Low-Complexity Real-Time Single-Tone Phase and Frequency Estimation

D. Richard Brown III, Yizheng Liao, and Neil Fox

***Abstract*—This paper presents a low-complexity real-time single-tone phase and frequency estimation technique based on zero-crossing detection and linear regression. The proposed zero-crossing phase and frequency estimator fills a gap between low-complexity phase locked loop estimation and high-performance maximum likelihood estimation. Similar to a phase locked loop, the zero-crossing phase and frequency estimator offers low-complexity sample-by-sample operation appropriate for real-time applications like distributed transmit beamforming. Numerical results demonstrate, however, that the proposed technique significantly exceeds the performance of phase locked loop estimation, closely tracking the Cramer-Rao lower bound over a wide range of signal to noise ratios.**

I. Introduction

Distributed transmit beamforming is a technique in which two or more single-antenna transmitters simultaneously transmit with phase-aligned carriers such that the passband signals coherently combine at an intended destination. The transmitters in a distributed transmit beamformer form a *virtual antenna array* and, in principle, can achieve all of the gains of a conventional antenna array, e.g. increased range, rate, and/or energy efficiency, without the size, cost, and complexity of a conventional antenna array.

Unlike a conventional antenna array, a distributed transmit beamformer requires precise carrier synchronization among the transmitters with timing errors significantly smaller than a carrier period. Several carrier synchronization techniques have recently been proposed to facilitate distributed transmit beamforming including full-feedback closed-loop [1], one-bit closed-loop [2]–[4], master-slave open-loop [5], round-trip open-loop carrier synchronization [6], [7], and two-way open-loop carrier synchronization [8]. Each of these techniques has advantages and disadvantages in particular applications, as discussed in the survey article [9].

A common feature of several of these techniques is that the nodes participating in the beamformer must be able to accurately and quickly estimate the phase and frequency of one or more sinusoidal beacons received from other nodes in the network. It is often assumed that the nodes use maximum likelihood estimators (MLEs) or phase locked loops (PLLs) to obtain the necessary phase and frequency estimates. While the MLE is appealing because of its potential for near-optimal performance, computation of maximum likelihood estimates typically requires high computational complexity. The maximum likelihood estimator is also a batch processing technique in the sense that the processing begins only after all of the samples in the observation have been received. The PLL, on the other hand, is appealing for its low-complexity, sample-by-sample operation, but tends to provide phase and frequency estimates with worse performance than the MLE.

This paper presents a new technique for low-complexity real-time single-tone phase and frequency estimation based on zero-crossing detection and linear regression. The proposed technique is similar to the linear regression estimation techniques described in [10] and [11] except that it avoids the computational burden of inverse tangents and phase unwrapping by only using information contained in the zero-crossings of the observed signal. The appeal of the proposed *zero-crossing phase and frequency estimator* is that it offers low-complexity sample-by-sample operation similar to a phase locked loop. Numerical results demonstrate that the proposed technique also offers near-optimal performance over a wide range of signal to noise ratios (SNRs).

II. Observation Model

We consider the signal model of [12] with a received signal given as

$$z(t) = b\exp(j(\omega t + \theta)) + w(t) \tag{1}$$

where b, ω, and θ denote the unknown amplitude, frequency, and phase of the signal, respectively, and $w(t)$ denotes zero-mean proper complex additive white Gaussian noise. The received signal is sampled at a constant sampling rate $f_s = 1/T$ to produce the discrete-time observation

$$z[n] = z(t_0 + nT) = b\exp(j(\omega(t_0 + nT) + \theta)) + \eta[n] \tag{2}$$

for $n = 0, \ldots, N-1$ where t_0 denotes the time of the first sample and $\eta[n]$ is a zero-mean proper complex Gaussian random variable with $\mathrm{var}\{\mathrm{Re}(\eta[k])\} = \mathrm{var}\{\mathrm{Im}(\eta[k])\} = \sigma^2$ and $\mathrm{cov}\{\mathrm{Re}(\eta[k]), \mathrm{Im}(\eta[k])\} = 0$. It is assumed that $\eta[n]$ are independent and identically distributed (i.i.d.) for $n = 0, \ldots, N-1$.

The N-sample observation of (2) is provided as an input to a phase and frequency estimator. Depending on the type of estimator, the computations can be performed in real-time on a sample-by-sample basis, i.e. as each sample is received, or as a batch operation on the entire received signal vector after all

D.R. Brown III is an Associate Professor with the Electrical and Computer Engineering Department, Worcester Polytechnic Institute, Worcester, MA 01609 USA. e-mail: drb@ece.wpi.edu.

Yizheng Liao is an undergraduate student in the Electrical and Computer Engineering Department, Worcester Polytechnic Institute, Worcester, MA 01609 USA. e-mail: liaoyizheng@wpi.edu.

