$, guaranteeing initial synchronization.

In summary, the timing recovery convergence and sensitivity are dictated by the bin number and the constellation size. The proposed receiver can take advantage of this by selecting which bins to observe for timing recovery. Assuming the use of 4-QAM, observing only the lowest bin trades-off a sensitivity of $0.13UI$ to achieve a convergence of $\pm 4UI$. Alternatively, observing every bin trades-off a convergence of $\pm 0.27UI$ to achieve a sensitivity of $8.7 * 10^{-3}UI$. As such, when significant jitter is present, susceptibility to loss of synchronization can be minimized by 15 times. On the other hand, phase error sensitivity can be improved by 15 times when jitter is of little concern. This control is not possible using conventional pilot-carrier-based techniques.

VII. CONCLUSION

This paper proposes a discrete multi-tone timing-recovery technique with adaptive equalization that overcomes two fundamental trade-offs: channel estimation accuracy versus spectral efficiency and timing recovery accuracy versus spectral efficiency. Together, this overcomes the 27% overhead prominent when using four pilot carriers and a 32-point FFT. Two simplified timing-recovery models are examined and used to compare the proposed solution to its alternative. A complete system-level model is created and used to evaluate the link's overall performance. The design is synthesized at gate level, and implemented in an FPGA, where it is run in near real-time to validate functionality.

Results show improvement in phase-error sensitivity by 3.6 times, tracking bandwidth by 1.7 times, an increase in the jitter tolerance slope by 20dB per decade at low-frequency and the ability to remove residual equalization error. Convergence analysis of both the adaptive equalizer and clock synchronization shows the system's ability to mitigate error propagation and remain synchronized in the presence of impairments. Finally, we highlight the system's ability to trade-off convergence versus sensitivity. Either parameter can be adjusted by 15 times, optimizing the receiver over a broad range of signal conditions.

ACKNOWLEDGMENT

The authors thank the anonymous reviewers for their valuable feedback on the first draft of this article. They also thank Huawei Canada for their expertise and assistance throughout the course of this research. Access to CAD tools was provided by CMC Microsystems.

REFERENCES

- [1] M. LaCroix, "8.4 a 116GB/S DSP based wireline transceiver in 7nm CMOS achieving 6pj/bit at 45dB loss in PAM-4/DUO-PAM4 and 52dB in PAM-2," in *Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC)*, 2021, pp. 132–134.
- [2] B. Vatankhahghadim, N. Wary, and A. C. Carusone, "Discrete multi-tone signalling for wireline communication," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, 2020, pp. 1–5.
- [3] J. Bailey *et al.*, "8.8 a 112Gb/s PAM-4 low-power 9-tap sliding-block DFE in a 7nm FinFET wireline receiver," in *Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC)*, vol. 64, 2021, pp. 140–142.
- [4] C. E. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.*, vol. 27, no. 3, pp. 379–423, Jul. 1948.
- [5] G. Kim *et al.*, "30.2 a 161mW 56Gb/s ADC-based discrete multi-tone wireline receiver data-path in 14nm FinFET," in *Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC)*, 2019, pp. 476–478.
- [6] K. McCollough, S. D. Huss, J. Vandersand, R. Smith, C. Moscone, and Q. O. Farooq, "A 480Gb/s/mm 1.7pJ/b short-reach wireline transceiver using single-ended NRZ for die-to-die applications," in *Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC)*, 2021, pp. 184–185.
- [7] S. V. D. Heide, N. Eiselt, H. Griesser, J. J. Vegas Olmos, I. T. Monroy, and C. Okonkwo, "Experimental investigation of impulse response shortening for low-complexity MLSE of a 112-Gbit/s PAM-4 transceiver," in *Proc. ECOC 42nd Eur. Conf. Opt. Commun.*, 2016, pp. 1–3.
- [8] 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. 65, pp. 65–77, 2021, doi: 10.1109/OJCS.2020.3041239.
- [9] N. Eiselt *et al.*, "Performance comparison of 112-Gb/s DMT, Nyquist PAM4, and partial-response PAM4 for future 5G Ethernet-based fronthaul architecture," *J. Lightw. Technol.*, vol. 36, no. 10, pp. 1807–1814, May 15, 2018.
- [10] G. Kim, W. Kwon, T. Toifl, Y. Leblebici, and H.-M. 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.
- [11] J. Armstrong, "OFDM for optical communications," *J. Lightw. Technol.*, vol. 27, no. 3, pp. 189–204, Feb. 1, 2009.
- [12] R. Nguyen *et al.*, "8.6 A highly reconfigurable 40-97GS/s DAC and ADC with 40GHz afe bandwidth and sub-35fJ/conv-step for 400Gb/s coherent optical applications in 7nm FinFET," in *Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC)*, 2021, pp. 136–138.
- [13] K. Mizutani, T. Matsumura, and H. Harada, "A comprehensive study of universal time-domain windowed OFDM-based LTE downlink system," in *Proc. 20th Int. Symp. Wireless Pers. Multimedia Commun. (WPMC)*, 2017, pp. 28–34.

