

Improved Carrier Tracking for Low-Threshold Telemetry Using a Smoother

W. J. Hurd¹ and A. Mileant²

This work establishes the threshold performance of noncausal smoothing filters used for tracking residual carrier signals when performance is limited by both phase process noise and additive receiver noise. Previous work [1] based entirely on linear theory predicts significant improvements in phase estimation by using smoothing filters as compared with using causal phase-locked loops (PLLs). For phase process noise having a $1/f^3$ spectral density, which is typical of many oscillators, linear theory predicts a reduction in phase estimation error of 5 dB for the same phase-noise and receiver-noise magnitudes. Alternatively, linear theory predicts the ability of a smoother to track a 7.5-dB weaker signal with the same mean-square phase error, under the same conditions of phase process noise. Simulations show that the smoothers achieve an effective loop signal-to-noise ratio (SNR) within 0.5 dB of linear theory when the linear theory predicts an effective loop SNR as low as 11.4 dB. Under the same conditions, an optimum PLL would have a linear-model effective loop SNR of only 6.4 dB and would perform more than 1-dB worse than this, with many cycle slips.

I. Introduction

The general problem of estimating the phase of a carrier signal in noise, when the carrier phase is a random process, is highly nonlinear. When the conditions are such that the estimation error is fairly small, the performance can be well predicted by a linear model. For example, linear models for phase-locked loops (PLLs) are reasonably accurate when the mean-square phase error is less than approximately 0.1 rad^2 . The linear model is not adequate in the threshold region when cycle slipping occurs. To obtain a simple model to predict the performance of noncausal filters and to find optimum filters, previous work [1] has treated the problem as linear. The current work determines the valid range of the linear model and the useful range of the smoothing filters, i.e., the range above threshold.

The linear analysis models the received carrier signal as a random phase process plus additive white receiver noise. From Wiener filtering theory, the optimum linear estimator of the phase process is the noncausal filter, or smoother, with transfer function

¹ TMOD Systems Engineering and Standards Office.

² Communications Systems and Research Section.

The research described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

$$H_0(j2\pi f) = \frac{S_s(f)}{S_s(f) + S_n(f)} \quad (1)$$

where $S_s(f)$ and $S_n(f)$ are the two-sided spectral densities of the signal phase process being estimated and of the additive random-noise process, respectively. Since $S_s(f)$ and $S_n(f)$ are real-valued and even-symmetric about zero frequency, so is H_0 , and so is the corresponding weighting function, $h_0(t)$. Thus, the optimum estimator is a symmetric smoothing filter with an infinite delay.

In specifying the phase stability of oscillators, the spectral density of the phase noise typically is given in terms of $L(f)$, which is the phase-noise density f Hz from the oscillator nominal frequency, with f being either positive or negative. The natural units of $L(f)$ are rad^2/Hz , or, in decibels, $\text{dB-rad}^2/\text{Hz}$. When $L(f)$ is measured by spectral analysis, it often is expressed in decibels per Hertz relative to the total oscillator power, or dBc/Hz ; this is numerically equivalent to dBrad^2/Hz . In this work, we model the phase process as a low-pass process that phase modulates an ideal sinusoid. Since half of the power in the low-pass modulation process goes into each side of the oscillator nominal frequency, the one-sided spectral density of the low-pass phase process is $S_\theta(f) = 2L(f)$ rad^2/Hz . Expressing the additive noise in the same units requires normalization by the signal power, so the noise spectral density is N_0/P_c rad^2/Hz , where P_c is the oscillator power and N_0 is the spectral density of the original additive noise process.

The case considered in this and the previous work is that when the signal is a phase process, $\theta(t)$, with one-sided spectral density

$$S_\theta(f) = 2L(1)f^{-\alpha} \quad (2)$$

The simulations were conducted for $\alpha = 3$, because this is typical of oscillator noise, but the analytical results are presented for general $\alpha > 1$. The optimum filter has transfer function [1]

$$H_0(j2\pi f) = \frac{1}{1 + \frac{N_0|f|^\alpha}{2P_cL(1)}} \quad (3)$$

The optimum smoother is not realizable because it has infinite extent in time. One way to approximate it is to assume a filter of a fixed shape and to optimize only the bandwidth of the filter. This was done for several filter shapes. It was shown that a filter of any shape can be characterized by a parameter b_0 , which is defined as the optimum bandwidth for that shape filter when $N_0 = 2P_cL(1)$.