Neil Fox is a Program Manager with DARPA's Strategic Technology Office, Arlington, VA 22203. e-mail: Neil.Fox@darpa.mil.

This work was supported by NSF award CCF-0447743.

of the samples have been received. The phase and frequency estimates generated by the estimator are denoted as $\hat{\theta}$ and $\hat{\omega}$, respectively, and the resulting phase and frequency errors are denoted as $\tilde{\theta} := \theta - \hat{\theta}$ and $\tilde{\omega} := \omega - \hat{\omega}$, respectively.

III. Maximum Likelihood Phase and Frequency Estimation

Given a joint density of the observation, the maximum likelihood estimator seeks to find the value of the unknown parameter(s) that maximizes the likelihood equation, i.e.

$$\hat{\lambda} = \arg\max_{\lambda \in \Lambda} \log p_Z(z\,;\,\lambda)$$

where $p_Z(z\,;\,\lambda)$ is the joint density of the observation parameterized by λ. When the observations are i.i.d. Gaussian distributed, the maximum likelihood estimator possesses three desirable asymptotic properties as the number of samples in the observation becomes large: asymptotic unbiasedness, asymptotic efficiency, and asymptotically Gaussian estimation errors. In other words, when the number of samples in the observation becomes large, the joint statistics of the maximum likelihood phase and frequency estimation errors are approximately zero-mean Gaussian distributed with covariance given in (13).

Upon receiving the discrete time observations according to (2) for $n = 0, \ldots, N-1$, the maximum likelihood frequency estimate can be computed as [12]

$$\hat{\omega} = \arg\max_{\omega} |A(\omega)| \tag{3}$$

where

$$A(\omega) = \frac{1}{N} \sum_{n=0}^{N-1} z[n] \exp(-jn\omega T). \tag{4}$$

Once the maximum likelihood frequency estimate has been computed, the maximum likelihood phase estimate follows as

$$\hat{\theta} = \text{angle}\left\{\exp(-j\hat{\omega}t_0)A(\hat{\omega})\right\}. \tag{5}$$

Computation of (5) requires a four-quadrant $\arctan$ operation which can be achieved with reasonable complexity through the use of a lookup table. Computation of (3), however, requires maximization over a multimodal objective function. Simple single-parameter maximization techniques, e.g. gradient ascent, can converge to local maxima if not initialized close to the solution. This can lead to poor estimation accuracy.

Given the multimodal nature of $|A(\omega)|$ and recognizing that the M-point DFT is a sampled version of $A(\omega)$ at frequencies $\omega = \frac{2\pi k}{MT}$ for $k = 0, \ldots, M-1$, a common approach to maximum likelihood frequency estimation is to use the fast Fourier transform (FFT) to efficiently compute $A\left(\frac{2\pi k}{MT}\right)$ for $k = 0, \ldots, M-1$ and to then select the index k at which $A\left(\frac{2\pi k}{MT}\right)$ attains its maximum magnitude, i.e.,

$$\hat{\omega} = \arg\max_{\omega \in \Omega} |A(\omega)| \tag{6}$$

where $\Omega := \{0, \frac{2\pi}{MT}, \ldots, \frac{2\pi(M-1)}{MT}\}$. Using an M-point FFT to compute the maximum likelihood frequency estimate, however, establishes a tradeoff between estimation accuracy and computational complexity. The achievable estimation accuracy of the FFT-based maximum likelihood frequency estimator in (6) depends on the frequency resolution of the M-point FFT, i.e. $\frac{f_s}{M}$, whereas the asymptotic complexity of an M-point FFT is $M \log_2(M)$. In low-SNR scenarios, i.e. when $\frac{b}{2\sigma^2}$ is small, fine frequency resolution is not necessary because the performance of the MLE is dominated by noise. In high-SNR scenarios, however, the performance of the MLE may be limited by the frequency resolution of the M-point FFT.

Rather than selecting the frequency estimate from a discrete set as in (6), the FFT magnitude can be interpolated to improve the estimation accuracy. A full sinc interpolation of the FFT is computationally infeasible, but simpler interpolation techniques can often be used with acceptable results. For example, suppose the maximum in (6) occurs at FFT index $\hat{k}$. A quadratic fit $y = ax^2 + bx + c$ in the neighborhood of the maximum can be computed given the frequencies $x \in \left\{\frac{2\pi(\hat{k}-1)}{MT}, \frac{2\pi\hat{k}}{MT}, \frac{2\pi(\hat{k}+1)}{MT}\right\}$ and FFT magnitudes $y = |A(x)|$. The maximum likelihood frequency estimate can then be computed using standard calculus techniques as $\hat{\omega} = \frac{-b}{2a}$.