- [14] M. Li and W. Zhang, "A novel method of carrier frequency offset estimation for OFDM systems," *IEEE Trans. Consum. Electron.*, vol. 49, no. 4, pp. 965–972, Nov. 2003.
- [15] J. B. Anderson and A. Svensson, *Coded Modulation Systems*. New York, NY, USA: Kluwer Acad., 2003.
- [16] D. K. Kim, S. H. Do, H. B. Cho, H. J. Chol, and K. B. Kim, "A new joint algorithm of symbol timing recovery and sampling clock adjustment for OFDM systems," *IEEE Trans. Consum. Electron.*, vol. 44, no. 3, pp. 1142–1149, Aug. 1998.
- [17] J. Cioffi. "Fundamentals of synchronization." 2019. [Online]. Available: <https://cioffi-group.stanford.edu/doc/book/chap6.pdf>
- [18] S. Shahramian *et al.*, "30.5 a 1.41pJ/b 56gb/s PAM-4 wireline receiver employing enhanced pattern utilization CDR and genetic adaptation algorithms in 7nm CMOS," in *Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC)*, 2019, pp. 482–484.
- [19] J.-H. Kim and W.-W. Kim, "Frequency domain-DFE coupled with common phase error tracking loop in OFDM systems," in *Proc. IEEE 61st Veh. Technol. Conf.*, vol. 2, 2005, pp. 1248–1252.
- [20] B. Ai, Z.-X. Yang, C.-Y. Pan, J.-H. Ge, Y. Wang, and Z. Lu, "On the synchronization techniques for wireless OFDM systems," *IEEE Trans. Broadcast.*, vol. 52, no. 2, pp. 236–244, Jun. 2006.
- [21] D.-C. Chang, "Analysis and compensation of channel correction in pilot-aided OFDM systems with symbol timing offset," in *Proc. IEEE Int. Conf. Electro/Inf. Technol.*, 2006, pp. 324–329.
- [22] M. Zhao, A. Huang, Z. Zhang, and P. Qiu, "All digital tracking loop for OFDM symbol timing," in *Proc. IEEE 58th Veh. Technol. Conf. (VTC-Fall)*, vol. 4, 2003, pp. 2435–2439.
- [23] M. Engels, *Wireless OFDM Systems: How to Make Them Work?* Norwell, MA, USA: Kluwer Acad., 2002.
- [24] H.-D. Lin, T.-H. Sang, and S.-Y. Hsu, "SPC08-6: A histogram-based symbol timing synchronization algorithm for OFDM systems," in *Proc. IEEE Globecom*, 2006, pp. 1–5.
- [25] "IEEE P802.3CK Task Force—Tools and Channels." 2018. [Online]. Available: <https://www.ieee802.org/3/ck/public/tools/>
- [26] "Common electrical I/O (CEI)—Electrical and Jitter Interoperability Agreements for 6G+ bps, 11G+ bps, 25G+ bps I/O and 56G+ bps." 2017. [Online]. Available: <https://www.oiforum.com/wp-content/uploads/2019/01/OIF-CEI-04.0.pdf>
- [27] R. F. Ormondroyd and E. A. Al-Susa, "A high efficiency channel estimation and equalisation strategy for a broadband cofdm system," in *Proc. URSI Int. Symp. Signals Syst. Electron. Conf.*, 1998, pp. 471–475.
- [28] J. Rinne and M. Renfors, "Equalization of orthogonal frequency division multiplexing signals," in *Proc. Commun. Global Bridge (IEEE GLOBECOM)*, vol. 1, 1994, pp. 415–419.
- [29] S. Arar. "An Introduction to the cordic algorithm." AllAboutCircuits.com. 2017. [Online]. Available: <https://www.allaboutcircuits.com/technical-articles/an-introduction-to-the-cordic-algorithm/>
- [30] Gisselquist Technology. "Cordic part two: rectangular to polar conversion." 2017. [Online]. Available: <https://zipcpu.com/dsp/2017/09/01/topolar.html>
- [31] Quinapalus. "Calculate exp() and log() without multiplications." 2019. [Online]. Available: <https://www.quinapalus.com/efunc.html>
- [32] A. Amirkhany, "Basics of clock and data recovery circuits: Exploring high-speed serial links," *IEEE Solid State Circuits Mag.*, vol. 12, no. 1, pp. 25–38, Jan. 2020.