Then, for a filter with parameter b_o , the optimum bandwidth as a function of the phase and noise spectral densities is [1]

$$b_{opt} = \left[\frac{2P_cL(1)}{N_0} \right]^{1/\alpha} b_0 \quad (4)$$

and the minimum mean-square phase-estimation error is

$$\sigma_{opt}^2 = \left(\frac{\alpha}{1-\alpha} \right) \left(\frac{N_0}{P_c} \right) b_{opt} \quad (5)$$

Besides the oscillator noise case with $\alpha = 3$, the case of $\alpha = 8/3$ is of particular interest because it is characteristic of phase noise caused by propagation through charged particles, such as when a signal

passes close to the Sun. Although this case was not simulated, the results would be expected to be similar to the simulated cases, because 8/3 is close to 3. Tracking results for Pioneer 10 using optimum phase-locked loops are given elsewhere [2].

The minimum mean-square error also can be expressed in terms of $L(1)$, P_c/N_o , and the filter-shape bandwidth parameter, b_0 , as

$$\sigma_{opt}^2 = \left(\frac{\alpha}{1 - \alpha} \right) [2L(1)]^{1/\alpha} \left(\frac{P_c}{N_0} \right)^{(1-\alpha)/\alpha} b_0 \quad (6)$$

This is the minimum mean-square error that can be achieved by optimizing the bandwidth of a filter with shape characterized by b_0 , which in turn depends on the filter shape and α .

Examining Eq. (6), it is interesting that when the filter shape is fixed but the bandwidth is optimized, the mean-square error increases only as the $1/\alpha$ power of the phase-noise density and as the $(1 - \alpha)/\alpha$ power of the receiver-noise density. In particular, it increases less rapidly with receiver noise than when the bandwidth is fixed.

Table 1 presents the values of b_0 for some filters and smoothers for α of 3 and 8/3. Some of these values were presented incorrectly in [1] and are corrected here. The filter with a triangular weighting function has been added. This triangular filter results in a mean-square error only 1.01 times, or 0.04-dB, larger than the optimum linear smoother. It is this filter that is simulated throughout this work.

II. Simulation Objectives, Problems, and Approach

The primary objectives of the simulation were to establish the smoother performance in the linear region of operation and to determine where it approaches threshold. Because threshold effects are dependent on the details of implementation of the entire phase estimation and demodulation process, another

Table 1. Optimum bandwidths for smoothers and causal filters for $N_0 = 2L(1) = 1$.

Case	Description	b_0 for $\alpha = 3$	b_0 for $\alpha = 8/3$
Smoothers			
1	Optimum smoother	0.80613	0.797
2	Rectangular passband	1.0	1.0
3	Rectangular weighting function	0.900	0.895
4	Symmetric exponential weighting	0.851	0.817
5	Triangular weighting function	0.816	0.805
Causal filters			
6	Rectangular weighting function	Infinite	2.60
7	First-order (exponential weighting)	Infinite	2.47
8	Second-order PLL, critically damped	2.489	2.234
9	Third-order PLL, critically damped	2.752	2.490
10	Third-order PLL [2, Eqs. (6)–(8), with $r = 4, k = 0.25$]	2.629	2.354

objective of the simulation was to test different configurations of the overall smoothing system in order to isolate the best among them.

The triangular-shaped weighting-function approximation to the optimum smoother was used because of its finite extent and ease of implementation. The equations of the linear smoother do not define a complete implementable solution to the noncausal phase-tracking problem because of the sinusoidal phase nonlinearities. However, intuitively, one would expect the linear model to be accurate when the values of the phase being estimated are “small” within the smoother’s window. For example, for an observation window extending from $-T/2$ to $T/2$ s, one would expect the linear model to be reasonably accurate when the magnitude of the phase input to the smoother is less than about 0.3 rad everywhere within this region. The simulations support this rationale. Thus, one of the main problems in implementing the smoother is to keep the phase small throughout the estimation interval.

To understand the rationale for the method of reducing the phase error at the smoother input, note that the smoother is not sensitive to linear phase changes. In other words, because the optimum linear smoother is indeed linear, the same result is obtained if a linear change in the input is made, then the smoother is applied, and finally the change to the input is removed at the output. Thus, the optimality of the smoother is not changed if the input signal is shifted by a frequency that is constant over the time extent of the smoother.

Following the above rationale, the general approach to reducing phase error at the smoother input is to translate the input phase using the predicted phase and the estimated frequency at the center of each smoother interval. Several methods for estimating frequency and phase were tried, as discussed in Section III. The method that worked best and was used for the final simulations was to track the input signal with a phase-locked loop. The average numerically controlled oscillator (NCO) frequency and the previous smoother phase estimate are used to estimate the phase for the next smoother interval. A cycle-slip detector is used to detect cycle slips in the PLL. When these occur, the NCO values during the slips are ignored so that they will not corrupt the estimates of the PLL’s average frequency. More on this is mentioned in the next section.