An additional difficulty with applying the FFT-based MLE to carrier synchronization in distributed transmit beamforming systems is that the FFT is a batch operation that requires the entire observation to be present in order to begin processing. Hence, a high-resolution FFT may incur significant processing delay after the conclusion of the observation, leading to potential drift in open-loop synchronization systems. This is in contrast to real-time sample-by-sample estimation techniques such as the phase locked loop and zero-crossing phase and frequency estimator, described in the following sections.

IV. PLL Phase and Frequency Estimation

Phase locked loops (PLLs) are a well-known technique for synchronizing a local oscillator to an external reference signal [13]. Although PLLs are not usually considered in the context of phase and frequency estimation, they can be used as a computationally efficient method for extracting the phase and frequency from an observation of a single-tone signal in noise.

A block diagram of a quadrature PLL is shown in Figure 1. The PLL is composed of three elements: a phase detector, a loop filter, and a voltage controlled oscillator (VCO). In the absence of noise, the quadrature phase detector output is

$$\begin{aligned} v[n] &= K_d \text{Im}\left\{b\exp(j(\omega nT + \theta))\exp(-j\phi_o[n])\right\} \\ &= K_d b \sin\left(\Delta[n]\right) \end{aligned}$$

where $\Delta[n]$ is the phase difference between the input and feedback signals and K_d is the phase detector gain parameter.

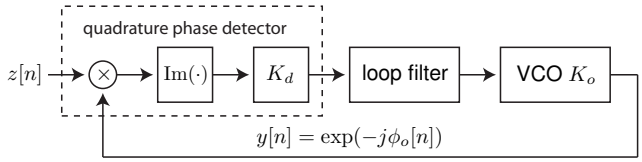

Fig. 1. Quadrature phase locked loop block diagram.

The loop filter serves to suppress noise at the output of the phase detector. The loop filter is typically parameterized by its 3dB bandwidth $\omega_{3\mathrm{dB}}$. Given $\omega_{3\mathrm{dB}}$, K_o, and K_d, the loop filter can be designed according to the procedures in [13] such that the PLL has the desired transient and steady-state characteristics.

The output of the loop filter $u[n]$ is connected to the input of VCO which generates a complex exponential output $y[n]$ for feedback to the phase detector at frequency $\omega_o[n] = \omega_{\mathrm{nom}} + K_o u[n]$ where ω_{nom} is the nominal frequency of the VCO in the absence of a control input and K_o is the VCO gain parameter. The VCO phase is then updated according to

$$\phi_o[n+1] = \phi_o[n] + T\omega_o[n]$$

and the feedback $y[n+1] = \exp(-j\phi_o[n+1])$ is then computed and fed back to the phase detector to be multiplied with the next input sample.

Given an N-sample observation, a frequency estimate can be obtained from the PLL by simply sampling the frequency of the VCO after the PLL has locked at the end of the observation, i.e.

$$\hat{\omega} = \omega_o[N-1] = \omega_{\mathrm{nom}} + K_o u[N-1]. \tag{7}$$

The phase estimate can be computed as

$$\hat{\theta} = \phi_o[N-1] - (N-1)T\hat{\omega}. \tag{8}$$

In other words, the PLL phase estimate is computed by subtracting the estimated accumulated phase over the observation (based on the PLL frequency estimate) from the final phase of the VCO.

When the PLL is locked and $\Delta[n]$ is small such that $\sin(\Delta[n]) \approx \Delta[n]$, the PLL can be analyzed as a linear system [13]. The discrete-time transfer function relating the PLL output phase and input phase can then be written as

$$H_\phi(z) = \frac{\Phi_o(z)}{\Phi_{in}(z)} = \frac{bK_dK_oF(z)L(z)}{1 + bK_dK_oF(z)L(z)} \tag{9}$$

where $L(z) := \frac{Tz^{-1}}{1-z^{-1}}$ and $F(z)$ is the discrete-time transfer function of the loop filter. Similarly, the discrete-time transfer function relating the PLL output frequency and input phase can be written as

$$H_\omega(z) = \frac{\Omega_o(z)}{\Phi_{in}(z)} = \frac{bK_dK_oF(z)}{1 + bK_dK_oF(z)L(z)}. \tag{10}$$

These transfer functions can be used to provide analytical predictions of the PLL phase and frequency estimation performance given a white Gaussian phase input.

The PLL operates on standard feedback control principles where the desired dynamics can be achieved through proper selection of the parameters $\omega_{3\mathrm{dB}}$, K_o, and K_d. By following the PLL design procedures in [13], however, it can be shown that the closed loop transfer function of the PLL depends only on $\omega_{3\mathrm{dB}}$. Intuitively, $\omega_{3\mathrm{dB}}$ is selected to achieve the desired tradeoff in convergence speed and the amount of noise passed through to the VCO. As shown by example in Figure 2, selecting $\omega_{3\mathrm{dB}}$ too low results in poor estimation performance because the PLL does not have enough time to converge. Setting $\omega_{3\mathrm{dB}}$ too high also results in poor estimation performance because, while the PLL converges quickly, the noise passed by the loop filter to the input of the VCO degrades the frequency and phase estimates. In Section VI, we numerically evaluate the mean squared estimation error of a third-order PLL with an active PI (proportional-integral) loop filter [13] as a function of $\omega_{3\mathrm{dB}}$ and SNR.