JEREMY COSSON-MARTIN (Graduate Student Member, IEEE) received the B.A.Sc. degree in electrical engineering from Queen's University, Kingston, ON, Canada, in the spring of 2018. In the fall of 2018, he began an M.A.Sc. program at the University of Toronto, ON, Canada, under the supervision of Prof. A. Sheikholeslami. In 2019, he transferred into a Ph.D. program. In the summer of 2018, he joined Huawei Canada, Toronto, ON, Canada, as an intern, where he was involved in creating in-lab measurement scripts for a prototype 56-Gb/s SerDes integrated chip. He also received experience in 7-nm FinFET layout. Currently, he is researching multi-tone schemes for ultra-high-speed wireline applications.



HOSSEIN SHAKIBA (Senior Member, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering from the Department of Electrical and Computer Engineering, Isfahan University of Technology, Iran, in 1985 and 1989, respectively, and the Ph.D. degree in electrical engineering from the Department of Electrical and Computer Engineering, University of Toronto, Canada, in 1997. He has over 35 years of teaching, research, design, and management experience in the area of analog circuit and system design for various

applications with focus on wireline communication in both the industry and academia. He is currently working on system and circuit development for next generation serial links at Huawei Canada in collaboration with the wireline industry with emphasis on link design, modeling, and analysis including statistical and signal integrity. He is also actively involved in conducting research with various universities and co-supervises several graduate students.



ALI SHEIKHOLESAMI (Senior Member, IEEE) received the B.Sc. degree in electrical engineering from Shiraz University, Iran, in 1990 and the M.A.Sc. and Ph.D. degrees in electrical engineering from the University of Toronto, Canada, in 1994 and 1999, respectively.

In 1999, he joined the Department of Electrical and Computer Engineering, University of Toronto, where he is currently a Professor. He was on research sabbatical with Fujitsu Labs in 2005–2006, and with Analog Devices, Toronto, ON, Canada, in 2012–2013. His research interests are in analog and digital integrated circuits, high-speed signaling, and CMOS annealing. He has co-authored over 70 journal and conference papers, 10 patents, and a graduate-level textbook entitled “Understanding Jitter and Phase Noise.”

Dr. Sheikholeslami has received numerous teaching awards, including the 2005–2006 Early Career Teaching Award and the 2010 Faculty Teaching Award, both from the Faculty of Applied Science and Engineering at the University of Toronto. He served on the Memory, Technology Directions, and Wireline Subcommittees of the ISSCC in 2001–2004, 2002–2005, and 2007–2013, respectively. He was an SSCS Distinguished Lecturer in 2018–2019. He currently serves as the Education Chair for ISSCC and the Vice President, Education, for SSCS. He is an Associate Editor for the *Solid-State Circuits Magazine*, in which he has a regular column entitled “Circuit Intuitions.” He was an Associate Editor for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—PART I: REGULAR PAPERS for 2010–2012, and the Program Chair for the 2004 IEEE ISMVL. He is a registered professional engineer in Ontario, Canada.