Robust and Simple Phase and Timing Synchronization for *M*-ary Partial-Response CPM

Erik Perrins

Department of Electrical Engineering & Computer Science
University of Kansas
Lawrence, KS 66049
E-mail: esp@ieee.org
June 17, 2014

ABSTRACT | We consider carrier phase and symbol timing synchronization for M-ary partial-response continuous phase modulation (CPM). We focus on developing a classical phase locked loop (PLL)-based method that is robust even for M-ary partial-response CPMs, which has proven to be elusive thus far in the literature. A key part of our design is a simple yet effective timing false lock detector, which solves the problems faced by M-ary partial-response CPMs in the past. The lock detector maintains a running count of successive, simple, short-term false lock decisions, rather than evaluating a single, long-term decision. Using a Markov chain model, we show that the lock detector can provide accurate and rapid timing corrections. We provide a comprehensive set of numerical performance results for three different M-ary partial-response CPM schemes, including S-curves, probability of false detection, acquisition time, steady-state error variance, and transient error tracking; we also consider the so-called *tilted phase* CPM model in our analysis, which has fundamentally different synchronization behavior from the traditional CPM model. Our results emphasize the low signal-to-noise ratio (SNR) regime, to show that our system can be used in modern, capacity-approaching, coded CPM applications.

1 | INTRODUCTION

Continuous phase modulation (CPM) [1] has long been appreciated for being bandwidth efficient when used with power-efficient nonlinear amplifiers. Recently [2], a large-scale study was undertaken to identify capacity-approaching CPMs under varying bandwidth and complexity constraints; many of the CPM schemes that were identified fall into the category of *M-ary partial-response CPM*, which presents the greatest challenge when it comes to symbol timing recovery and carrier phase synchronization (see [3] and the references therein for a longer discussion of these challenges).

In sorting through the existing works on CPM synchronization, they can be categorized as data aided (DA), non data aided (NDA), decision directed (DD), etc. For example, the DA approach in [3] is very effective, even for M-ary partial-response CPMs. Our focus in this work is to develop a classical phase locked loop (PLL)-based DD scheme—of the type shown in Fig. 1—that is targeted for use with capacity-approaching CPMs operating at very low signal-to-noise ratios (SNRs). Although the basic scheme in Fig. 1 is of widespread interest and has broad applicability, it has proven challenging for M-ary partial-response CPMs because of *timing false locks*. We tackle the false-lock challenge in this work, because it represents the "last piece"

in the otherwise well-understood system shown in Fig. 1. We first treated this problem in [4], but here we give expanded coverage and a much larger set of numerical results. We readily acknowledge that our solution is adapted from previous efforts in [5] and [6]; however, we assemble our system in a novel, reduced-complexity manner, and demonstrate that accuracy and rapid synchronization can be achieved at the lowest SNRs required by the capacity-approaching schemes in [2]. As such, our approach enables an important class of receivers to be effective when used with modern, capacity-approaching CPMs.

This paper is organized as follows. In Section 2 we outline the CPM signal model. In Section 3 we develop the receiver architecture shown in Fig. 1. In Section 4 we develop the timing false lock detector, including the false lock detection algorithm—which is modeled as a Markov chain—and the quantized timing error correction that is inserted when a false lock is detected. In Section 5 we give a comprehensive set of performance results for three capacity-approaching, M-ary, partial-response CPM schemes; these results include S-curves, probability of false detection, acquisition time, steady-state error variance, and transient error tracking. An important contribution of our work is that we give full consideration to the so-called *tilted phase* model [7] of CPM, which has fundamentally different synchronization behavior from the traditional CPM model. Our results

show that the receiver in Fig. 1 achieves excellent overall performance with modest complexity for *M*-ary partial-response CPMs.

2 | SIGNAL MODEL

We consider CPM signals with complex envelope

$$s(t;\boldsymbol{\alpha}) = \exp\left\{j2\pi h \sum_i \alpha_i q(t - iT_s)\right\} \tag{1}$$

where $\alpha_i \in \{\pm 1, \pm 3, \cdots, \pm(M-1)\}$ is an *M*-ary data symbol, T_s is the duration of each α_i, and h is the modulation index. The phase response $q(t)$ is the time-integral of a frequency pulse $f(t)$ with area 1/2 and duration LT_s. When $L = 1$ the signal is *full response* and when $L > 1$ it is *partial response*. Popular frequency pulse shapes are length-*L* rectangular (*L*REC), raised cosine (*L*RC), and Gaussian (*L*G) [1, p. 52].

The complex envelope of the received signal is modeled as

$$r(t) = \sqrt{E_s/T_s}\, s(t - \tau;\boldsymbol{\alpha}) e^{j\phi} + w(t) \tag{2}$$

where E_s is the energy per symbol, τ is the symbol timing offset, ϕ is the carrier phase offset, and $w(t)$ is complex-valued additive white Gaussian noise (AWGN) with zero mean and power spectral density N_0. The received signal is passed through an anti-aliasing filter (AAF) that is assumed not to distort the signal component of the received waveform. The output of the AAF is sampled at a rate $1/T$, which we assume is an integer multiple N of the symbol rate $1/T_s$ (we have used $N = 4$ herein). The samples of $r(t)$ at the instants $t = nT$ are denoted as $r_{\mathrm{U}}[n]$, where the subscript indicates that they are *unsynchronized* with respect to symbol timing and carrier phase. The relationship between the sample index n and the symbol index k is $kN \leq n < (k+1)N$ with $n = kN + m$ and $0 \leq m \leq N - 1$.

Our focus is on estimating and correcting for ϕ and τ for CPM schemes with $M > 2$ and $L > 1$ (*M*-ary, partial response), especially at the low ratios of E_s/N_0 that are encountered in modern, capacity approaching coding schemes. To that end, we provide numerical results and examples for 4 specific CPM schemes:

Scheme 1: $M = 4$, $h = 1/4$, 2RC;
Scheme 2: $M = 4$, $h = 1/4$, 2REC;
**Scheme 3:** $M = 8$, $h = 1/8$, 2RC;
Scheme 4: $M = 2$, $h = 1/2$, 4G.

The last scheme (a version of Gaussian minimum shift keying, or GMSK) is not an *M*-ary scheme, but is included to illustrate the effectiveness of our final results.

In what follows, we refer to estimated and hypothesized values of a generic quantity x as $\hat{x}$ and $\tilde{x}$, respectively. Also, $\hat{x}$ and $\tilde{x}$ can assume the same values as x itself.

3 | RECEIVER ARCHITECTURE

A block diagram of the receiver is shown in Fig. 1. Many of the receiver modules are described in the existing literature and are summarized as follows:

- The phase corrector applies the phase estimate $\hat{\phi}[k]$ via the operation $e^{-j\hat{\phi}[k]} r_{\mathrm{U}}[n]$. The interpolator applies the timing estimate $\hat{\tau}[k]$, which results in the *synchronized*
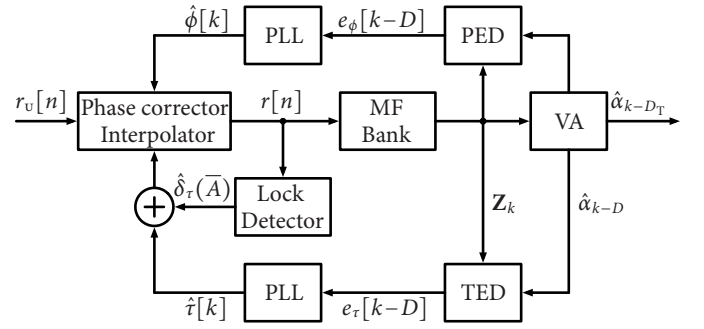
FIGURE 1 | Block diagram of a CPM receiver with decision-directed (DD) PLLs.

samples of the received signal, $r[n]$; we have implemented the piecewise-parabolic interpolator described in [8], [9].

- The matched filter (MF) bank can be implemented in a number of ways. We have applied the standard MF bank, e.g. [1, Ch. 7], and the pulse amplitude modulation (PAM) MF bank [10]–[12], although other reduced-complexity options are equally applicable, e.g. [3], [13]–[17].

- The Viterbi algorithm (VA), e.g. [1, Ch. 7], makes use of the set of MF samples, $\mathbf{Z}_k$, to update its path metrics (over a possibly reduced trellis) and produce a detected symbol $\hat{\alpha}_{k-D_{\mathrm{T}}}$, where D_{T} is the *traceback delay*.

- The DD phase error detector (PED) and timing error detector (TED) are developed in [5]; these are selected due to their excellent steady-state performance, which approaches the modified Cramér-Rao bound (MCRB) [18]. They make use of *tentative* decisions $\hat{\alpha}_{k-D}$ (and the corresponding MF samples $\mathbf{Z}_{k-D}$), where $D < D_{\mathrm{T}}$ in general and we have adopted $D = 1$ [5]. The expression for the DD PED output is

$$e_\phi[k - D] = \mathrm{Im}\left\{\gamma_{k-D}(\hat{e}_{k-D})\right\} \tag{3}$$

where $\hat{e}_{k-D}$ denotes the edge (branch) in the trellis at time step $k - D$ with the best overall metric (i.e. the *global survivor*), and $\gamma_{k-D}(\hat{e}_{k-D})$ is the complex-valued metric increment associated with that edge. The expression for the DD TED output is

$$e_\tau[k - D] = \mathrm{Re}\left\{\dot{\gamma}_{k-D}(\hat{e}_{k-D})\right\} \tag{4}$$

where $\dot{\gamma}$ is the time derivative of γ; this can be obtained from derivative MFs or it can be approximated by taking the difference of early/late samples of the regular MFs [5].

- The phase locked loops (PLLs) are standard in their design, e.g. [19, Appx. C]. Their outputs constitute the final carrier phase and symbol timing estimates, $\hat{\phi}[k]$ and $\hat{\tau}[k]$, respectively. In a more general model of (2) with time-varying $\phi(t)$ and $\tau(t)$, a second-order PLL should be used to resolve residual frequency offsets.

In [5] it was identified that the DD TED is susceptible to false locks when used with *M*-ary, partial-response CPMs. This is verified by the so-called *S-curve* shown in Fig. 2, which was obtained by simulation for Scheme 1; the lock points manifest themselves as zero-crossings with a positive slope, and thus the *false* lock points occur at $\delta_\tau \approx \pm 0.35 T_s$, where $\delta_\tau \triangleq \tau - \hat{\tau}$ is the timing error. Thus, we define $\delta_{\mathrm{F}} = 0.35 T_s$ as the false lock point
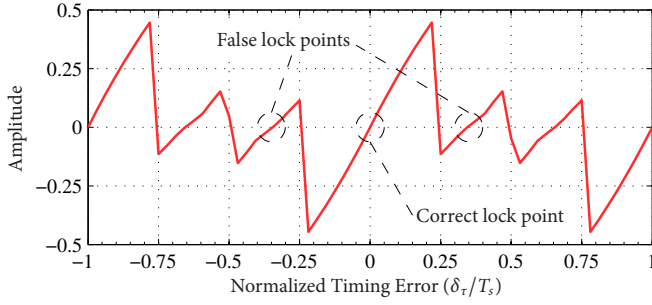
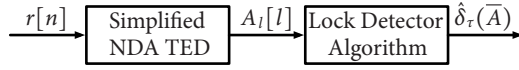**FIGURE 2 |** S-curve for the DD TED in [5] for Scheme 1.

FIGURE 3 | Block diagram of the timing false lock detector.

FIGURE 5 | The unit circle divided into two regions by the condition $C(A)$, with $-\text{sgn}(\text{Im}\{A\})$ used to differentiate between the two false lock points of $\pm 0.35 T_s$ for Scheme 1.

ming operation is

$$A_1[l] \triangleq \sum_{n=lNL_0}^{(l+1)NL_0-1} a_1[n]. \tag{5}$$

As was done in [20], we achieve a major reduction in complexity by quantizing the impulse response $h_1[n]$ using the function $Q_1(\cdot)$ defined in [20, Eq. (12)], which returns only three values, i.e. $Q_1(h_1[n]) \in \{-M_{h_1}, 0, M_{h_1}\}$, where $M_{h_1} \triangleq \max_n(|h_1[n]|)$. This quantization obviates the need for multiplications within the filter. This reduces the complexity of the NDA TED to 8 multiplications per sample $a_1[n]$. For the special case of $N = 4$, the number of multiplications is only 5 per $a_1[n]$ due to the mixers in Fig. 4 assuming values of $\{\pm 1, \pm j\}$ half of the time.

By comparison, when the original (unquantized) $h_1[n]$ is used, $2[(L_h - 1)/2 + 1]$ additional multiplications per sample $a_1[n]$ are needed for the most efficient discrete-time implementation (i.e. exploiting even symmetry), where L_h is the number of non-zero samples in $h_1[n]$. For example, Fig. 11 gives a plot of $h_1(t)$ and $Q_1(h_1(t))$ for Scheme 1; because $h_1[n]$ has $L_h = 19$ with $N = 4$, this amounts to 20 additional multiplications per $a_1[n]$.

4 | TIMING FALSE LOCK DETECTOR

4.1 | Simplified NDA TED

The timing false lock detector consists of the two modules shown in Fig. 3. For convenience, the block diagram of the NDA TED from [6, Fig. 2] is reproduced here in Fig. 4 using the following notation:

- The input to the NDA TED is $r[n]$, which is segmented into non-overlapping intervals of L_0 symbol times (NL_0 samples); each segment is indexed by l.
- The impulse response of the internal filter block is $h_1[n]$ (see [6]), which is real-valued and typically has a duration of four or more symbol times.
- The final output of the NDA TED, which follows the summing operation shown in Fig. 4, is $A_1[l]$.
- The input to the summing operation is $a_1[n]$.
- The relationship between the input and output of the sum-

4.2 | Quantization of the NDA TED Output

Because the NDA TED processes $r[n]$, and because $r[n]$ has already been synchronized by the receiver's primary method of timing recovery (i.e., the phase corrector and the interpolator in Fig. 1), we recognize that the NDA TED estimates any *residual timing error* that may be present. If this residual error is "small," then the receiver is assumed to have locked correctly; if it is "large," then a false lock is assumed.

Adapting [6, Eq. (29)] to the present context, an estimate of the *residual timing error* is obtained as

$$\hat{\delta}_\tau = -\frac{T_s}{2\pi} \arg\{A_1[l]\}. \tag{6}$$

Because the $\arg\{\cdot\}$ function is non-trivial in hardware, we are interested in simple-to-compute quantities involving $A_1[l]$ that can be used to divide the unit circle into "correct lock" and "false lock" regions. This question was entertained briefly in [4], but we give additional results here.

FIGURE 4 | Block diagram of the NDA TED from [6], labeled with the notation used herein. The TED is greatly simplified by using the quantized filter response $Q_1(h_1[n])$ in place of the original response $h_1[n]$.
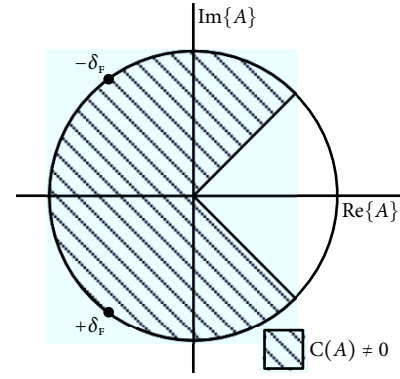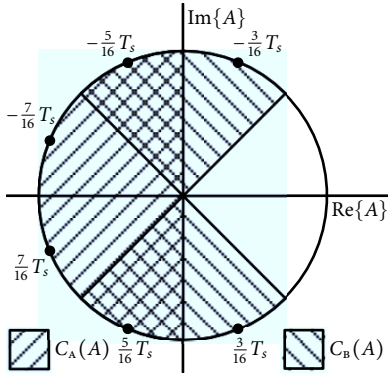
FIGURE 6 | The unit circle divided into eight "phase sectors" based on the three binary-valued conditions: $C_A(A)$, $C_B(A)$, and $-\text{sgn}(\text{Im}\{A\})$.

FIGURE 7 | S-curves for the tilted phase model (red solid line) vs. the traditional model (gray dashed line). The period for the tilted phase model is $2T_s$, whereas it is T_s for the traditional model. Both S-curves are for the DD TED in [5] for Scheme 1.