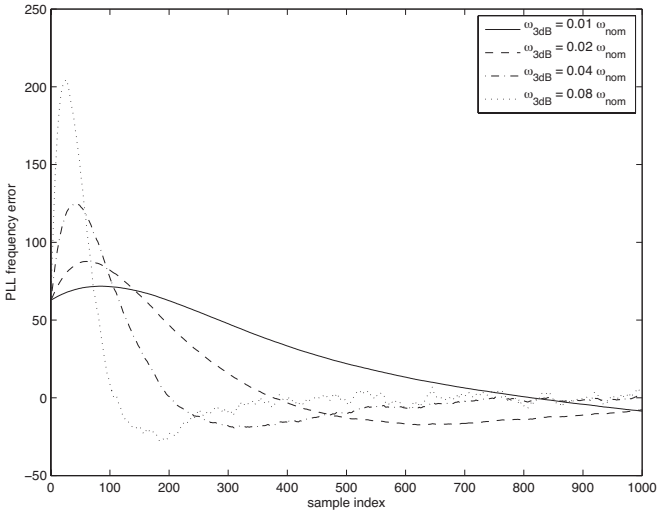

Fig. 2. An example of PLL convergence as a function of the loop filter bandwidth $\omega_{3\mathrm{dB}}$. When $\omega_{3\mathrm{dB}} = 0.01\omega_{\mathrm{nom}}$ or $\omega_{3\mathrm{dB}} = 0.02\omega_{\mathrm{nom}}$, the PLL suppresses the noise but converges too slowly. When $\omega_{3\mathrm{dB}} = 0.08\omega_{\mathrm{nom}}$, the PLL converges quickly but lets through excessive noise.

V. Zero-Crossing Phase and Frequency Estimation

The zero-crossing phase and frequency estimator is based on the observation that the phase of the signal in (1) is a first-order polynomial in t, i.e. $z(t) = b\exp(j\phi(t)) + w(t)$ with $\phi(t) = \omega t + \theta$. The slope and intercept of this line correspond to the radian frequency and phase of the observed signal. If $L \geq 2$ points on this line can be measured, i.e. $\{(t_1, \phi_1), \ldots, (t_L, \phi_L)\}$, then the slope and intercept of the line can be determined exactly. In the presence of noise, a linear regression can be performed to find the least-squares fit for the slope and intercept of the line.

The zero-crossing phase and frequency estimator generates the time/phase coordinate set $\{(t_1, \phi_1), \ldots, (t_L, \phi_L)\}$ by detecting zero crossings in the real and imaginary components of (2) according to the state machine shown in Figure 3. The hysteresis parameter $\alpha \geq 0$ sets the threshold at which zero crossings are detected.

To illustrate the principle of operation, consider the example shown in Figure 4 with hysteresis parameter $\alpha = 0.5$. Both the real and imaginary parts of the discrete-time observation in (2) are plotted. Looking first at the real part of the signal, we see that the signal begins in state 2 and transitions to state 4 when $n = 2$. The signal remains in either state 2 or state 4 until $n = 13$, at which point the $4 \rightarrow 1$ state transition triggers the detection of positive-slope zero crossing. The time of this first

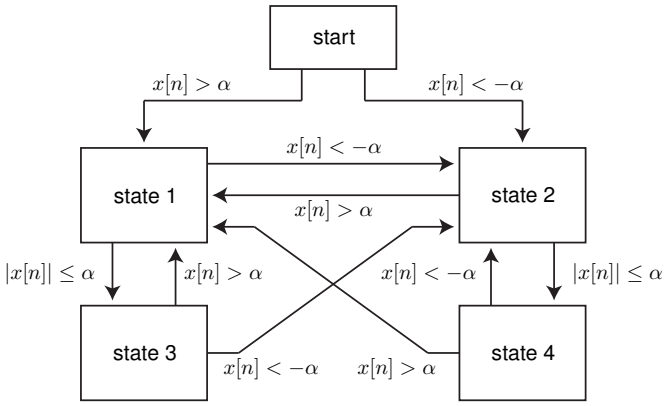

Fig. 3. State machine implementation of the zero crossing detector. The variable $x[n]$ denotes either the real or imaginary part of $z[n]$. A negative-slope zero crossing is detected on state transitions $1 \rightarrow 2$ and $3 \rightarrow 2$. A positive-slope zero crossing is detected on state transitions $2 \rightarrow 1$ and $4 \rightarrow 1$.

zero crossing is estimated using a simple linear interpolator between the last sample in state 2, i.e. $n = 11$, and the first sample in state 1, i.e. $n = 13$. To determine the phase of the signal at the time of this zero crossing, we note that positive-slope zero crossings of the real part of the signal must occur at phase $k2\pi - \pi/2$ for integer k. Similarly, negative-slope zero crossings of the real part of the signal must occur at phase $k2\pi + \pi/2$ for integer k. Hence, the phase of this first positive-slope zero crossing is set to $-\pi/2$.