III. System Description

Figure 1 depicts the main components of the smoother simulation. The PLL tracks the input signal using a type II PLL with the optimum bandwidth for the conditions of signal-to-noise ratio (SNR) and phase-noise density. To minimize the phase error inside the smoother observation windows, linear curve fitting is performed by modeling the instantaneous phase process in the interval $-T/2 \leq t \leq T/2$ s as being of the form $\theta(t) = \theta_0 + \hat{\theta}t$. The PLL frequency is averaged over the smoothing interval to obtain an estimate, $\hat{\theta}$, of the average NCO frequency. An estimate of phase, $\hat{\theta}_0$, is obtained from the previous smoother output, extrapolated to the center of the present smoothing interval using $\hat{\theta}$. Thus, a linear phase prediction is obtained for the current smoothing phase interval expressed by Eq. (A-11) of the Appendix in the discrete implementation of this algorithm.

The simulations have shown that the performance of the noncausal receiver was very dependent on the accuracy of the $\hat{\theta}$ and $\hat{\theta}_0$ estimates. To quantify the impact of these estimates, a case was run wherein the estimates were derived from linear fits to the actual input phase process (which is possible only in a simulation). As expected, this gave the best performance, very close to the theoretical predictions.

A total of four frequency-tracking methods were simulated, including the PLL. These were (1) a frequency-locked loop (FLL); (2) a fast Fourier transform (FFT)-based estimator; (3) differencing of the last two smoother phase estimates; and (4) a PLL-based estimator. In our simulations, the frequency estimation was performed in parallel with the smoothing process so that it did not distort the spectrum of the signal samples fed into the smoother array. In an actual situation, any shifts in phase performed