1) "Binary" Quantization: We begin by partitioning the unit circle into two regions. The lock detector does this by testing the following condition for the generic complex number A:

Let $C(A) \neq 0$, if $(\text{Re}\{A\} < 0)$ or $(|\text{Im}\{A\}| > |\text{Re}\{A\}|)$;

$\quad C(A) = 0$, otherwise. $\quad$ (7)

When this condition is false[1] ($C(A) = 0$) we have $|\hat{\delta}_\tau| < \frac{1}{8}T_s$, which is well inside the region of the S-curve in Fig. 2 where the primary timing recovery system operates correctly. When this condition is true ($C(A) \neq 0$) we have $|\hat{\delta}_\tau| > \frac{1}{8}T_s$, which is the region of the S-curve in Fig. 2 that contains the false lock points.

Motivated by the above arguments, we propose a simple estimate of the residual timing error based on $C(A)$ for M-ary partial response CPMs in general:

$$\hat{\delta}_\tau(A) \triangleq \begin{cases} 0, & C(A) = 0 \\ -\text{sgn}(\text{Im}\{A\}) \times \delta_F, & C(A) \neq 0 \end{cases} \quad (8)$$
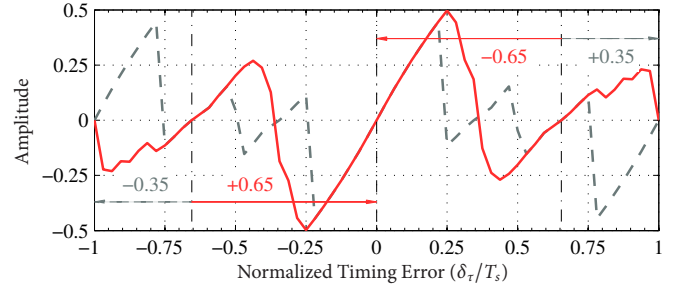
where δ_F is determined by the false lock points on the S-curve for the given CPM scheme (once again, $\delta_F = 0.35 T_s$ for Scheme 1 in Fig. 2). Fig. 5 illustrates the timing estimate in (8). The signum function is defined as

$$\text{sgn}(x) \triangleq \begin{cases} +1, & x > 0 \\ 0, & x = 0 \\ -1, & x < 0 \end{cases} \quad (9)$$

where the $x = 0$ case almost never occurs when x is real [as it is in (8)], but occurs regularly when x is an integer [as will be seen later].

2) Expanded Quantization: The estimate in (8) favors extreme simplicity, which can come at the expense of accuracy if there is more than one false lock point, or if there is a false lock "region." Additionally, it requires calibration toward a specific value of δ_F (which is, admittedly, straightforward to accomplish).

As an alternative, the lock detector can test the following two

[1] We note that the polarity of Eq. (7) is reversed from that found in [4, Eq (4)].

sub-conditions of (7), which sub-divide the true ($C(A) \neq 0$) case:

$$C_A(A) \triangleq \begin{cases} 1, & \text{Re}\{A\} < 0 \\ 0, & \text{otherwise} \end{cases} \quad (10)$$

$$C_B(A) \triangleq \begin{cases} 1, & |\text{Im}\{A\}| > |\text{Re}\{A\}| \\ 0, & \text{otherwise.} \end{cases} \quad (11)$$

These can be combined to form an expanded version of (7):

$$C(A) \triangleq 2C_A(A) + C_B(A), \quad C(A) \in \{0, 1, 2, 3\}. \quad (12)$$

The 4-ary condition $C(A)$ in (12), along with $-\text{sgn}(\text{Im}\{A\})$, divides the unit circle into eight "phase sectors," as shown in Fig. 6. A quantized timing correction is then obtained as the center-point of the sectors with $C(A) \neq 0$:

$$\hat{\delta}_\tau(A) \triangleq \begin{cases} 0, & C(A) = 0 \\ -\text{sgn}(\text{Im}\{A\})\frac{3}{16}T_s, & C(A) = 1 \\ -\text{sgn}(\text{Im}\{A\})\frac{5}{16}T_s, & C(A) = 3 \\ -\text{sgn}(\text{Im}\{A\})\frac{7}{16}T_s, & C(A) = 2 \end{cases} \quad (13)$$

3) Quantization for the Tilted Phase Model: The tilted phase model for CPM [7] is advantageous because it reduces the number of phase states (and therefore the overall number of trellis states) by a factor of two. However, another consequence of the tilted phase model is that it fundamentally alters the synchronization behavior of the receiver.

For the traditional CPM receiver, the timing recovery S-curve has a period of T_s, as seen in Fig. 2. The tilted phase model introduces a notion of even and odd symbol indexes, which causes the timing recovery S-curve to have a period of $2T_s$. This is illustrated in Fig. 7, which shows S-curves for the tilted phase model vs. the traditional model for Scheme 1. The expanded period of $2T_s$ poses a problem for the lock detector, which is completely outside of the VA block in Fig. 1 and is thus "unaware" of which model (tilted phase or traditional) is being used for the trellis. Ideally, what is needed is a lock detector that is sensitive to the entire length-$2T_s$ interval, i.e., $-T_s < \delta_\tau < T_s$. However, because the lock detector is sensitive only to δ_τ in the interval $-\frac{1}{2}T_s < \delta_\tau < \frac{1}{2}T_s$, some adjustments are needed.

As with all of the above cases, we assign the interval $0 < |\delta_\tau| < \frac{1}{8}T_s$ to the "correct lock" case, where no timing correction is needed. Because of the limited range of the lock detector, this
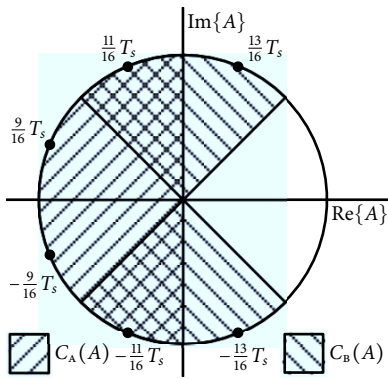
FIGURE 8 | The unit circle for the tilted phase model. The timing correction is designed to accommodate timing errors only in the range $\frac{1}{2}T_s < |\delta_\tau| < \frac{7}{8}T_s$.

FIGURE 9 | State diagram for timing false lock detector algorithm.

means that the interval $\frac{7}{8}T_s < |\delta_\tau| < T_s$ is also (unavoidably) assigned to the "correct lock" case. We must now decide what to do with the intervals $\frac{1}{8}T_s < |\delta_\tau| < \frac{1}{2}T_s$ and $\frac{1}{2}T_s < |\delta_\tau| < \frac{7}{8}T_s$. In Fig. 7, we see that the tilted phase TED has a false lock at $\delta_F = 0.65T_s$, *which the lock detector perceives as being at* $-0.35T_s$. These findings are typical of other CPM schemes, as we shall see. Therefore, the interval $\frac{1}{2}T_s < |\delta_\tau| < \frac{7}{8}T_s$ is given priority by the timing correction rule, which means it is not possible to accommodate the interval $\frac{1}{8}T_s < |\delta_\tau| < \frac{1}{2}T_s$ within this rule for the tilted phase model. The timing correction rule that thus emerges is

$$\hat{\delta}_\tau(A) \triangleq \begin{cases} 0, & C(A) = 0 \\ \text{sgn}(\text{Im}\{A\})\frac{13}{16}T_s, & C(A) = 1 \\ \text{sgn}(\text{Im}\{A\})\frac{11}{16}T_s, & C(A) = 3 \\ \text{sgn}(\text{Im}\{A\})\frac{9}{16}T_s, & C(A) = 2 \end{cases} \quad (14)$$

which is pictured in Fig. 8 and is basically $T_s - \delta_\tau$ with respect to the previous rule in (13).

With several options for the timing correction rule now defined, we now address the challenge of making the false lock decision more robust. This is necessary because the NDA TED is known to be quite noisy for M-ary, partial-response CPM schemes and small L_0.

4.3 | False Lock Detector Algorithm

In order to reduce the probability of false detection, and also to reduce the noise in the estimated timing correction, we introduce a counting algorithm. The state of the count at index l is $S[l]$, where $S[l] \in \{0, \pm 1, \cdots, \pm N_s\}$. When a new value of $A_1[l]$ becomes available, if $C(A_1[l]) \neq 0$ then the count is incremented in the direction of $\text{sgn}(\text{Im}\{A_1[l]\})$ in order to strengthen the hypothesis of a false lock in that direction on the unit circle. If $C(A_1[l]) = 0$, then the "correct lock" hypothesis is strengthened and the count is incremented toward zero (whichever direction that may be), or it remains at zero if it is already there, i.e. the increment in this case is $-\text{sgn}(S[l-1])$. When the count is non-zero, the algorithm stores a running sum of $A_1[l]$ in the variable $\overline{A}$, which is reset if the count ever returns to zero. In the event that the count overflows/underflows, i.e. $|S[l]| > N_s$, then a "timing false lock" is declared; the lock
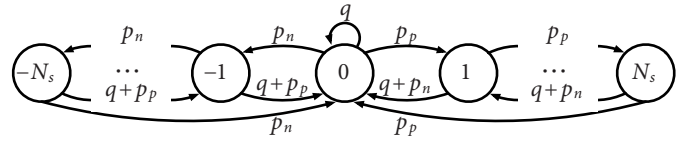
detector then inserts a timing correction into the receiver's primary timing recovery system based on the running sum $\overline{A}$, i.e. $\hat{\delta}_\tau(\overline{A})$—based on one of the timing correction rules in (8), (13), or (14)—and the count is returned to the zero state. We have also observed that the path metrics of the Viterbi algorithm (VA) within the demodulator are "biased" during a timing false lock; therefore, we also reset the VA path metrics to zero when a false lock is detected.

The above steps are summarized in Algorithm 1. The counter is modeled in Fig. 9 as a time-homogeneous Markov chain. There are three probabilities that describe the state transitions: p_p is the probability of transitioning in the positive direction due to $C(A_1[l]) \neq 0$, p_n is the probability of transitioning in the negative direction due to $C(A_1[l]) \neq 0$, q is the probability of $C(A_1[l]) = 0$, and we have $p_p + p_n + q = 1$. These probabilities do not vary with the particular timing correction rule that is employed [(8), (13), or (14)], and thus Algorithm 1 and the analysis below are applicable to all three cases.

5 | PERFORMANCE

5.1 | List of Figures

The following is a list of the types of figures that are presented for Schemes 1, 2, and 3:

- The filter response $h_1(t)$ and quantized version $Q_1(h_1(t))$.
- For the traditional model, the S-curves for the DD PED and TED in [5].

Algorithm 1 Timing False Lock Detector

1: Initialize $S[-1] = 0, \overline{A} = 0$;
2: **for** $l = 0, 1, 2, \cdots$ **do**
3:     Compute $A_1[l]$;
4:     **if** $C(A_1[l]) \neq 0$ **then**,
5:         Update $S[l] = S[l-1] + \text{sgn}(\text{Im}\{A_1[l]\})$;
6:         Update $\overline{A} = \overline{A} + A_1[l]$;
7:     **else**
8:         Update $S[l] = S[l-1] - \text{sgn}(S[l-1])$;
9:     **end if**;
10:     **if** $S[l] = 0$ **then**,
11:         Set $\overline{A} = 0$;
12:     **end if**;
13:     **if** $|S[l]| > N_s$ **then**,
14:         Update $\hat{\tau}[k] = \hat{\tau}[k] + \hat{\delta}_\tau(\overline{A})$;
15:         Set VA path metrics to zero;
16:         Set $S[l] = 0$;
17:         Set $\overline{A} = 0$;
18:     **end if**;
19: **end for**

- For the tilted phase model, the S-curves for the DD PED and TED in [5].
- The probabilities p_p and p_n when the receiver is in the correct lock state ($\delta_\tau \approx 0$), where $p_n = p_p$.
- The probabilities p_p and p_n when the receiver is in the false lock state of $\delta_\tau \approx +\delta_F$, where $p_n \gg p_p$. The reverse situation of $\delta_\tau \approx -\delta_F$, where $p_n \ll p_p$, is not shown because of redundancy.
- The probability of false detection, P_{FD}.
- The acquisition time, t_D/T_s.
- The bit error rate (BER).
- The phase error variance, $\mathrm{Var}(\phi)$, and the normalized timing error variance, $\mathrm{Var}(\tau)$—both of which are compared with their respective MCRBs.
- For the traditional model, the phase error (δ_ϕ, in cycles) and the normalized timing error (δ_τ/T_s) vs. time at three different E_s/N_0, with the lowest E_s/N_0 near channel capacity.
- For the tilted phase model, the phase error and the normalized timing error vs. time at the same three different E_s/N_0.

For Scheme 4 (the GMSK scheme), the lock detector is not necessary, as confirmed by the S-curves for this scheme in Figs. 44 and 45. As such, only the S-curves, BER plot, variance plot, and phase/timing error vs. time plots are given. We now discuss the entire body of results in greater detail.

5.2 | S-Curves

S-curves for the traditional model are shown in Figs. 12, 23, and 34 for Schemes 1, 2, and 3, respectively; the S-curve for the PED in [5] is shown on the top and the S-curve for the TED in [5] is shown on the bottom. For all three schemes, the S-curve for the PED shows that—as expected—there are $2p$ correct lock points around the unit circle, or that the S-curve has a period of $\frac{1}{2p}$ cycles, where p is the denominator of the modulation index when it is expressed as a rational number, i.e., $h = k/p$. Also, for all three schemes, the S-curve for the TED shows that it has a period of T_s, but that in-between these symbol-spaced correct lock points there are *false lock points*, which are of course the main problem addressed by this work.

S-curves for the tilted phase model are shown in Figs. 13, 24, and 35 for Schemes 1, 2, and 3, respectively, using the same top/bottom format for the PED/TED. As was stated previously, the tilted phase model fundamentally alters the synchronization behavior of the receiver. For the PED, this means that there are only p correct lock points around the unit circle, or that the S-curve has a period of $\frac{1}{p}$ cycles. For the TED, this means that the correct lock points are spaced two symbols apart (i.e., the period of the S-curve is $2T_s$) and that the false lock points are spaced differently than before.

As we mentioned above, the S-curves for Scheme 4 (the GMSK scheme) are shown in Figs. 44 and 45 for the traditional model and the tilted phase model, respectively. Because Scheme 4 is a *binary* CPM ($M = 2$), it does not suffer from timing false locks.

5.3 | Probabilities p_p and p_n

In order to evaluate the usefulness of the false lock detector, the performance impact of the parameters p_p, p_n, N_s, and L_0

must be understood. Because no analytical method is available to evaluate p_p (or p_n), we resort to computer simulations.

We first examine the case where the receiver is in a state of correct timing lock (i.e., $\delta_\tau \approx 0$). These results are shown in Figs. 14, 25, and 36 for Schemes 1, 2, and 3, respectively. We have evaluated p_p vs. E_s/N_0 for four different observation intervals L_0. During these simulations, p_p is determined at each E_s/N_0 by counting the occurrences of the joint events $C(A_1[l]) \neq 0$ and $\mathrm{Im}\{A_1[l]\} > 0$, and then dividing this count by the number of trial values of $A_1[l]$ observed; p_n is determined in a similar fashion except with $\mathrm{Im}\{A_1[l]\} < 0$. Each simulation is conducted until at least 1,000 counts is observed. As would be expected, p_p decreases with increasing E_s/N_0 and increasing L_0; also, as would be expected, $p_n = p_p$ when the receiver is in the correct lock state ($\delta_\tau \approx 0$).

We next examine the case where the receiver is in the false lock state of $\delta_\tau \approx +\delta_F$. These results are shown in Figs. 15, 26, and 37 for Schemes 1, 2, and 3, respectively, where the same four values of L_0 are used and the same simulation methodology is employed. As expected, because $\delta_\tau \approx +\delta_F$, the values of p_n approach unity and the values of p_p vanish with increasing E_s/N_0. The reverse situation of $\delta_\tau \approx -\delta_F$, where $p_n \ll p_p$, is not shown because of redundancy.