The next zero crossing of the real part of the signal is detected in the $1 \rightarrow 2$ state transition between samples $n = 15$ and $n = 16$. The time of this negative-slope zero crossing is again estimated using a simple linear interpolator and the phase of this zero crossing is set to $\pi/2$. Two more zero crossings in the real part of the signal are evident in this example, resulting in the coordinate set for the real part of the signal of $\mathcal{S}_R = \{(t_{R1}, -\pi/2), (t_{R2}, \pi/2), (t_{R3}, 3\pi/2), (t_{R4}, 5\pi/2)\}$.

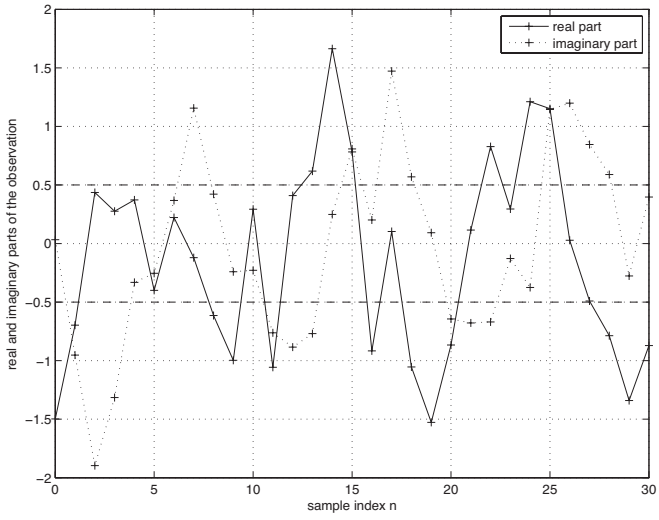

Fig. 4. Real and imaginary parts of a noisy discrete-time observation used to illustrate the zero-crossing phase and frequency estimator with hysteresis parameter $\alpha = 0.5$.

The same process is applied to the imaginary part of the signal. We note that positive-slope zero crossings of the imaginary part of the signal must occur at phase $k2\pi$ for integer k. Similarly, negative-slope zero crossings of the imaginary part of the signal must occur at phase $k2\pi + \pi$ for integer k. In the example shown in Figure 4, a total of five zero crossings are detected in the imaginary part of the signal. Since the first detected zero crossing has positive slope, the coordinate set for the imaginary part of the signal is $\mathcal{S}_I(\ell) = \{(t_{I1}, \ell 2\pi), (t_{I2}, \ell 2\pi + \pi), (t_{I3}, \ell 2\pi + 2\pi), (t_{I4}, \ell 2\pi + 3\pi), (t_{I5}, \ell 2\pi + 4\pi)\}$. Note that this coordinate set is paremeterized by the integer parameter ℓ to facilitate alignment with the coordinate set obtained from the real part of the signal, as described below.

To compute the phase and frequency estimates, we form the union of the coordinate sets

$$\mathcal{S}(\ell) = \mathcal{S}_R \cup \mathcal{S}_I(\ell) \tag{11}$$

and perform a linear regression [14] on $\mathcal{S}(\ell)$ to determine the least-squares slope and intercept of the phase line. The slope and intercept are set to $\hat{\omega}(\ell)$ and $\hat{\theta}(\ell)$, respectively. The total squared error of the linear regression is then

$$\mathcal{E}(\ell) = \sum_{n=0}^{N-1} \left| b \exp\left\{ j \left(\hat{\omega}(\ell) nT + \hat{\theta}(\ell) \right) \right\} - z[n] \right|^2. \tag{12}$$

Note that (11) and (12) should be computed at least over $\ell \in \{-1, 0, 1\}$ to find the value of ℓ such that $\mathcal{E}(\ell)$ is minimized. In other words, the integer parameter ℓ should be selected to best align the independently obtained phases in the real and imaginary coordinate sets. It is often the case that the total squared error is minimized when $\ell = 0$, i.e. the coordinate sets are often naturally aligned. In the example shown in Figure 4, however, $\mathcal{E}(\ell)$ is minimized when $\ell = -1$.