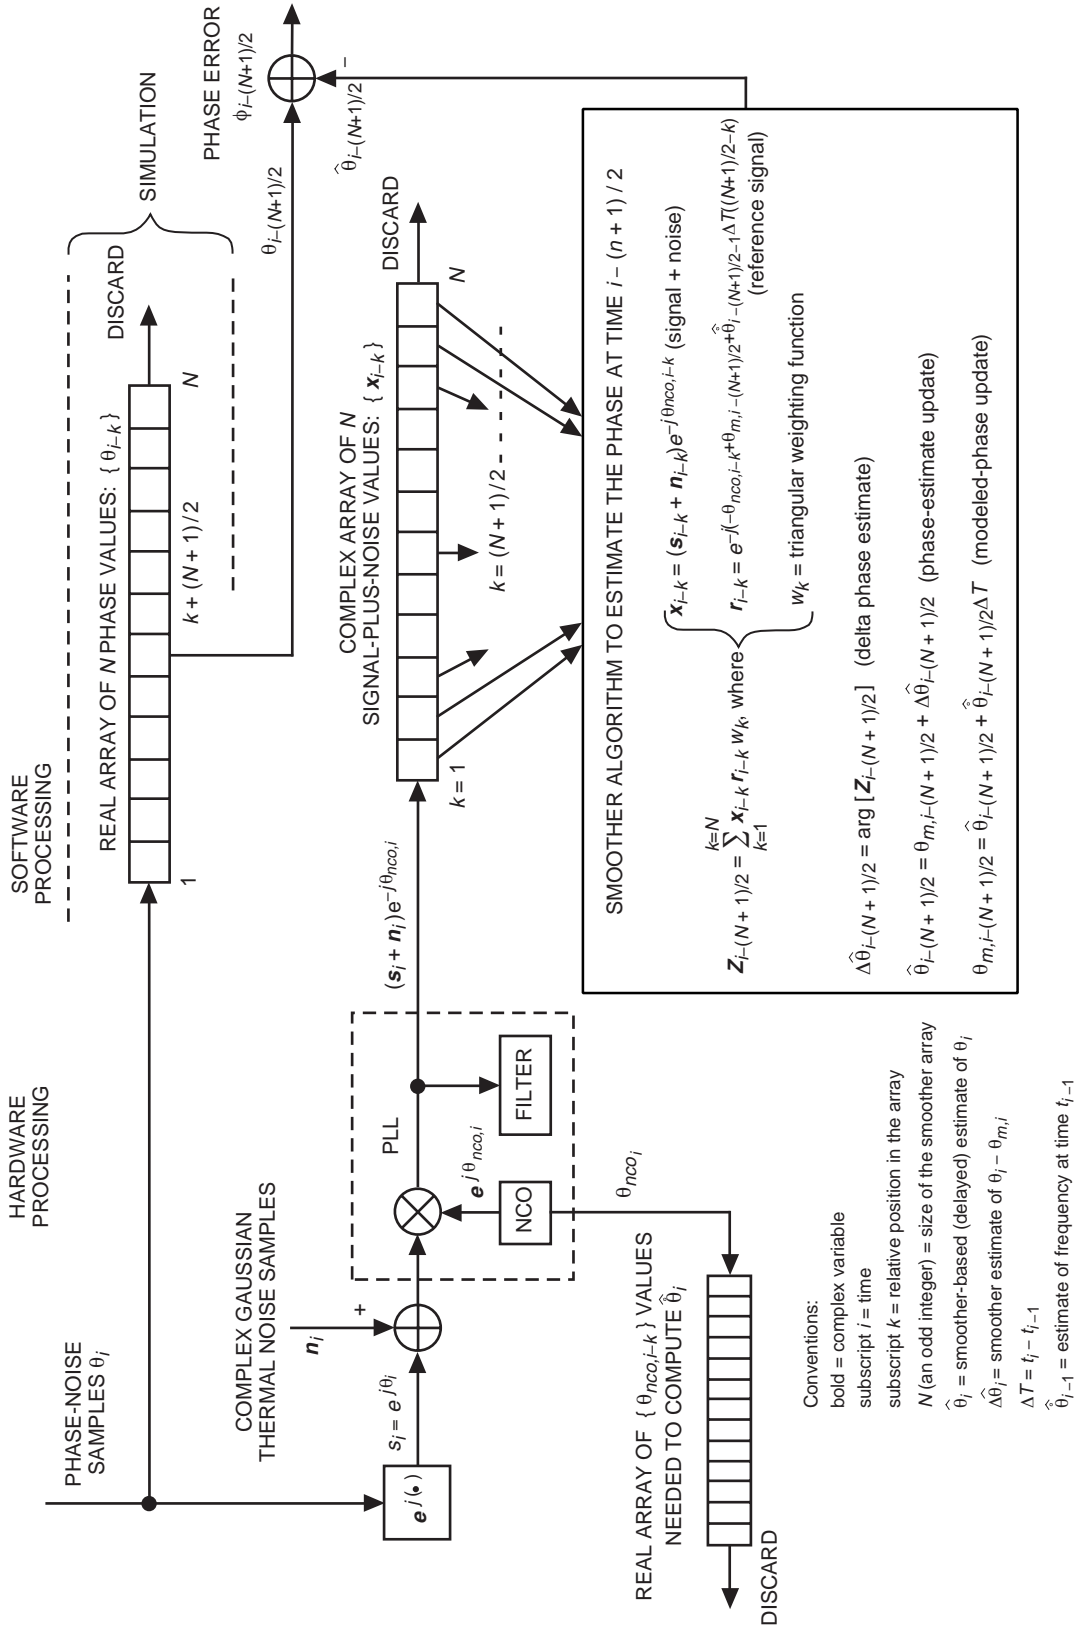


Fig. 1. Block diagram of a noncausal receiver with a PLL preceding it.

before the smoother, say by a PLL, have to be canceled in the smoother algorithm in order to preserve the spectrum of the tracked signal. See Fig. 1 and Eq. (A-7) of the Appendix. Details of the algorithms used in the simulation are summarized in the Appendix.

Applying the user-defined values of P_c/N_0 and of the phase-noise spectrum, the simulating program sets the optimum bandwidths of the smoother and of the other tracking devices. At the end of the run, the program computes the statistics of the phase error and of the cycle slips in different loops and compares these with the theoretical predictions.

The best results in terms of the smoother's minimum phase-error variance were obtained when the smoother was aided by PLL-based frequency estimates. The limiting factor in this method (thresholding) comes from cycle slips that the PLL experiences at low SNR values. Because of noncausality of the smoother, some of these cycle slips can be detected and accounted for when computing the average frequency. In the simulation, an algorithm was used to bypass the PLL's cycle-slip regions, which typically lasted about 20 NCO phase samples out of about 400 or more samples in the observation window. This strategy significantly improved the performance of the smoother in the region of cycle slipping. There was no difficulty in handling two or more consecutive cycle slips. It was observed that, at low SNR values, the smoother also experiences "cycle slips" when its phase estimates become off by 2π rad relative to the input phase process. Assessment of the impact of these cycle slips on the demodulated data was outside the scope of our simulations.