5.4 | Probability of False Detection, P_{FD}, and the Acquisition Time, t_D/T_s

We turn our attention to the competing design objectives of minimizing the probability of false detection, P_{FD}, while simultaneously minimizing the time needed for correct detection, t_D. With the availability of p_p and p_n, these quantities can be evaluated analytically. Let the stationary distribution $\boldsymbol{\pi}$ of the Markov chain in Fig. 9 be a length-$(2N_s + 1)$ row vector, with the i-th element π_i equal to $\Pr(S[l] = i)$ at equilibrium, and let $\mathbf{P}$ be the state transition matrix with the (i, j)-th element equal to $p_{ij} = \Pr(S[l + 1] = j|S[l] = i)$. $\boldsymbol{\pi}$ is the solution to the eigenvalue/eigenvector equation $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$ corresponding to the eigenvalue of unity [21]. In terms of Fig. 9, P_{FD} is the probability of being in states $\pm N_s$ and transitioning directly back to state zero, given that the receiver is in the correct lock state; it is normalized by L_0 so that it conveys the probability of false detection *per symbol*. Therefore,

$$P_{FD} = \frac{p_p \cdot \pi_{N_s} + p_n \cdot \pi_{-N_s}}{L_0}. \tag{15}$$

Similarly, t_D/T_s is L_0 times the expected number of time steps l until a transition occurs from state $\pm N_s$ directly back to state zero, given that the receiver is in the false lock state; therefore,

$$\frac{t_D}{T_s} = \frac{L_0}{p_p \cdot \pi_{N_s} + p_n \cdot \pi_{-N_s}}. \tag{16}$$

We emphasize the fact that p_p and p_n in (15) are obtained in a simulation where the receiver is in the correct lock state (as in Figs. 14, 25, and 36), and p_p and p_n in (16) are obtained in a separate simulation where the receiver is in the false lock state (as in Figs. 15, 26, and 37).

In [2], $E_s/N_0 \geq 2$ dB is found to be the region where Scheme 1 is optimal, thus $E_s/N_0 = 2$ dB is the lowest SNR that must be considered for this scheme. Likewise, the lowest SNR that need

TABLE 1 | Design pairs (L_0, N_s) for Cases 1 and 2 for each Scheme. The designs in bold result in good P_{FD} and t_{D}/T_s and thus appear in both cases.

Scheme 1				
Case 1	(8,13)	(16,11)	(32,9)	**(64,7)**
Case 2	(8,9)	(16,9)	(32,8)	**(64,7)**
Scheme 2				
Case 1	(32,13)	(64,11)	(128,9)	**(256,7)**
Case 2	(32,9)	(64,9)	(128,8)	**(256,7)**
Scheme 3				
Case 1	(8,12)	(16,10)	(32,8)	**(64,7)**
Case 2	(8,8)	(16,8)	(32,7)	**(64,7)**

be considered for Scheme 2 is $E_s/N_0 = 2$ dB and for Scheme 3 is $E_s/N_0 = 5$ dB. We now design lock detector schemes for use at these *target SNRs*. Each design consists of a pair of parameters (L_0, N_s). The designs are grouped according to two different design rules, or design *cases*:

Case 1: These designs are chosen to yield *tightly-grouped values of P_{FD} at the target SNR*, as shown Figs. 16, 27, and 38 for Schemes 1, 2, and 3, respectively. The design pairs for Case 1 are listed in Table 1. Although P_{FD} at the target SNR for these designs is nearly identical, t_{D}/T_s varies significantly, as shown in Figs. 17, 28, and 39 for Schemes 1, 2, and 3, respectively.

Case 2: These designs are chosen to yield *tightly-grouped values of t_{D}/T_s at the target SNR*, as shown in Figs. 17, 28, and 39 for Schemes 1, 2, and 3, respectively. The design pairs for Case 2 are also listed in Table 1. Although t_{D}/T_s at the target SNR for these designs is nearly identical, P_{FD} varies significantly, as shown in Figs. 16, 27, and 38 for Schemes 1, 2, and 3, respectively.

These results show that when L_0 is decreased, N_s must be increased (and vice versa), in order to maintain steady performance for P_{FD} (or t_{D}/T_s). There are values that are too extreme (e.g. $L_0 = 8$ for Scheme 1 does a poor job of balancing the tradeoff between P_{FD} and t_{D}/T_s); but there are also designs that provide very low P_{FD} while maintaining a rapid acquisition time. For example, the design with $(64, 7)$ belongs to both Cases for Scheme 1 and signifies that a balance can be achieved between competing tradeoffs; its P_{FD} is at or below 10^{-6}, with an acquisition time in the range $500 < t_{\text{D}}/T_s < 1500$, which is comparable to a PLL with a normalized loop bandwidth in the range $3 \times 10^{-4} < BT_s < 1 \times 10^{-3}$.

5.5 | Steady-State Performance of the PLL-Based Receiver

1) BER Performance: Figs. 18, 29, 40, and 46 show the BER for Schemes 1–4, respectively. Each plot shows the theoretical maximum likelihood sequence detection (MLSD) bound, the BER performance of a receiver with perfect synchronization, and the BER performance of the proposed receiver in Fig. 1; the phase and timing PLLs in Fig. 1 have normalized loop bandwidths of $BT_s = 10^{-3}$ for all Schemes, and the lock detector uses the design pair shown in bold in Table 1 for each Scheme. The BER plots show that the proposed receiver in Fig. 1 achieves a steady-state

BER that is essentially the same as perfect synchronization, even at low SNRs.

2) Phase and Timing Error Variance: Figs. 19, 30, 41, and 47 show the phase and normalized timing error variances for Schemes 1–4, respectively. Because the performance of the phase and timing PLLs has already been studied in [5], these plots are simply an extension for Schemes 1–4 of the data reported in [5].

3) Phase and Timing Error vs. Time: Figs. 20, 31, 42, and 48 show the phase and timing error vs. time (i.e., *transient behavior*) for Schemes 1–4, respectively, using the conventional CPM model; the results are repeated for three different SNRs for each scheme, as noted in the figure captions. The exact same conditions are repeated for the tilted phase model in Figs. 21, 32, 43, and 49 for Schemes 1–4. In each of these plots, 64 trial operations of the proposed receiver in Fig. 1 were conducted; in each trial, the receiver was initialized with a random phase and timing offset before being set into operation. The figures also show a red envelope, which depicts the ideal operation of a PLL with normalized loop bandwidth $BT_s = 10^{-3}$.

These data clearly show the PLLs settling into false timing locks, during which time the phase error remains large. The timing error corrections appear in the plots as *step functions* and are very noticeable; we used the timing correction rule in (13) for the traditional model and (14) for the tilted phase model. As one would expect, the transient period is longer for low SNRs. It is also slightly longer for the tilted phase receivers. At high SNRs, the overall acquisition time—including false lock correction and PLL settling time—is in line with the ideal PLL operation.

### 5.6 | Additional Discussion on the False Lock Detector

There are some added results given in [4, Figs. 5–6] for Scheme 1 with the extreme designs $(2048, 0)$ and $(1536, 0)$, which do away with the counting algorithm all together (i.e., $N_s = 0$) and simply increase L_0 until a sufficiently low $P_{\text{FD}} = (p_p + p_n)/L_0$ is achieved (or until the desired t_{D}/T_s is not exceeded). These designs correspond to the lock detector solution that was proposed in [5]. The results in [4] show that a low P_{FD} can be achieved with this approach, but that the large value of L_0 results in a large value of t_{D}/T_s that is more or less "fixed" and cannot decrease with increasing SNR. This helps underscore the contribution of our approach of counting shorter observations.

These final results are perhaps counterintuitive and prompt this important question: How is it possible that the lock detector performs better by assembling many brief observations (i.e., smaller L_0 and $N_s > 0$), than it does by using one long observation (i.e., larger L_0 and $N_s = 0$)? In other words, how can better performance be obtained with a shorter observation interval? This important question is answered by the data presented in Fig. 10.

Fig. 10 (a) plots the running accumulation of the variable $a_1[n]$:

$$\sum_{m=0}^{kN} a_1[m] \qquad (17)$$

and thus it is similar to (5). The observation interval extends out to 1536 symbols. The value at the end of this observation interval corresponds to the value of $A_1[l]$ for the $L_0 = 1536$ case. Note that the imaginary part of the accumulation exceeds the real part
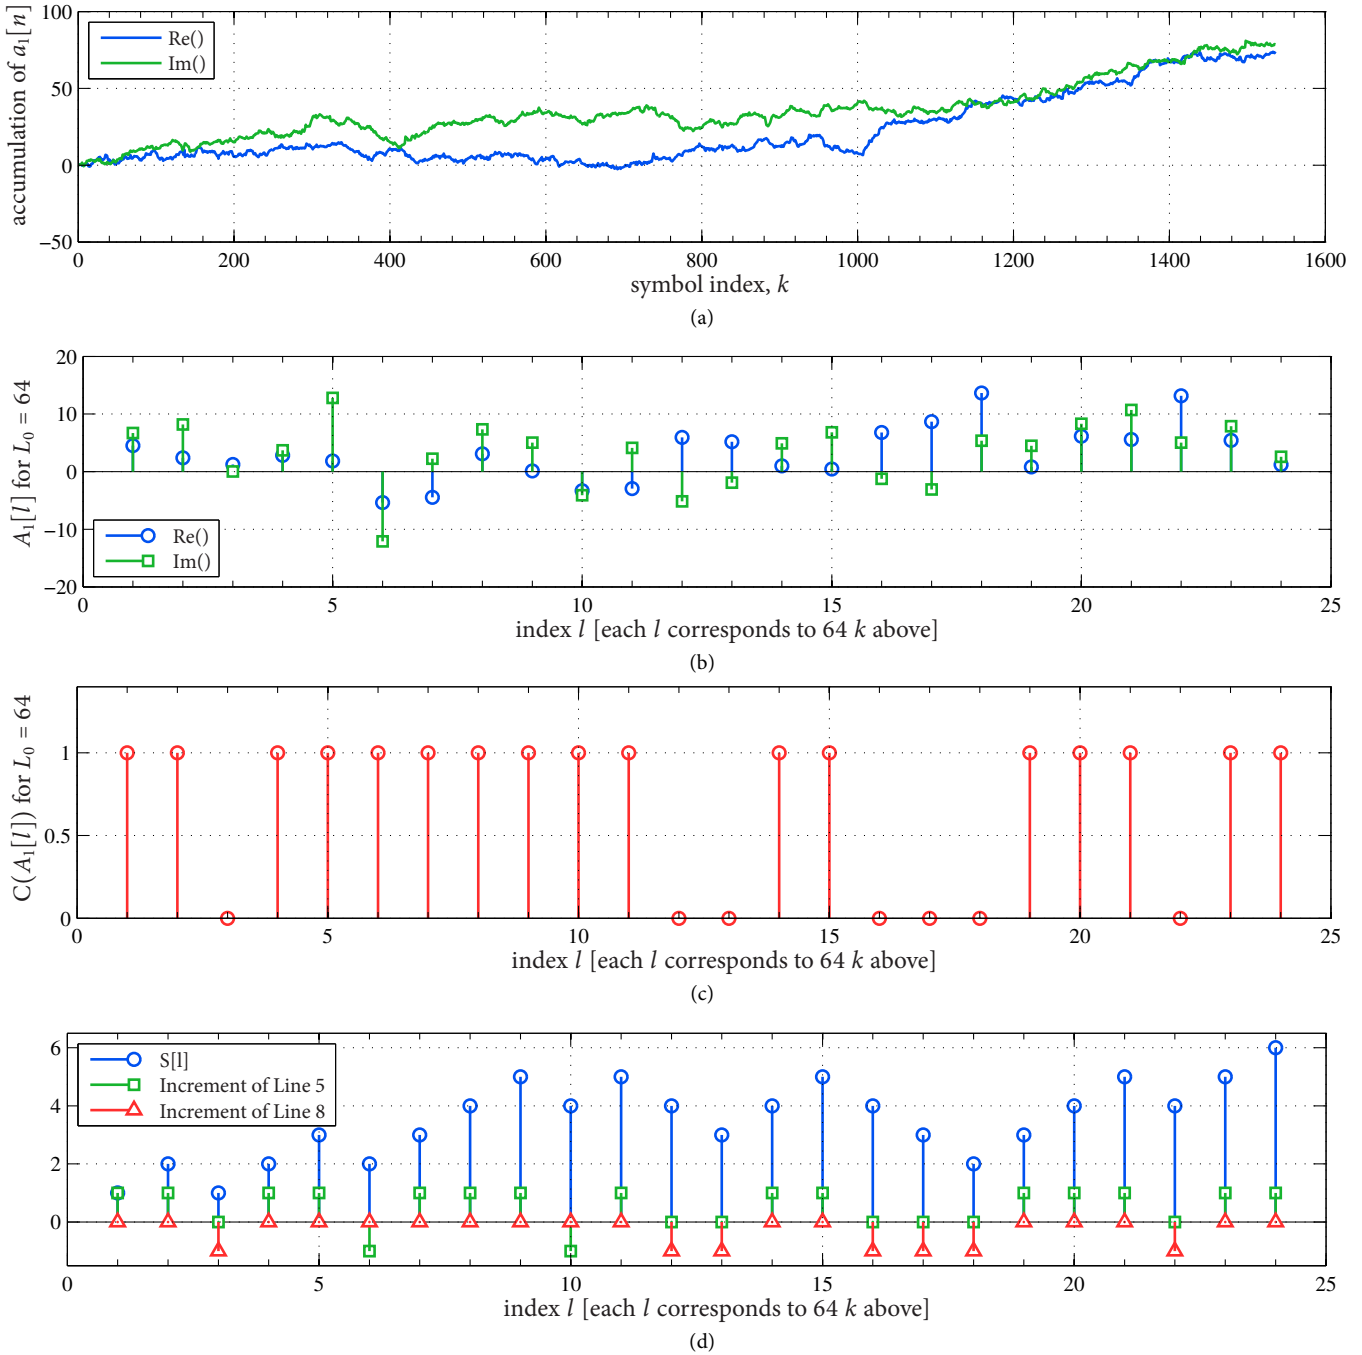
FIGURE 10 | Detailed time sequences for Scheme 1 for: (a) the accumulation of $a_1[n]$; (b) $A_1[l]$ for the $L_0 = 64$ case; (c) the condition $C(A_1[l])$ for the $L_0 = 64$ case; and (d) state of the counter, $S[l]$, and the values of the counter increments.

at symbol index $k = 1536$ (and the real part is positive), which means that $C(A) \neq 0$. Therefore, this data set corresponds to a "false detect" for the $L_0 = 1536$, $N_s = 0$ configuration.

Fig. 10 (b) plots $A_1[l]$ for the $L_0 = 64$ case. Thus the raw data $\{a_1[n]\}$ are segmented into non-overlapping intervals of 64 symbols and summed, as indicated in (5). Note that the abscissa of Fig. 10 (b) is the index l, which indexes 64 symbols at a time; however, we emphasize that all four subfigures in Fig. 10 are in *time alignment*.

Fig. 10 (c) plots the condition $C(A_1[l])$ for the $L_0 = 64$ case [i.e. it is the condition in (7) applied to the data plotted in

Fig. 10 (b)]; a 1 is used to represent $C(A_1[l]) \neq 0$. It is easy to visually confirm that the seven occasions where the condition evaluates to 0 (false), the real value in Fig. 10 (b) is positive $(\text{Re}\{A_1[l]\} > 0)$ *and* the imaginary value has a smaller magnitude than the real value $(|\text{Im}\{A\}| < |\text{Re}\{A\}|)$.

Fig. 10 (d) contains the key data for the final explanation. The sequence with circle markers (blue) corresponds to the state of the counter, $S[l]$, for $L_0 = 64$. As stated in Algorithm 1, $S[l]$ can be incremented according to Line 5 or Line 8, depending on the result of the condition $C(A_1[l])$ [or equivalently, using Figs. 5, 6, or 8, depending on whether or not the observed angle of $A_1[l]$

falls in the shaded (true) or unshaded (false) regions].

The sequence with square markers (green) in Fig. 10 (d) corresponds to the counter increment of Line 5; this increment is non-zero only when $C(A_1[l]) \neq 0$ (true). This sequence behaves more or less as one would expect. It is based on shorter observations of the data. The general trend shown in Fig. 10 (a) is that the imaginary part (green) is positive and greater than the real part (blue). This is reflected more coarsely in Fig. 10 (a), as would be expected, and this trend (i.e. positive increments) follows in Fig. 10 (d) with the sequence with square markers (green).