A. Real-Time Sample-by-Sample Implementation

Unlike the FFT-based MLE which requires all N samples of the observation to be present before processing can begin, the zero-crossing phase and frequency estimator can be implemented on a sample-by-sample basis. Prior to the first sample of the observation, we initialize the tracking variables $A = B = C = D = E = 0$. Samples are taken until a zero crossing is detected in the real or imaginary part of the observation. Upon the detection of a zero crossing, a new coordinate (t_i, ϕ_i) is added to the coordinate set $\mathcal{S}$ and the tracking variables are updated as follows:

$$\begin{aligned} A &= A + t_i^2, \\ B &= B + 1, \\ C &= C + t_i, \\ D &= D - t_i \phi_i, \text{ and} \\ E &= E - \phi_i. \end{aligned}$$

Note that the D and E tracking variables may need to be indexed and updated according to the alignment parameter ℓ in order to select the real/imaginary phase alignment that

minimizes the total squared error as discussed above. In any case, updating the tracking variables requires only a modest fixed amount of processing per detected zero crossing. As shown in [14], the least-squares phase and frequency estimates can be calculated from the tracking variables as

$$\hat{\theta} = -\frac{EA - CD}{BA - C^2}$$
$$\hat{\omega} = -\frac{C}{A}\hat{\theta} - \frac{D}{A}.$$

Hence, the least-squares phase and frequency estimates of the zero-crossing estimator can be computed at any time with two reciprocal operations and eight multiplications.

In practice, it is often the case that the unknown frequency ω will be known within a certain range. In these types of scenarios, the zero-crossing phase and frequency estimator can pre-filter the observation to reduce the overall noise variance while passing the desired frequencies. While adding some additional computational complexity, the filtering can also be performed also on a sample-by-sample basis and the order of the filter can be selected to maintain real-time operation. Moreover, the phase offset caused by pre-filtering is known and can be corrected at the output of the zero-crossing phase and frequency estimator.

VI. Numerical Results

This section presents numerical results comparing the FFT-based MLE, PLL, and zero-crossing phase and frequency estimators for different values of M and $\omega_{3\mathrm{dB}}$. All of the results in this section assume an observation with $N = 1000$ samples at $f_s = 16$ kHz. 10000 realizations of the complex exponential signal and AWGN were generated with fixed $b = 1$ and random independent uniformly distributed phase and frequency centered at 1020 Hz according to

$$\phi \sim \mathcal{U}(-\pi, \pi) \text{ and}$$
$$\omega \sim \mathcal{U}\left(2\pi \cdot 1010, 2\pi \cdot 1030\right).$$

The independent AWGN in each realization was generated with independent real and imaginary components, each zero mean, white, and with variance σ^2.

The FFT-based MLE, PLL, and zero-crossing phase and frequency estimation techniques are compared in terms of their mean squared estimation error. In order to understand the relative performance of each estimator with respect to the theoretically achievable performance, we also compute the Cramer-Rao bound (CRB) [15]. The CRB for the covariance of the frequency and phase estimates of a complex exponential in AWGN when both the phase and frequency are unknown is given as [12]

$$\mathrm{cov}\left\{\left[\tilde{\omega}, \tilde{\theta}\right]^\top\right\} \geq \frac{\sigma^2}{b^2}\begin{bmatrix} \frac{1}{T^2 N(R-S^2)} & \frac{-(n_0+R)}{TN(R-S^2)} \\ \frac{-(n_0+R)}{TN(R-S^2)} & \frac{n_0^2+2n_0R+S}{N(R-S^2)} \end{bmatrix} \quad (13)$$

where $n_0 := t_0 f_s$ is the index of the first sample, $R := (N-1)/2$, $S := (N-1)(2N-1)/6$, and the notation $\boldsymbol{A} \geq \boldsymbol{B}$ means that $\boldsymbol{A} - \boldsymbol{B}$ is positive semidefinite.

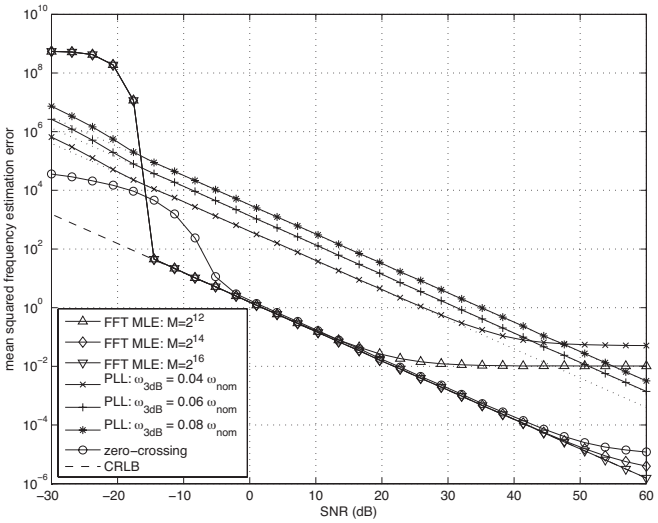

Fig. 5. Mean squared frequency estimation error as a function of SNR.

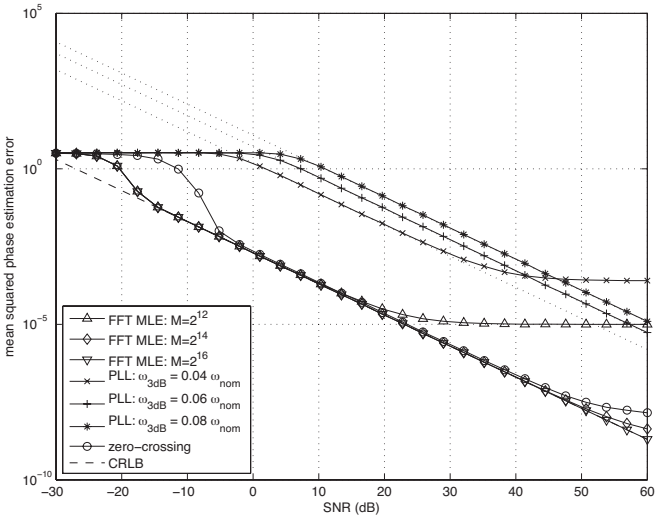

Fig. 6. Mean squared phase estimation error as a function of SNR.

Figures 5 and 6 show the mean squared frequency and phase estimation error, respectively, of the FFT-based MLE, PLL, and zero-crossing phase and frequency estimators as a function of $\mathsf{SNR} := 10\log_{10}(b^2/(2\sigma^2))$ for three values of M and three values of $\omega_{3\mathrm{dB}}$. The FFT-based MLE in these examples is implemented with quadratic interpolation as discussed in Section III. The PLL phase and frequency estimator in these examples uses a phase detector with $K_d = 1$, a VCO with $K_o = 10$ and $\omega_{\mathrm{nom}} = \omega$ (no frequency offset), and a second-order active PI loop filter [13] designed with several different values of $\omega_{3\mathrm{dB}}$ and converted to discrete-time using the bilinear transform. The zero-crossing phase and frequency estimator in these examples uses a hysteresis parameter $\alpha = 0.1$ and a 64^{th} order FIR bandpass input pre-filter with first stopband below 800 Hz, second stopband above 1250 Hz, and passband between 920 Hz and 1120 Hz designed using MATLAB's `firls` filter design function.

In these examples, the PLLs tend to provide the worst estimation performance of all of the techniques considered except with respect to the FFT-based MLE at very low SNR. To understand why this is the case, recall that the PLL frequency estimate is obtained as $\hat{\omega} = \omega_0[N-1] = \omega_{\rm nom} + K_0 u[N-1]$ where $u[N-1]$ is the final output of the loop filter in Figure 1. In the locked state with a white Gaussian phase input, (10) can be used to predict the variance of the frequency estimate as

$$\mathrm{var}(\hat{\omega}) = \mathrm{var}(\phi_{in}) \cdot \frac{1}{2\pi} \int_{-\pi}^{\pi} |H_\omega(e^{j\beta})|^2 \, d\beta.$$

The variance of the PLL phase estimate in the locked state can be predicted similarly. The dotted lines in Figures 5 and 6 show the PLL estimation performance prediction based on the linear locked PLL model in (9) and (10). The PLL phase and frequency estimation performance closely tracks the analytical predictions except at very low SNR (where the small angle approximation in the linear PLL model does not hold) and at very high SNR for the PLL with loop filter bandwidth $\omega_{\rm 3dB} = 0.04\omega_{\rm nom}$. The flattening of the MSE curve for the PLL with $\omega_{\rm 3dB} = 0.04\omega_{\rm nom}$ at SNR ≥ 40 dB is the result of slow PLL convergence. In this case, the the transient effects of the PLL dominate the effects of the noise on the phase and frequency estimates. Overall, as long as the PLL loop filter bandwidth is sufficiently wide to allow PLL convergence, narrower loop filter bandwidths tend to give better performance since the variances of the phase and frequency estimates are increasing functions of the loop filter bandwidth.

The FFT-based MLE phase and frequency estimators with quadratic interpolation tend to perform well when SNR ≥ -15 dB. At very low SNR, the FFT-based MLE tends to perform poorly because the desired signal does not emerge from the noise in the FFT. When $M = 2^{16}$ and $-15 \leq$ SNR ≤ 60 dB, the FFT-based MLE phase and frequency estimator closely tracks the CRLB. The effects of limited frequency resolution in the FFT-based MLE phase and frequency estimators can be seen in the flattening of the phase and frequency mean squared estimation error curves for $M = 2^{12}$ and $M = 2^{14}$ at higher SNR. Performance in this regime could be improved with more sophisticated interpolation techniques, at the expense of additional computational complexity.