The sequence with the triangle markers (red) in Fig. 10 (d) corresponds to the counter increment of Line 8, and is non-zero for the seven occasions where the observed angle of $A_1[l]$ falls in the unshaded (false) region of Figs. 5, 6, and 8. As can be seen by all of the data in Fig. 10 [i.e., the gap between the real and imaginary parts in Fig. 10 (a), the final margin of victory in Fig. 10 (a) of the imaginary part over the real part, the many positive square-marker (green) increments in Fig. 10 (d)], everything is pointing toward a false detection. However, as can be seen by the circle-marker (blue) sequence of $S[l]$ in Fig. 10 (d), the counter never exceeds the value of $N_s = 7$, and thus the $L_0 = 64$, $N_s = 7$ configuration does not falsely detect. This is because the triangle-marker (red) increments *set the counter back*. In other words, when the angle of $A_1[l]$ falls in the unshaded (false) region of Figs. 5, 6, or 8, the counter is penalized and thus the probability of false detection is reduced. In order for the counter to overflow, there must be a regular and consistent trend of square-marker (green) increments in the same direction, and few triangle-marker (red) increments. Figs. 17, 28, and 39 demonstrate the unavoidable consequence, that when we want the counter to overflow quickly, the triangle-marker (red) increments can prolong the process; this is the reason for the larger t_D/T_s at low SNRs.

This "intuitive" discussion is meant to motivate why it is that a low P_{FD} can be achieved in Figs. 16, 27, and 38 for short observation intervals. We emphasize that *all* of these dynamics are fully captured by the analysis of the Markov chain model. The probability of a triangle-marker (red) increment is measured by simulation as q. The probabilities of the positive and negative square-marker (green) increments are, respectively, p_p and p_n. Once we have measured these probabilities, we can design L_0 and N_s to jointly minimize P_{FD} and t_D/T_s.

6 | CONCLUSION

We have developed a classical PLL-based receiver architecture that is robust especially for M-ary partial-response CPMs, which has proven to be an elusive task in previous studies. The main new element to this system is a timing false lock detector for CPM that is adapted from an existing basic scheme. Our approach maintains a running count of successive, short-term false lock decisions, rather than evaluating a single, long-term decision. Using analysis of a Markov chain model for our scheme, we have demonstrated that it can achieve low probability of false detection and rapid synchronization for capacity-approaching CPMs over their SNR operating range. We have provided a comprehensive set of numerical results, which demonstrate the effectiveness of our design. A key aspect of our analysis is the inclusion of the tilted-phase CPM model, which must be treated separately due to its distinct synchronization behavior.

REFERENCES

[1] J. B. Anderson, T. Aulin, and C.-E. Sundberg, *Digital Phase Modulation*. New York: Plenum Press, 1986.

[2] A. Perotti, A. Tarable, S. Benedetto, and G. Montorsi, "Capacity-achieving CPM schemes," *IEEE Trans. Inform. Theory*, vol. 56, pp. 1521–1541, Apr. 2010.

[3] Q. Zhao and G. L. Stüber, "Robust time and phase synchronization for continuous phase modulation," *IEEE Trans. Commun.*, vol. 54, pp. 1857–1869, Oct. 2006.

[4] E. Perrins, "A timing false lock detector for M-ary partial response CPM," *IEEE Wireless Commun. Letters*, vol. 2, pp. 671–674, Dec. 2013.

[5] M. Morelli, U. Mengali, and G. M. Vitetta, "Joint phase and timing recovery with CPM signals," *IEEE Trans. Commun.*, vol. 45, pp. 867–876, Jul. 1997.

[6] A. N. D'Andrea, U. Mengali, and M. Morelli, "Symbol timing estimation with CPM modulation," *IEEE Trans. Commun.*, vol. 44, pp. 1362–1372, Oct. 1996.

[7] B. E. Rimoldi, "A decomposition approach to CPM," *IEEE Trans. Inform. Theory*, vol. 34, pp. 260–270, Mar. 1988.

[8] F. Gardner, "Interpolation in digital modems—Part I: Fundamentals," *IEEE Trans. Commun.*, vol. 41, pp. 501–507, Mar 1993.

[9] L. Erup, F. Gardner, and R. A. Harris, "Interpolation in digital modems—Part II: Implementation and performance," *IEEE Trans. Commun.*, vol. 41, pp. 998–1008, Jun. 1993.

[10] P. A. Laurent, "Exact and approximate construction of digital phase modulations by superposition of amplitude modulated pulses (AMP)," *IEEE Trans. Commun.*, vol. 34, pp. 150–160, Feb. 1986.

[11] U. Mengali and M. Morelli, "Decomposition of M-ary CPM signals into PAM waveforms," *IEEE Trans. Inform. Theory*, vol. 41, pp. 1265–1275, Sep. 1995.

[12] G. K. Kaleh, "Simple coherent receivers for partial response continuous phase modulation," *IEEE J. Sel. Areas Commun.*, vol. 7, pp. 1427–1436, Dec. 1989.

[13] M. H. M. Costa, "A practical demodulator for continuous phase modulation," in *Proc. Int. Symp. Inform. Theory*, (Trondheim, Norway), Jun. 1994.

[14] J. Huber and W. Liu, "An alternate approach to reduced complexity CPM receivers," *IEEE J. Select. Areas Commun.*, vol. 7, pp. 1437–1449, Dec. 1989.

[15] P. Moqvist and T. Aulin, "Orthogonalization by principal components applied to CPM," *IEEE Trans. Commun.*, vol. 51, pp. 1838–1845, Nov. 2003.

[16] S. J. Simmons, "Simplified coherent detection of CPM," *IEEE Trans. Commun.*, vol. 43, pp. 726–728, Feb./Mar./Apr. 1995.

[17] W. Tang and E. Shwedyk, "A quasi-optimum receiver for continuous phase modulation," *IEEE Trans. Commun.*, vol. 48, pp. 1087–1090, Jul. 2000.

[18] A. N. D'Andrea, U. Mengali, and R. Reggiannini, "The modified Cramer-Rao bound and its application to synchronization problems," *IEEE Trans. Commun.*, vol. 42, pp. 1391–1399, Feb./Mar./Apr. 1994.

[19] M. Rice, *Digital Communications: A Discrete-Time Approach*. New York: Prentice Hall, 2009.

[20] P. Chandran and E. Perrins, "Symbol timing recovery for CPM with correlated data symbols," *IEEE Trans. Commun.*, vol. 57, pp. 1265–1270, May 2009.

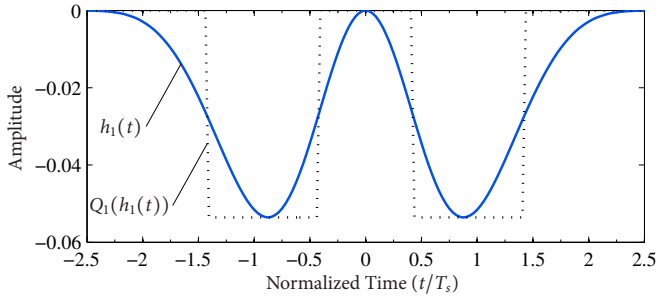[21] J. R. Norris, *Markov Chains*. Cambridge University Press, 1998.

FIGURE 11 | Scheme 1 ($M = 4$, $h = 1/4$, 2RC): Filter response $h_1(t)$ and quantized version $Q_1(h_1(t))$.
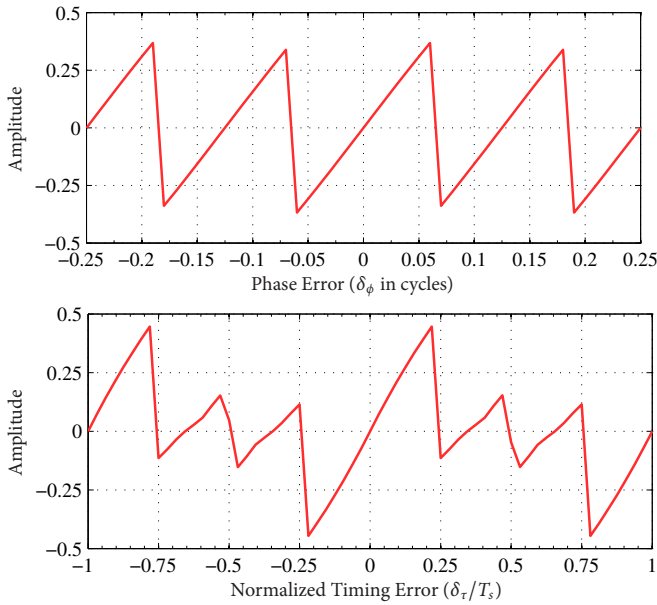

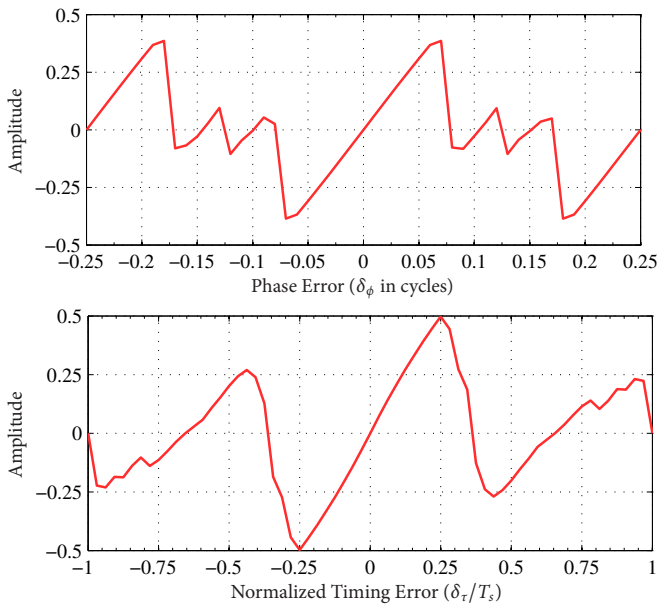**FIGURE 12 |** Scheme 1: S-curves for the DD PED (top) and TED (bottom) in [5].

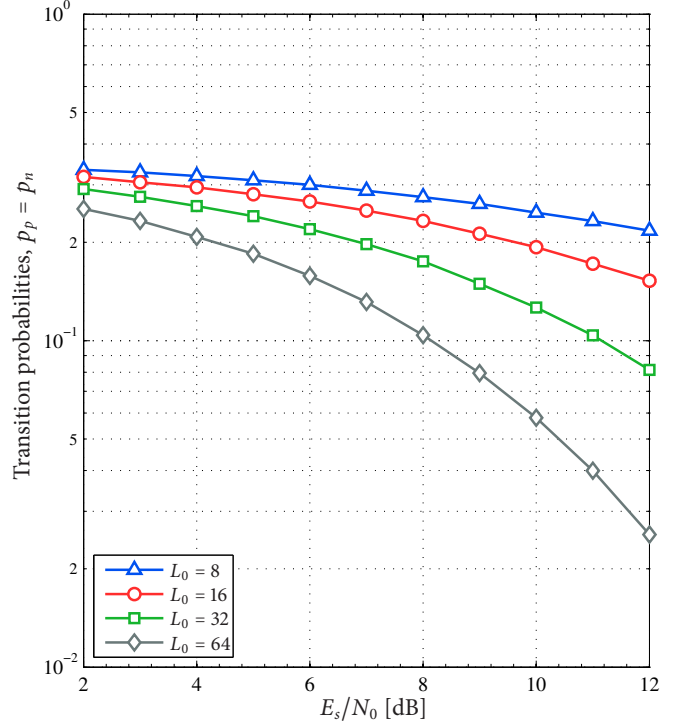**FIGURE 13 |** Scheme 1: S-curves for the DD PED (top) and TED (bottom) in [5] for the tilted-phase model.

FIGURE 14 | Scheme 1: The probabilities p_p and p_n when the receiver is in the correct lock state ($\delta_\tau \approx 0$), where $p_n = p_p$.
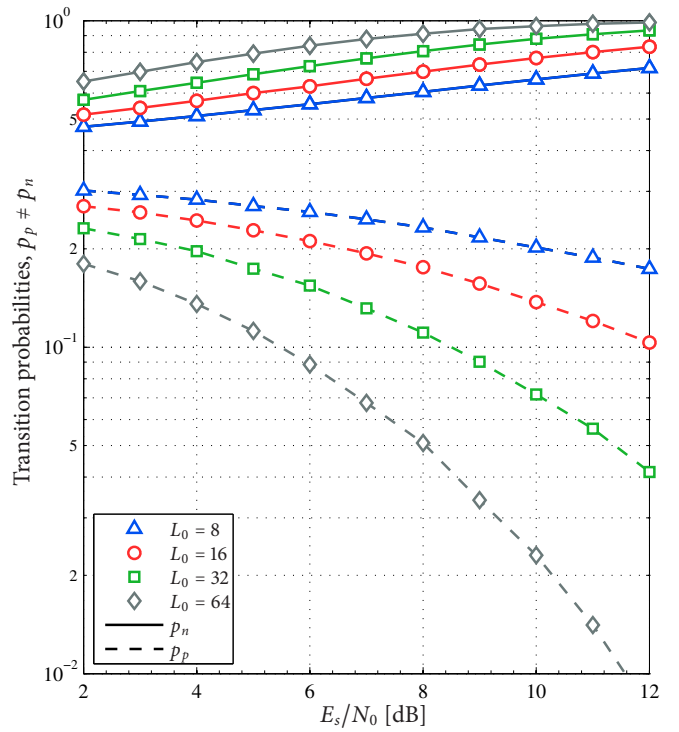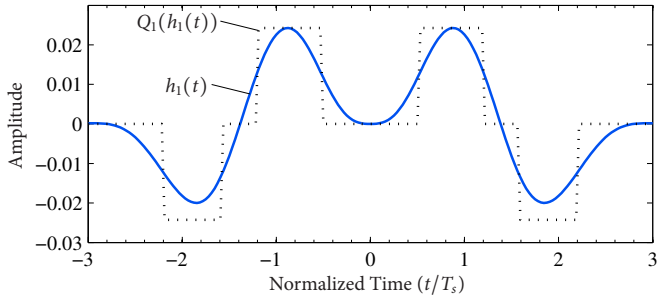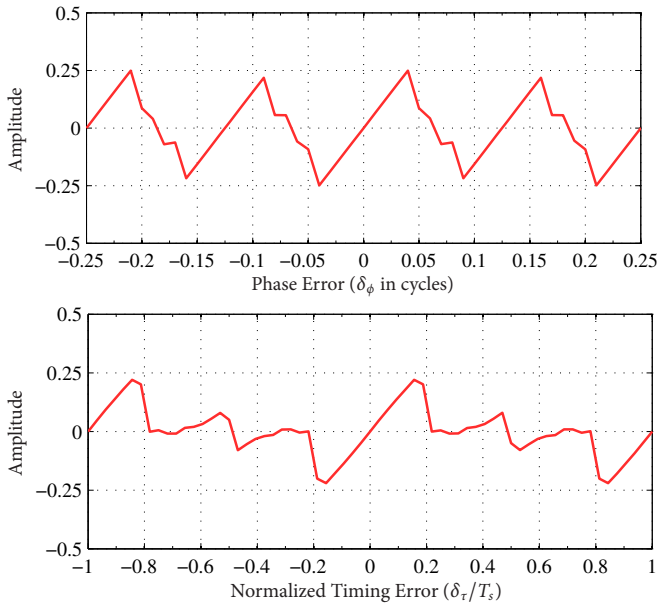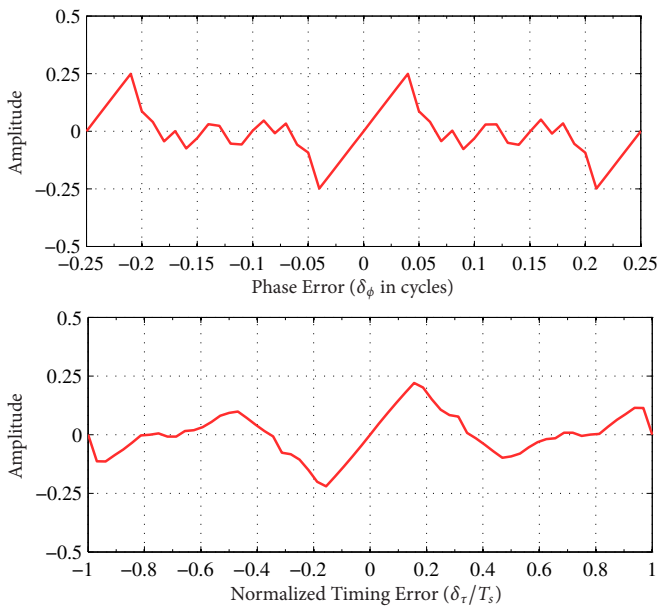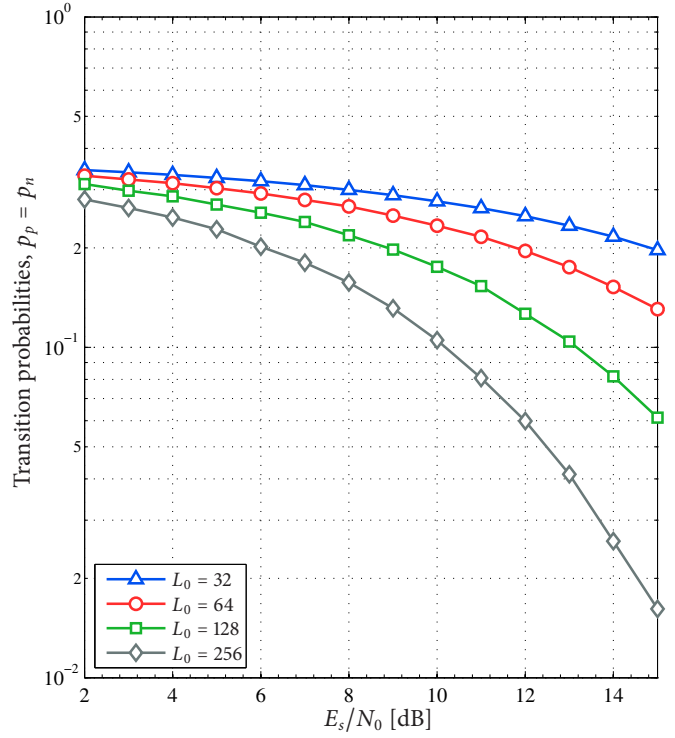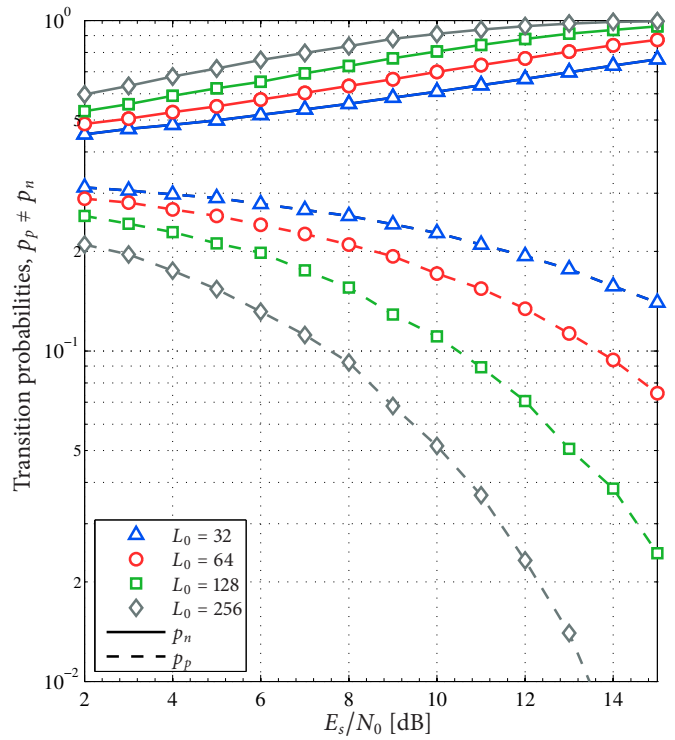

FIGURE 15 | Scheme 1: The probabilities p_p and p_n when the receiver is in the false lock state of $\delta_\tau \approx +\delta_F$, where $p_n \gg p_p$. The situation is reversed in the other false lock state of $\delta_\tau \approx -\delta_F$, where $p_n \ll p_p$ (not shown).
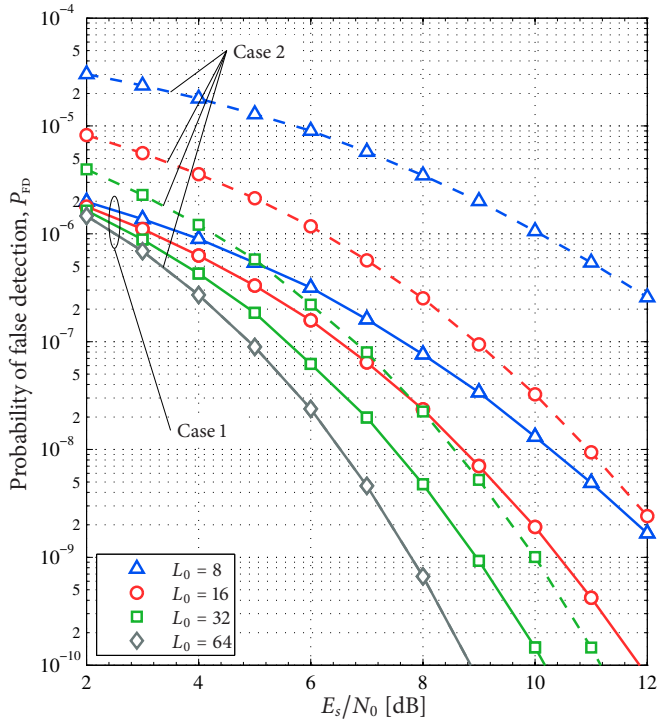
FIGURE 16 | Scheme 1: The probability of false detection, P_{FD}.

FIGURE 18 | Scheme 1: The bit error rate (BER).
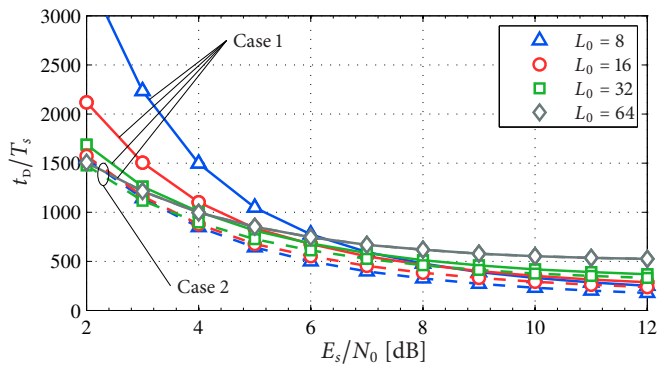
**FIGURE 17 |** Scheme 1: Acquisition time, t_{D}/T_s.
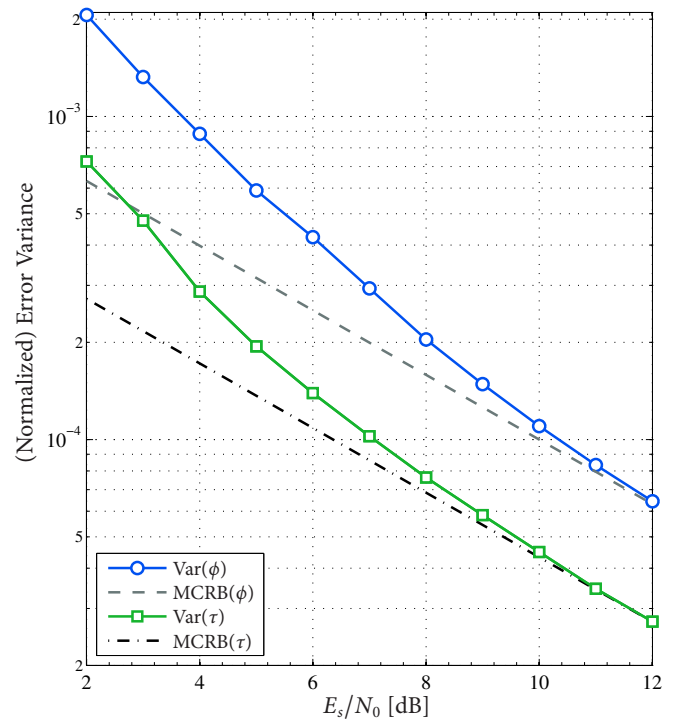
**FIGURE 19 |** Scheme 1: The phase error variance $\mathrm{Var}(\phi)$ and normalized timing error variance $\mathrm{Var}(\tau)$.
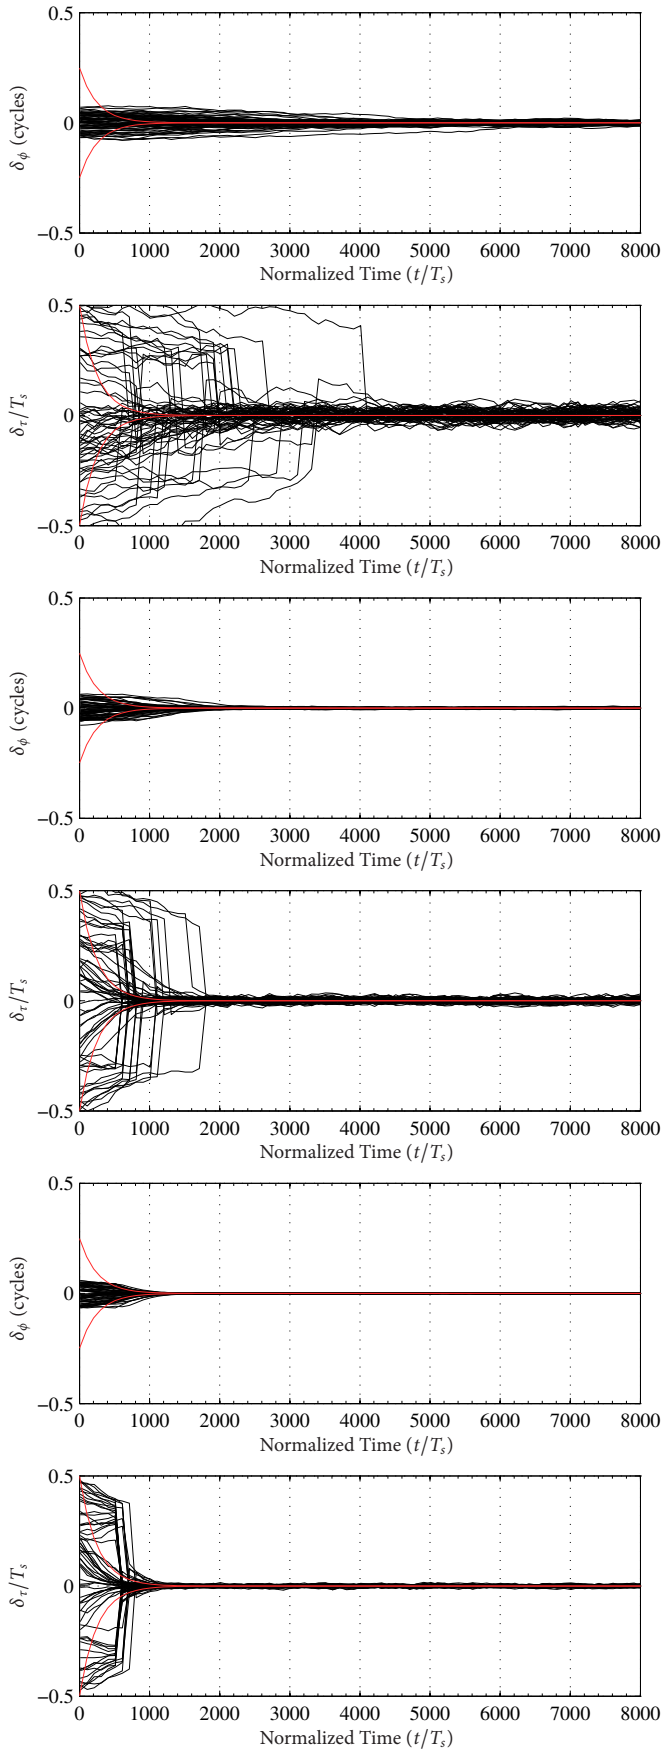
FIGURE 20 | Scheme 1: Phase error (δ_ϕ, in cycles) and normalized timing error (δ_τ/T_s) at three different E_s/N_0: 2 dB, 7 dB, and 12 dB (top to bottom, respectively).
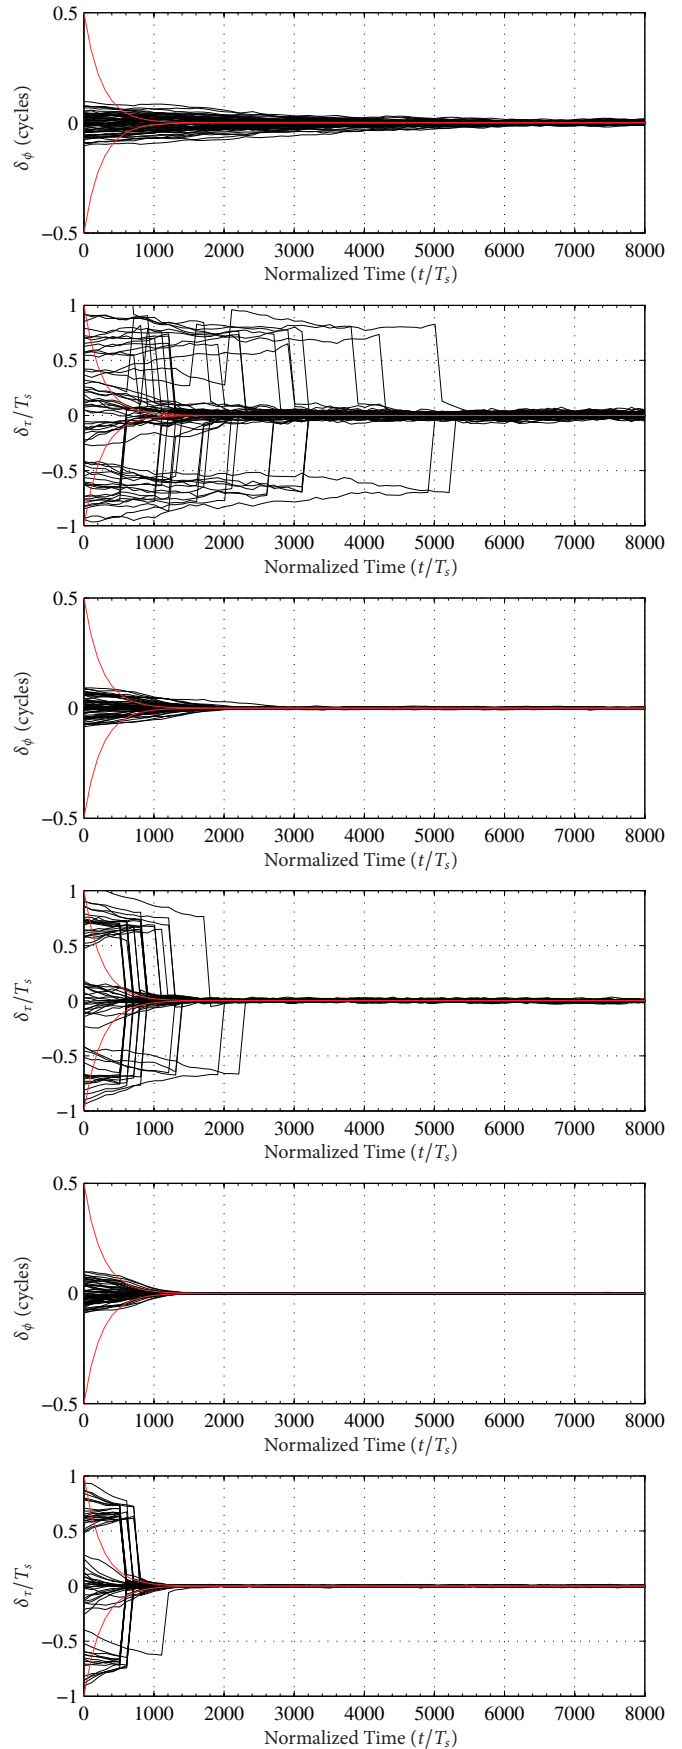
FIGURE 21 | Scheme 1: Phase error (δ_ϕ, in cycles) and normalized timing error (δ_τ/T_s) with tilted phase at three different E_s/N_0: 2 dB, 7 dB, and 12 dB (top to bottom, respectively).

FIGURE 22 | Scheme 2 ($M = 4$, $h = 1/4$, 2REC): Filter response $h_1(t)$ and quantized version $Q_1(h_1(t))$.


FIGURE 23 | Scheme 2: S-curves for the DD PED (top) and TED (bottom) in [5].

FIGURE 24 | Scheme 2: S-curves for the DD PED (top) and TED (bottom) in [5] for the tilted-phase model.