The zero-crossing phase and frequency estimator tends to perform well when SNR ≥ -3 dB, closely tracking the CRLB for SNR up to 45 dB. High SNR performance can be further improved by decreasing the hysteresis parameter α, increasing the sampling rate, and/or using more sophisticated interpolation techniques to estimate the zero crossing times. The zero-crossing estimator performs better than all of the FFT-based MLE estimators at very low SNR and offers equivalent or better performance than all of the FFT-based MLE estimators for $0 \leq$ SNR ≤ 40 dB. The advantage of the zero-crossing estimator, however, is that it is able to perform the necessary computations on a sample-by-sample basis and requires only a small amount of additional computation at the end of the observation to produce the necessary phase and frequency estimates.

VII. Conclusion

This paper presented a performance comparison among three different phase and frequency estimation techniques in the context of carrier synchronization for distributed transmit beamforming. We proposed a low-complexity real-time single-tone phase and frequency estimation technique based on zero-crossing detection and linear regression. The proposed technique is similar to the linear regression estimation techniques described in [10] and [11] except that it avoids the computational burden of inverse tangents and phase unwrapping by only using information contained in the zero-crossings of the observed signal. The proposed zero-crossing phase and frequency estimator offers near-optimal performance over a wide range of signal-to-noise ratios as well as low-complexity sample-by-sample operation appropriate for real-time applications. Numerical results were provided demonstrating that the proposed technique closely tracks the Cramer-Rao lower bound over a wide range of signal to noise ratios.

References

[1] Y. Tu and G. Pottie, "Coherent cooperative transmission from multiple adjacent antennas to a distant stationary antenna through AWGN channels," in *IEEE Vehicular Technology Conf. (VTC)*, vol. 1, Birmingham, AL, Spring 2002, pp. 130–134.

[2] R. Mudumbai, J. Hespanha, U. Madhow, and G. Barriac, "Scalable feedback control for distributed beamforming in sensor networks," in *IEEE International Symp. on Information Theory (ISIT)*, Adelaide, Australia, September 2005, pp. 137–141.

[3] R. Mudumbai, B. Wild, U. Madhow, and K. Ramchandran, "Distributed beamforming using 1 bit feedback: from concept to realization," in *44th Allerton Conf. on Comm., Control, and Computing*, Monticello, IL, Sep. 2006, pp. 1020 – 1027.

[4] R. Mudumbai, J. Hespanha, U. Madhow, and G. Barriac, "Distributed transmit beamforming using feedback control," *IEEE Trans. on Information Theory*, in review.

[5] R. Mudumbai, G. Barriac, and U. Madhow, "On the feasibility of distributed beamforming in wireless networks," *IEEE Trans. on Wireless Communications*, vol. 6, no. 5, pp. 1754–1763, May 2007.

[6] D.R. Brown III, G. Prince, and J. McNeill, "A method for carrier frequency and phase synch. of two autonomous cooperative transmitters," in *IEEE Signal Proc. Advances in Wireless Comm. (SPAWC)*, New York, NY, June 5-8, 2005, pp. 278–282.

[7] D.R. Brown III and H.V. Poor, "Time-slotted round-trip carrier synchronization for distributed beamforming," *IEEE Trans. on Signal Processing*, vol. 56, no. 11, pp. 5630–5643, November 2008.

[8] R. D. Preuss and D.R. Brown III, "Retrodirective distributed transmit beamforming with two-way source synchronization," in *44th Annual Conference on Information Sciences and Systems (CISS 2010)*, March 2010, p. tbd.

[9] R. Mudumbai, D.R. Brown III, U. Madhow, and H.V. Poor, "Distributed transmit beamforming: Challenges and recent progress," *IEEE Communications Magazine*, vol. 47, no. 2, pp. 102–110, February 2009.

[10] S. A. Tretter, "Estimating the frequency of a noisy sinusoid by linear regression," *IEEE Transactions on Information Theory*, vol. 31, no. 6, pp. 832–835, November 1985.

[11] S. Kay, "A fast and accurate single frequency estimator," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 37, no. 12, pp. 1987–1990, December 1989.

[12] D. Rife and R. Boorstyn, "Single-tone parameter estimation from discrete-time observations," *IEEE Trans. on Information Theory*, vol. 20, no. 5, pp. 591–598, September 1974.

[13] R. Best, *Phase-Locked Loops : Design, Simulation, and Applications*. New York: McGraw-Hill, 2003.

[14] G. Allaire and S. M. Kaber, *Numerical Linear Algebra*. Springer Science + Business Media, LLC, 2008.

[15] H.V. Poor, *An Introduction to Signal Detection and Estimation*, 2nd ed. New York: Springer-Verlag, 1994.