FIGURE 25 | Scheme 2: The probabilities p_p and p_n when the receiver is in the correct lock state ($\delta_\tau \approx 0$), where $p_n = p_p$.

FIGURE 26 | Scheme 2: The probabilities p_p and p_n when the receiver is in the false lock state of $\delta_\tau \approx +\delta_F$, where $p_n \gg p_p$. The situation is reversed in the other false lock state of $\delta_\tau \approx -\delta_F$, where $p_n \ll p_p$ (not shown).
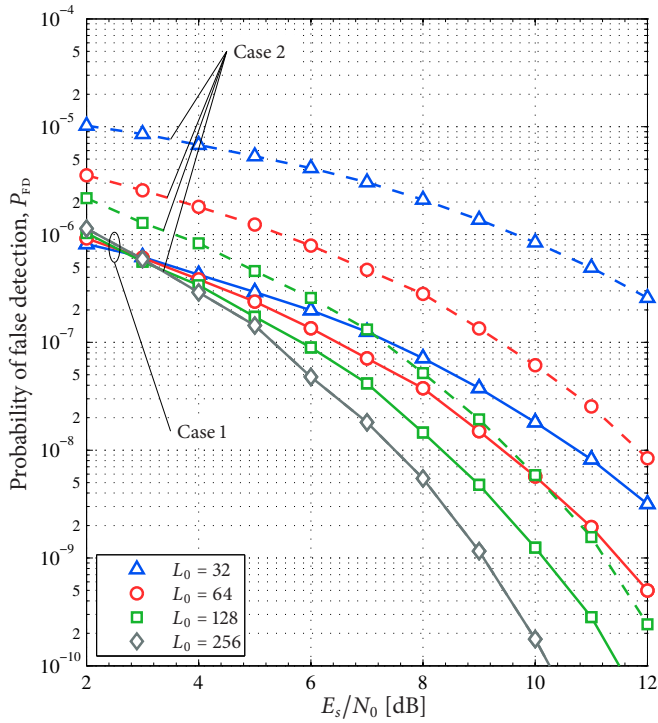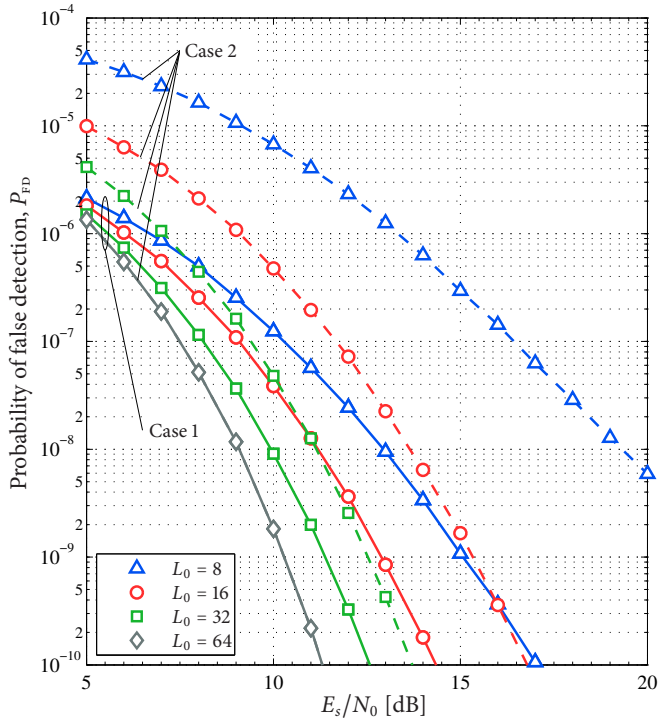
**FIGURE 27 |** Scheme 2: The probability of false detection, P_{FD}.
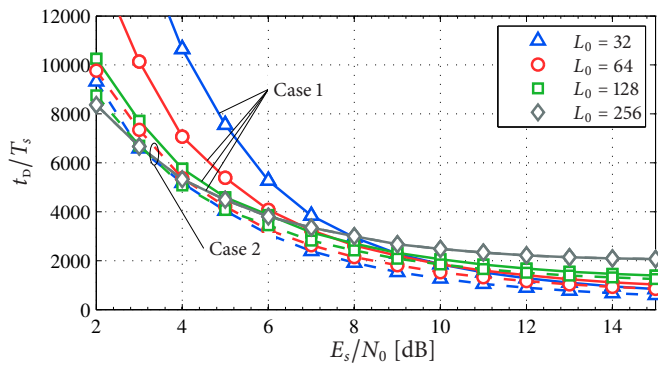
FIGURE 28 | Scheme 2: Acquisition time, t_{D}/T_s.
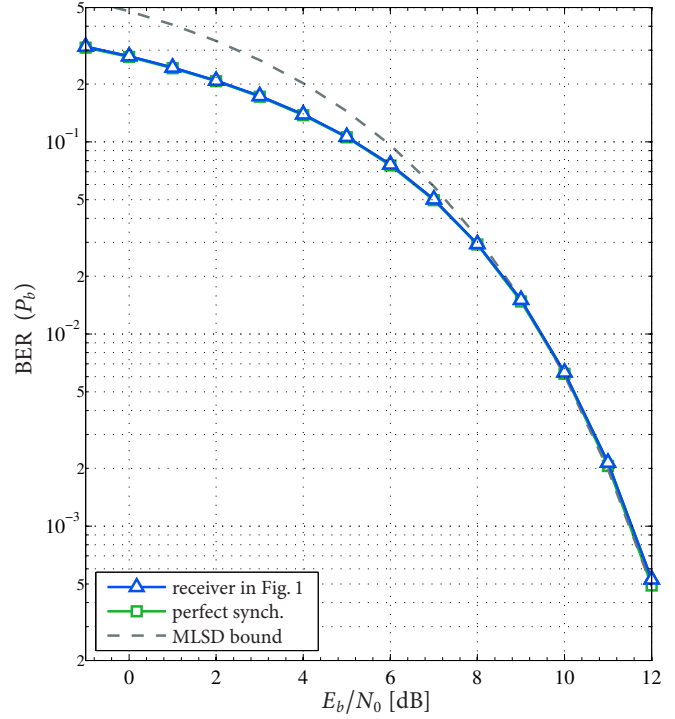
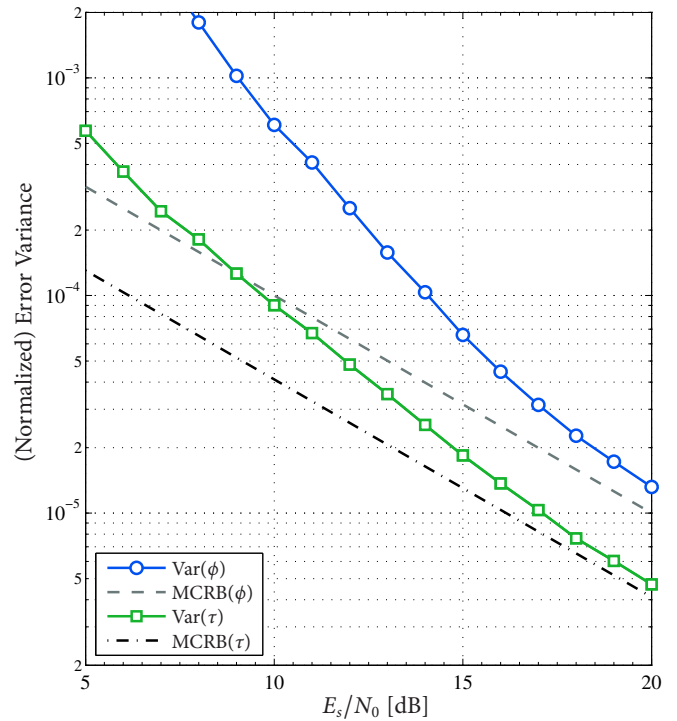

FIGURE 29 | Scheme 2: The bit error rate (BER).
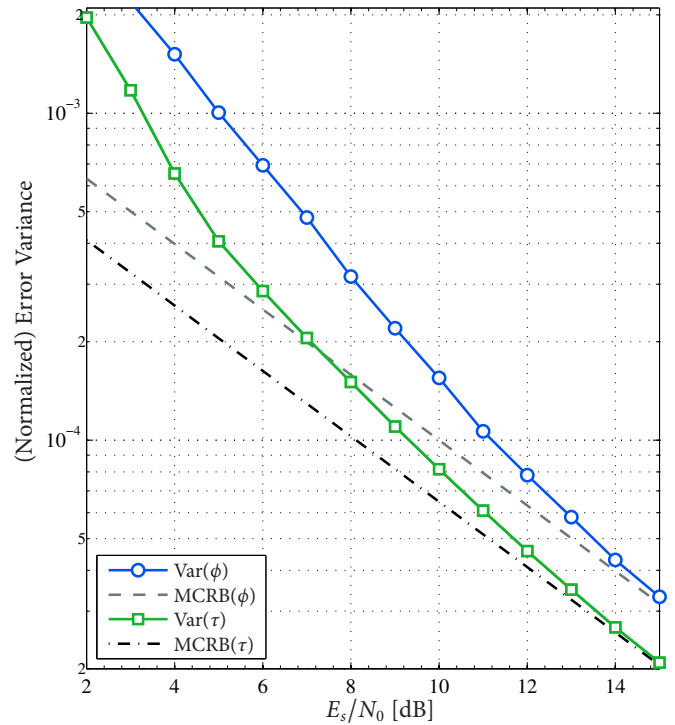
FIGURE 30 | Scheme 2: The phase error variance $\mathrm{Var}(\phi)$ and normalized timing error variance $\mathrm{Var}(\tau)$.
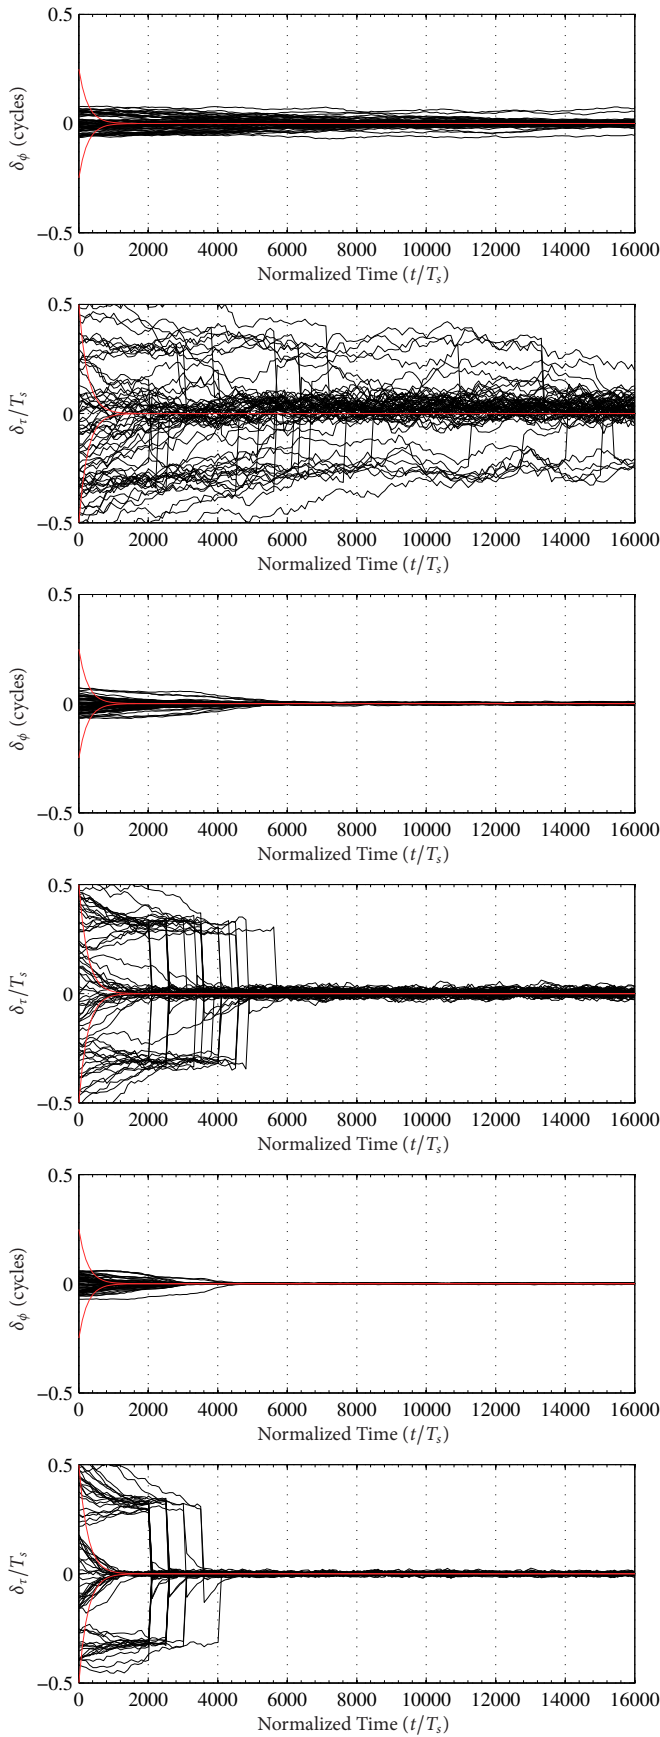
FIGURE 31 | Scheme 2: Phase error (δ_ϕ, in cycles) and normalized timing error (δ_τ/T_s) at three different E_s/N_0: 2 dB, 7 dB, and 12 dB (top to bottom, respectively).
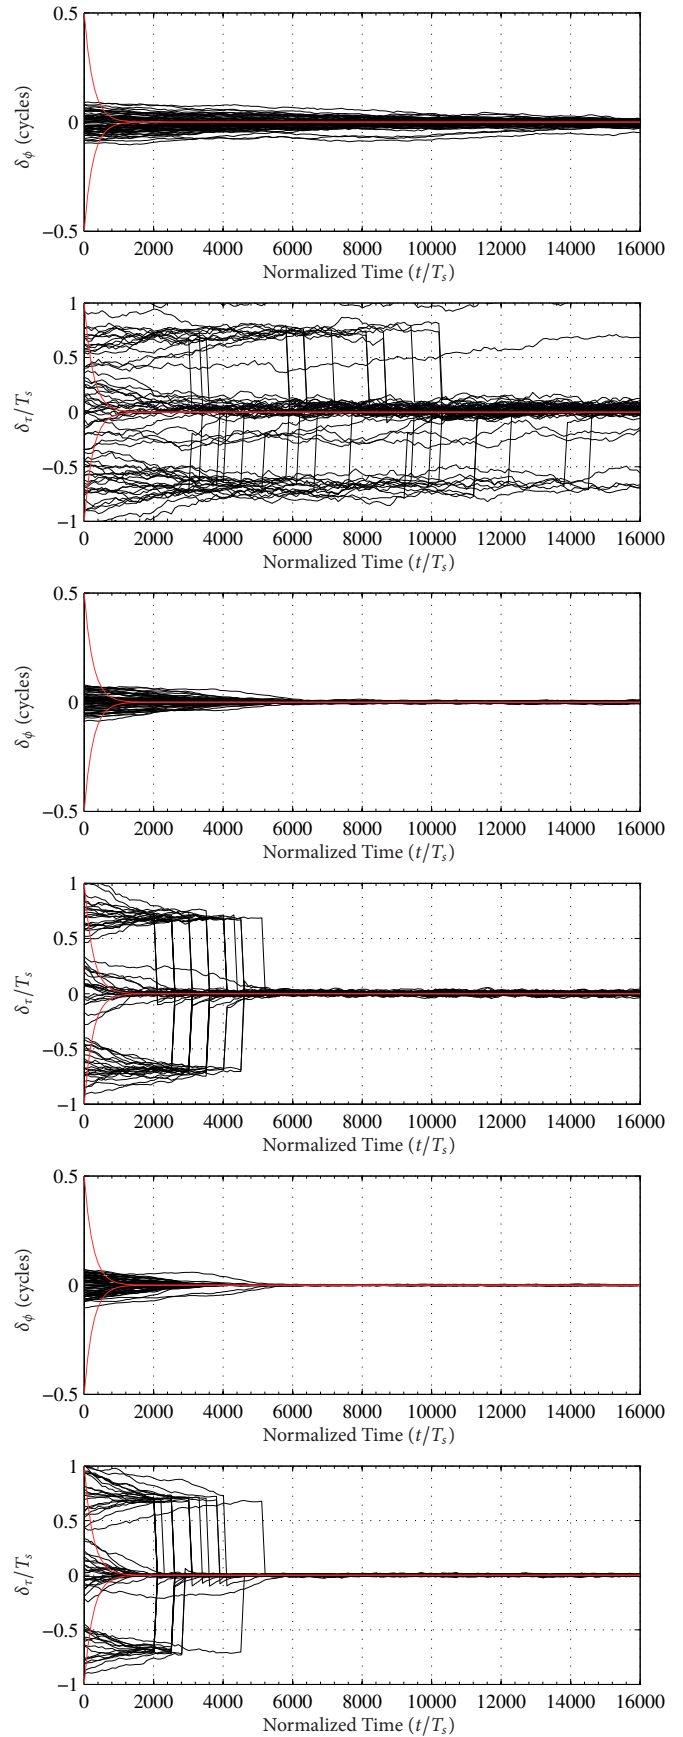
FIGURE 32 | Scheme 2: Phase error (δ_ϕ, in cycles) and normalized timing error (δ_τ/T_s) with tilted phase at three different E_s/N_0: 2 dB, 7 dB, and 12 dB (top to bottom, respectively).
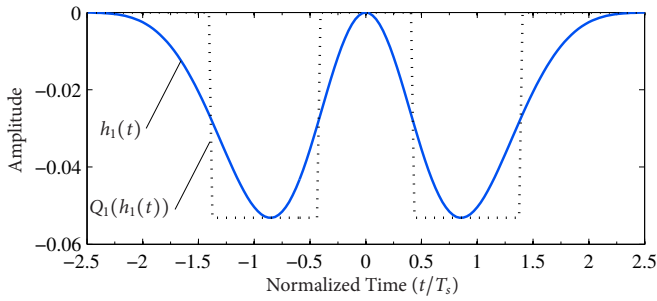
FIGURE 33 | Scheme 3 ($M = 8$, $h = 1/8$, 2RC): Filter response $h_1(t)$ and quantized version $Q_1(h_1(t))$.
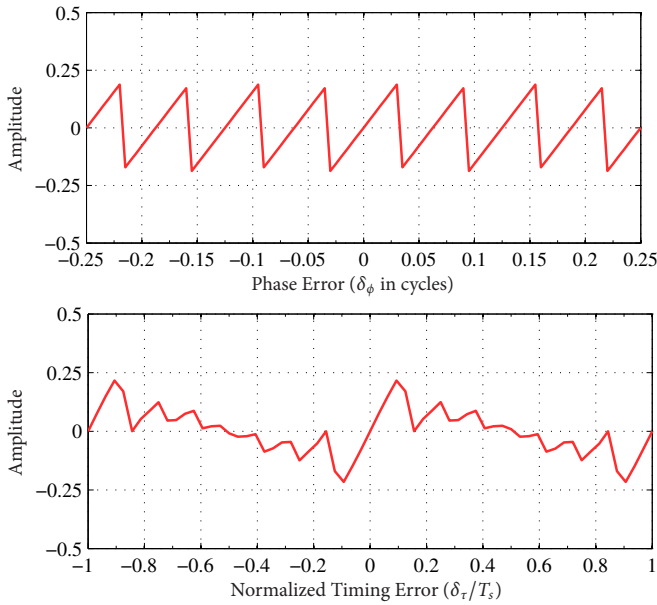
FIGURE 34 | Scheme 3: S-curves for the DD PED (top) and TED (bottom) in [5].
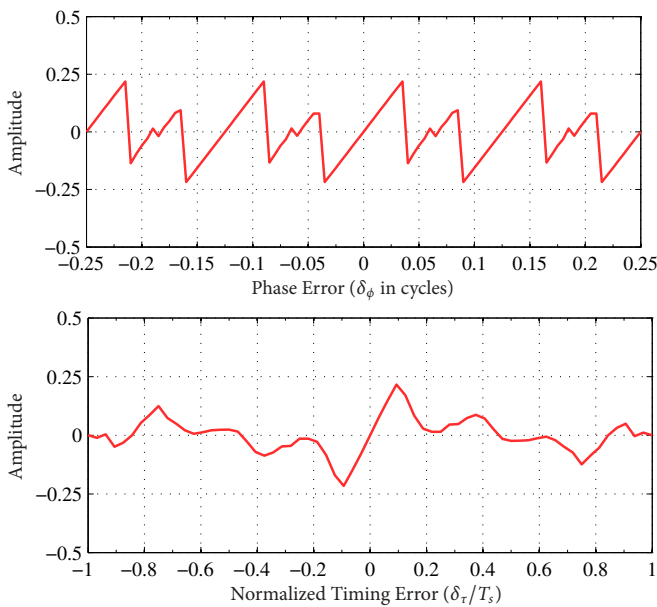
FIGURE 35 | Scheme 3: S-curves for the DD PED (top) and TED (bottom) in [5] for the tilted-phase model.
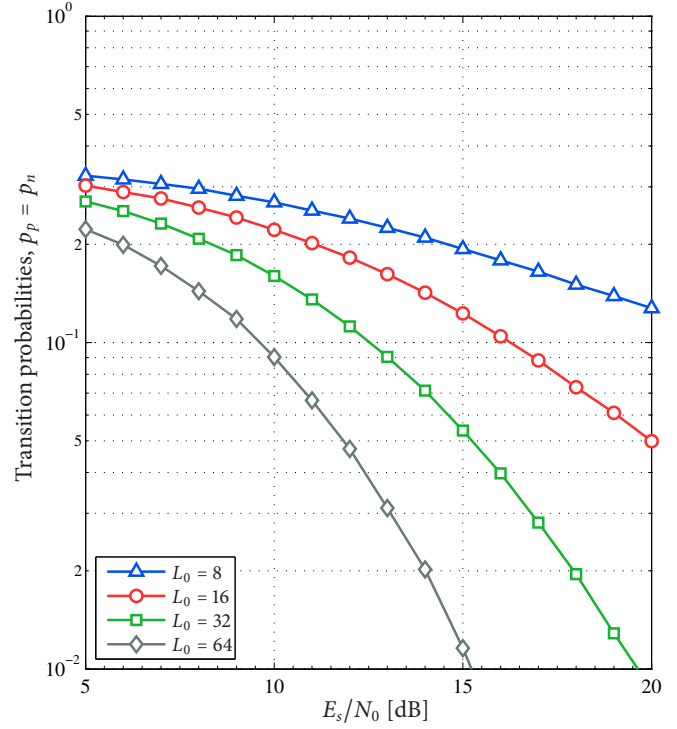
FIGURE 36 | Scheme 3: The probabilities p_P and p_n when the receiver is in the correct lock state ($\delta_\tau \approx 0$), where $p_n = p_P$.
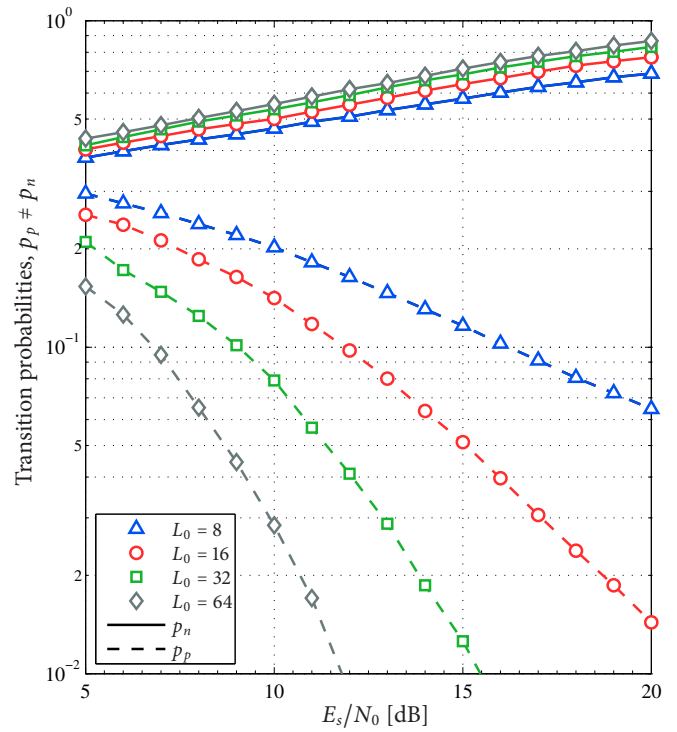
FIGURE 37 | Scheme 3: The probabilities p_P and p_n when the receiver is in the false lock state of $\delta_\tau \approx +\delta_F$, where $p_n \gg p_P$. The situation is reversed in the other false lock state of $\delta_\tau \approx -\delta_F$, where $p_n \ll p_P$ (not shown).

FIGURE 38 | Scheme 3: The probability of false detection, P_{FD}.

FIGURE 39 | Scheme 3: Acquisition time, t_{D}/T_s.
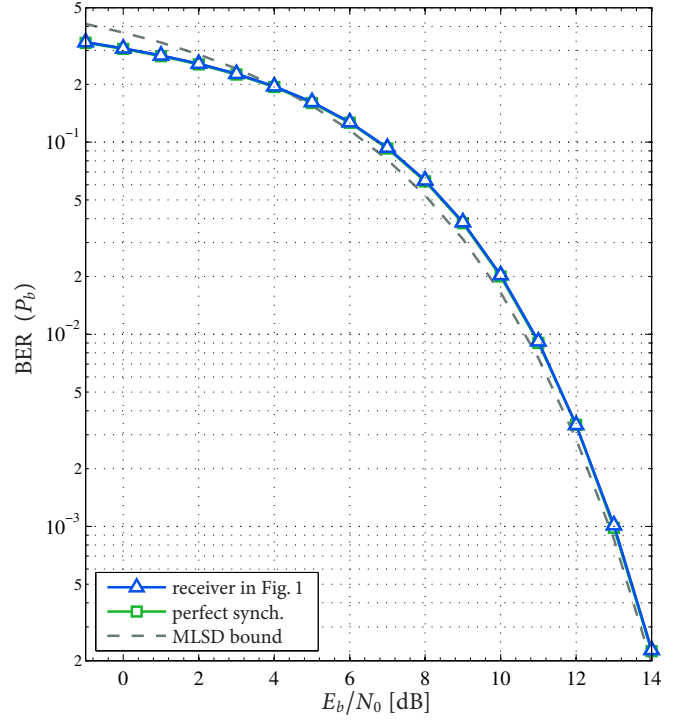
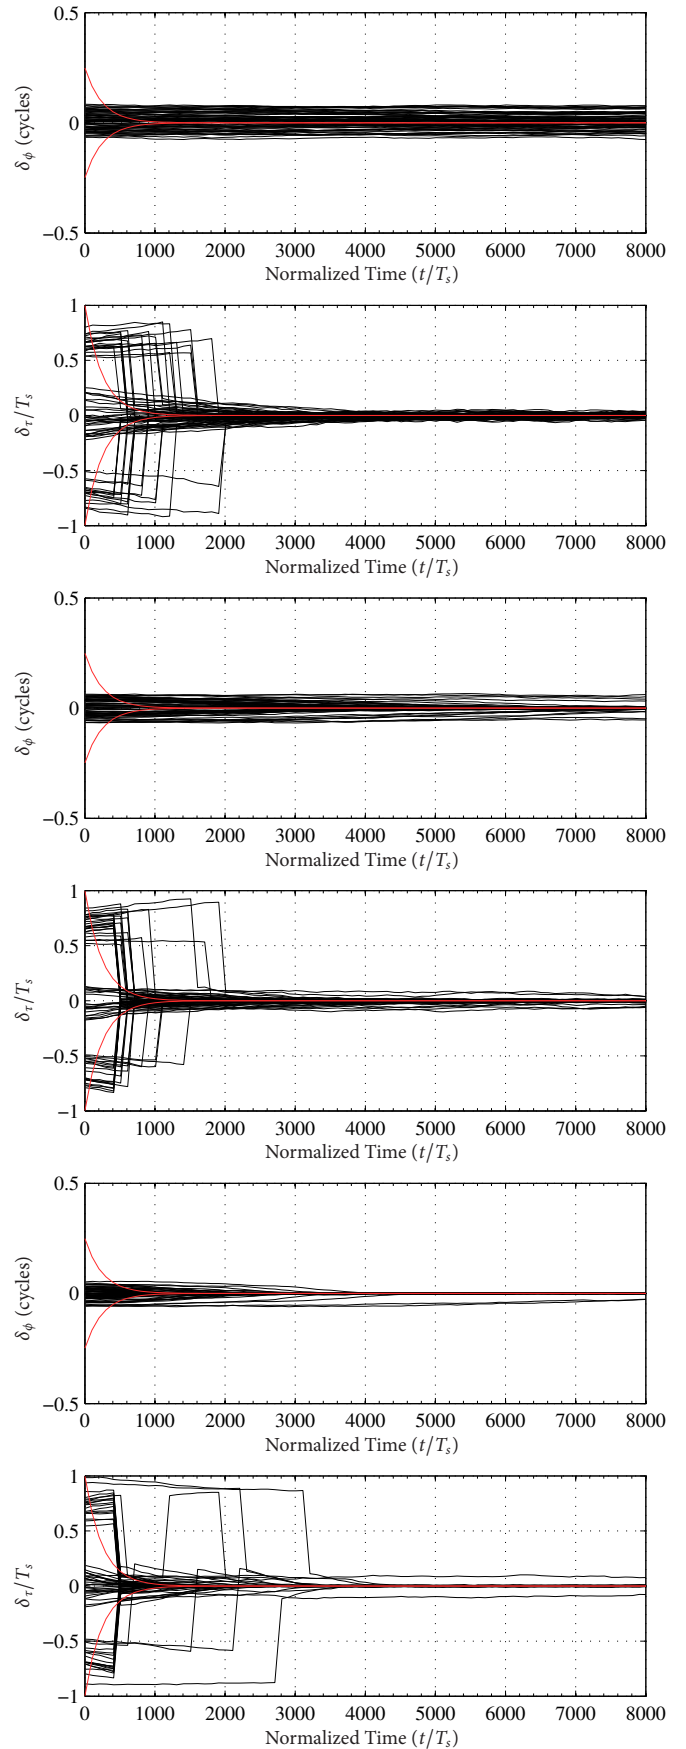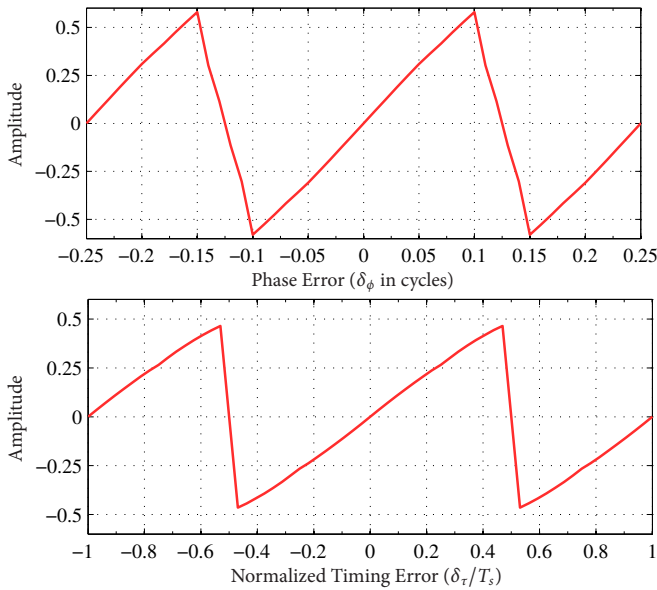**FIGURE 40 |** Scheme 3: The bit error rate (BER).

FIGURE 41 | Scheme 3: The phase error variance $\text{Var}(\phi)$ and normalized timing error variance $\text{Var}(\tau)$.
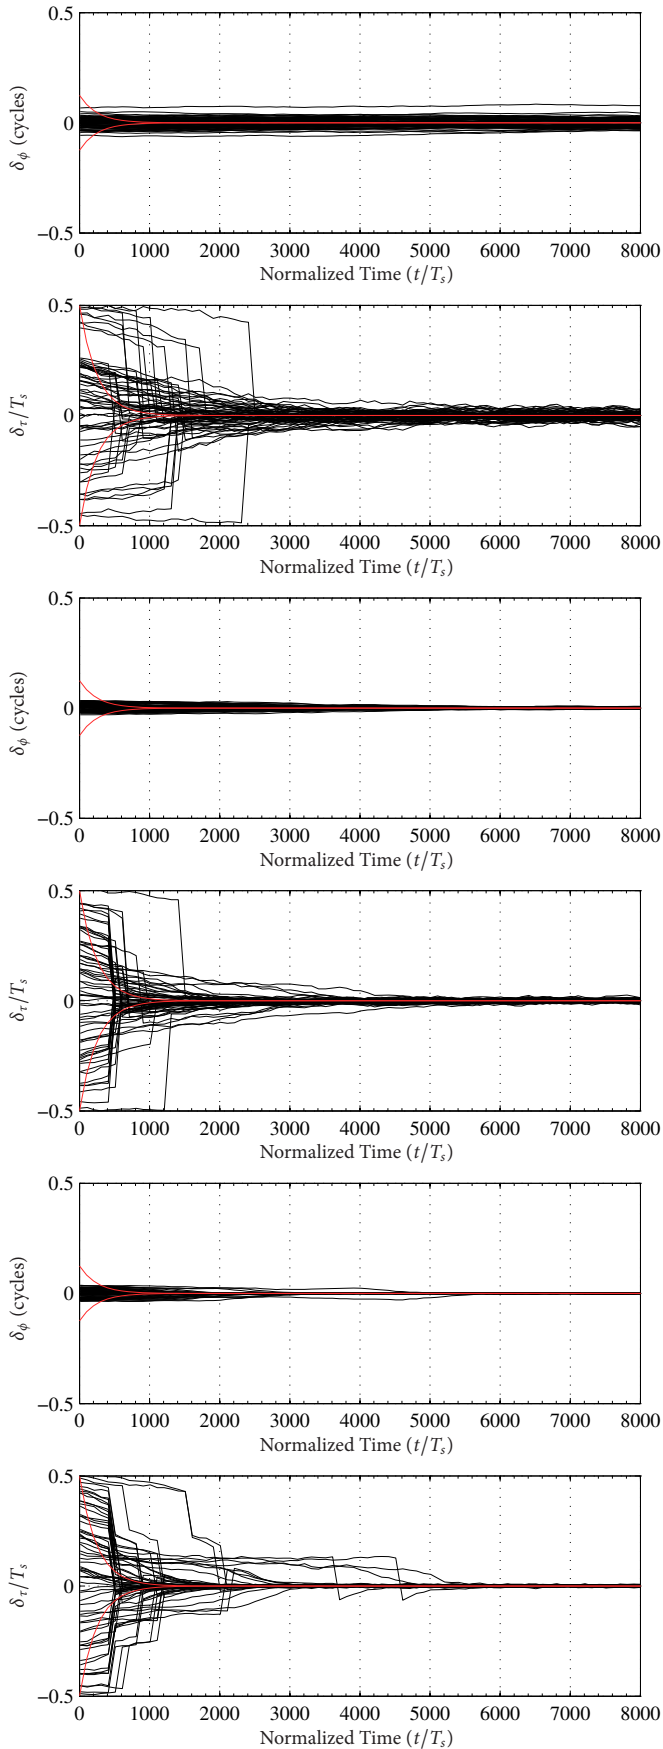
FIGURE 42 | Scheme 3: Phase error (δ_ϕ, in cycles) and normalized timing error (δ_τ/T_s) at three different E_s/N_0: 5 dB, 10 dB, and 15 dB (top to bottom, respectively).

FIGURE 43 | Scheme 3: Phase error (δ_ϕ, in cycles) and normalized timing error (δ_τ/T_s) with tilted phase at three different E_s/N_0: 5 dB, 10 dB, and 15 dB (top to bottom, respectively).

FIGURE 44 | Scheme 4 ($M = 2$, $h = 1/2$, 4G): S-curves for the DD PED (top) and TED (bottom) in [5].
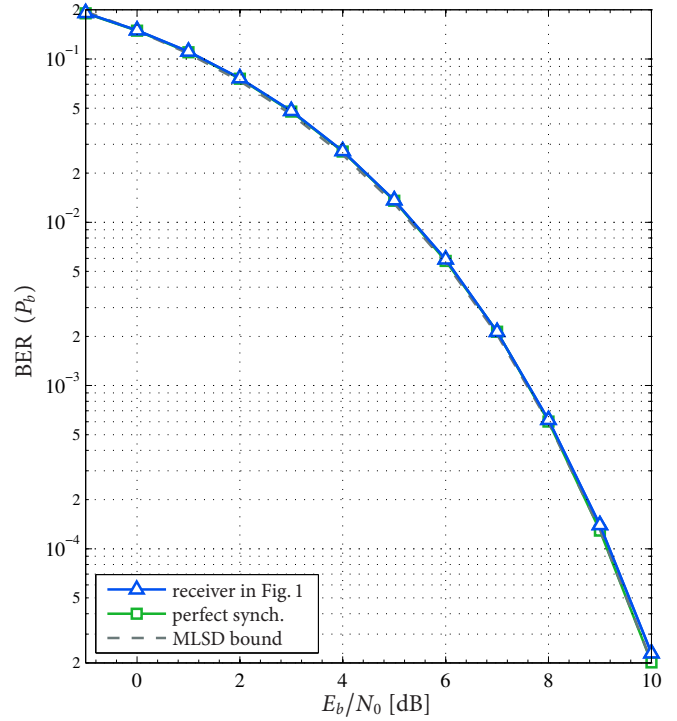
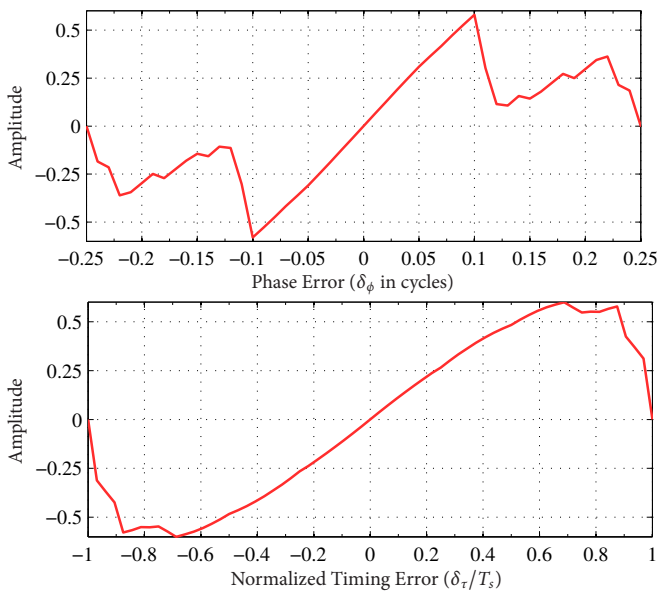**FIGURE 46 |** Scheme 4: The bit error rate (BER).

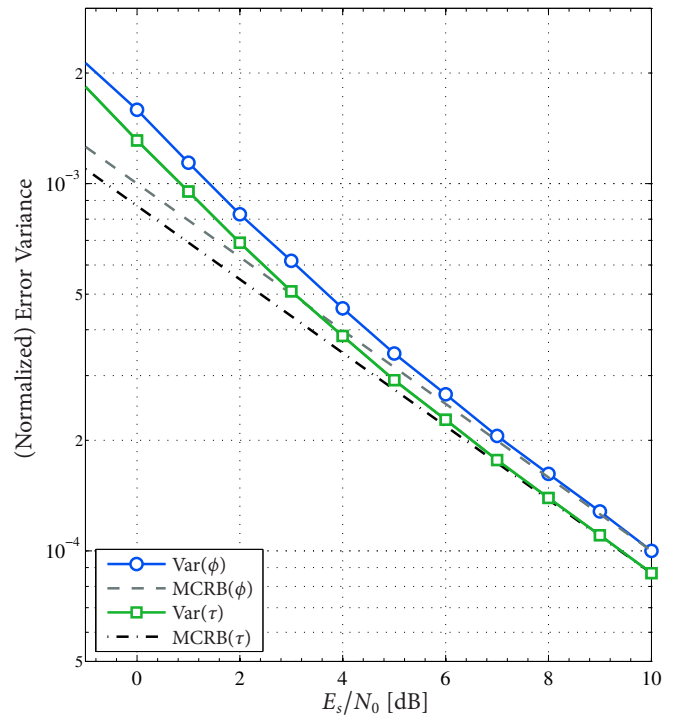**FIGURE 45 |** Scheme 4: S-curves for the DD PED (top) and TED (bottom) in [5] for the tilted-phase model.
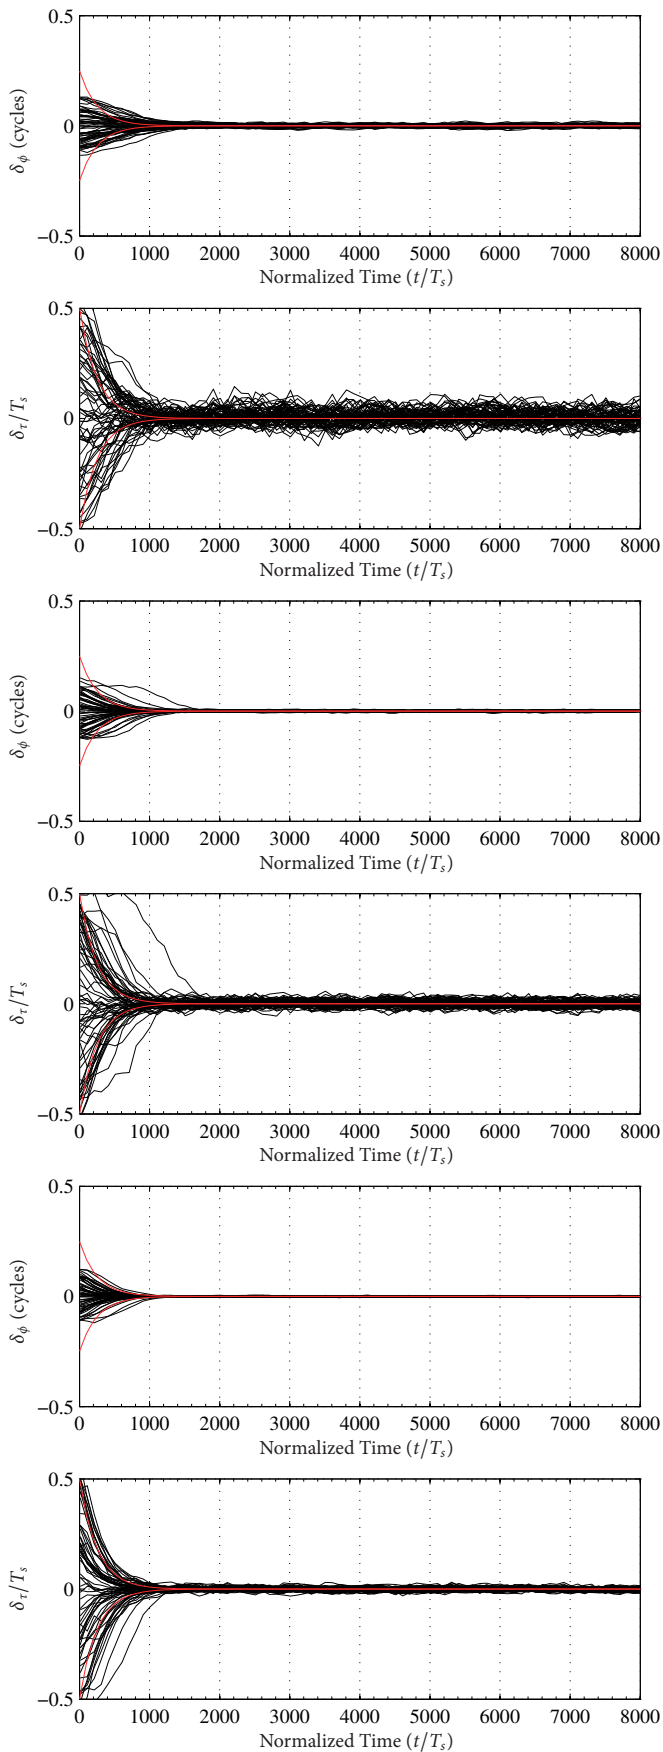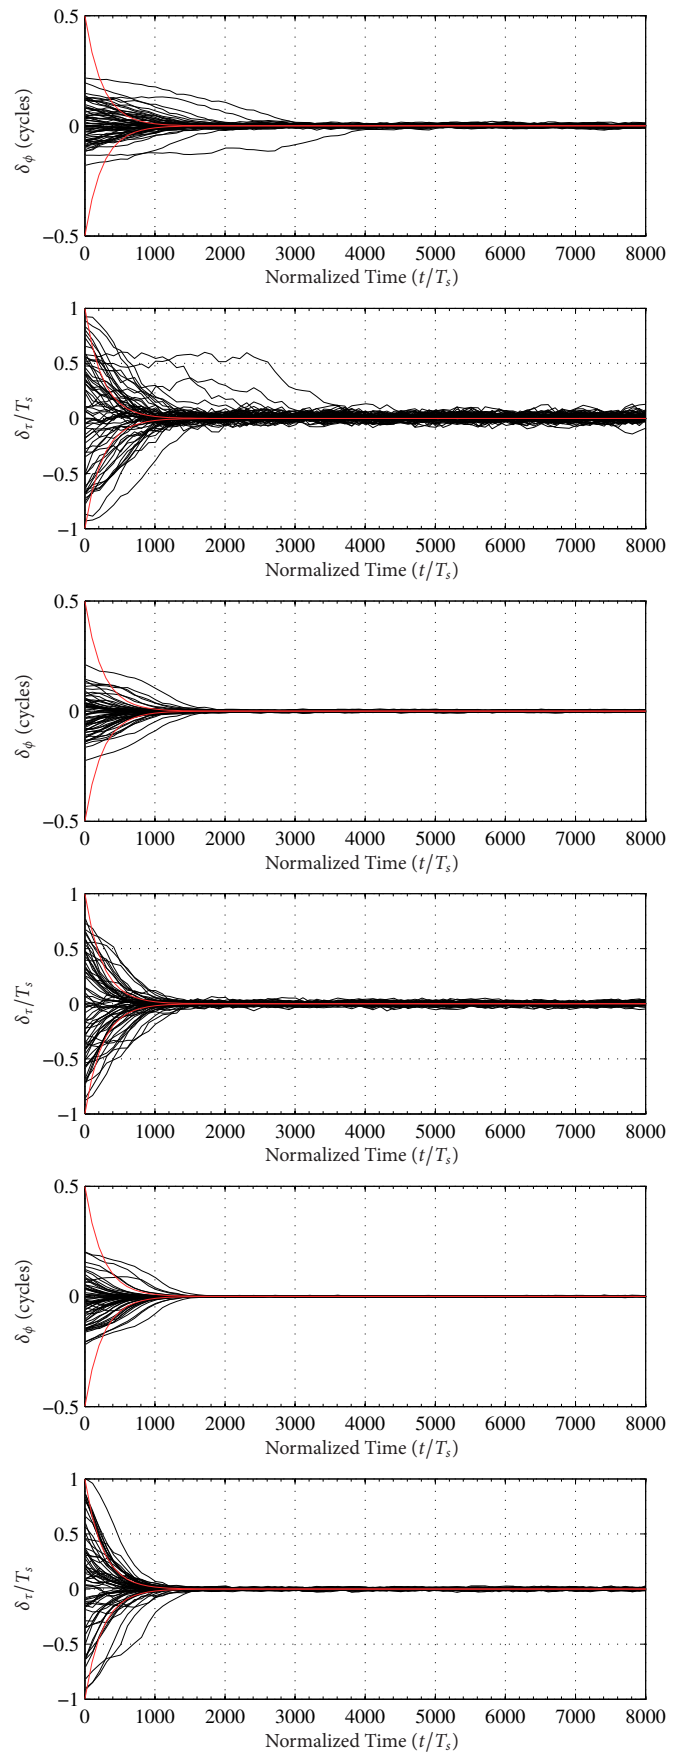

FIGURE 47 | Scheme 4: The phase error variance Var(ϕ) and normalized timing error variance Var(τ).

FIGURE 48 | Scheme 4: Phase error (δ_ϕ, in cycles) and normalized timing error (δ_τ/T_s) at three different E_s/N_0: 0 dB, 5 dB, and 10 dB (top to bottom, respectively).

FIGURE 49 | Scheme 4: Phase error (δ_ϕ, in cycles) and normalized timing error (δ_τ/T_s) with tilted phase at three different E_s/N_0: 0 dB, 5 dB, and 10 dB (top to bottom, respectively).