IV. Simulation Results

Because the PLL frequency-estimation method yielded the best results, detailed results are presented only for this case. Then the results for other cases are discussed briefly. These other cases were not studied sufficiently that the results should be considered to be conclusive. In other words, there is still a possibility that some implementation using one of these methods may be better than the PLL method.

A. Results for Frequency Estimation by PLL

Figures 2(a) and 2(b) show the simulation results for the system using the PLL to estimate the frequency for the smoother interval; they present the results for this system for phase process noise spectral densities at 1 Hz from the carrier of $L(1) = 10^{-4}, 10^{-3}$, and 10^{-2} rad²/Hz. The simulation results also are presented in Tables 2(a) through 2(f).

The figures show the effective ρ (effective loop SNR) of the causal (PLL) and noncausal receivers as functions of P_c/N_0 . The effective ρ is defined as the inverse of the phase-error variance. This is consistent with the usual definition of loop SNR for PLLs operating in the linear region and with negligible phase process noise.

Each simulation point is based on 6×10^6 samples, which is equivalent to 16 hours of continuous tracking. Triangular windowing was used in all cases. No reinitialization of loop parameters was done after cycle slips, so that the degradation values reported here should correspond to real-life statistics. Each phase error was taken modulo 2π , in the region $\pm\pi$, in calculation of the phase-error statistics. Thus, the results are appropriate for studying the loss in telemetry SNR or radio loss but are not applicable to Doppler navigation.

1. Results for Middle Value of Phase Noise. Figures 2(a) and 2(b) and Tables 2(c) and 2(d) show the results for $L(1) = 10^{-3}$ rad²/Hz. In the figures, the lower and upper straight lines without dots are the theoretical effective ρ for the optimum bandwidths of PLLs and smoothers, respectively.

Details for one SNR: First, consider the results for $P_c/N_0 = 10$ dB-Hz. At this point, the linear model for the PLL yields an effective ρ of 9.94 dB. Since 10 dB is the minimum effective ρ recommended for use

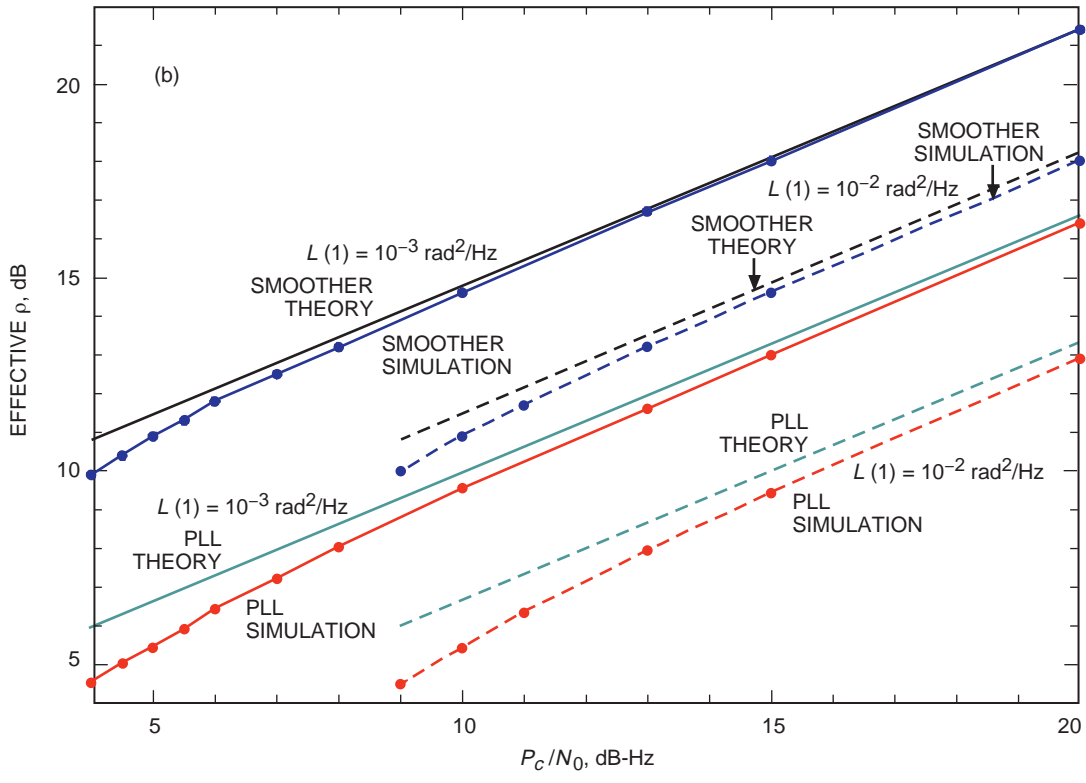
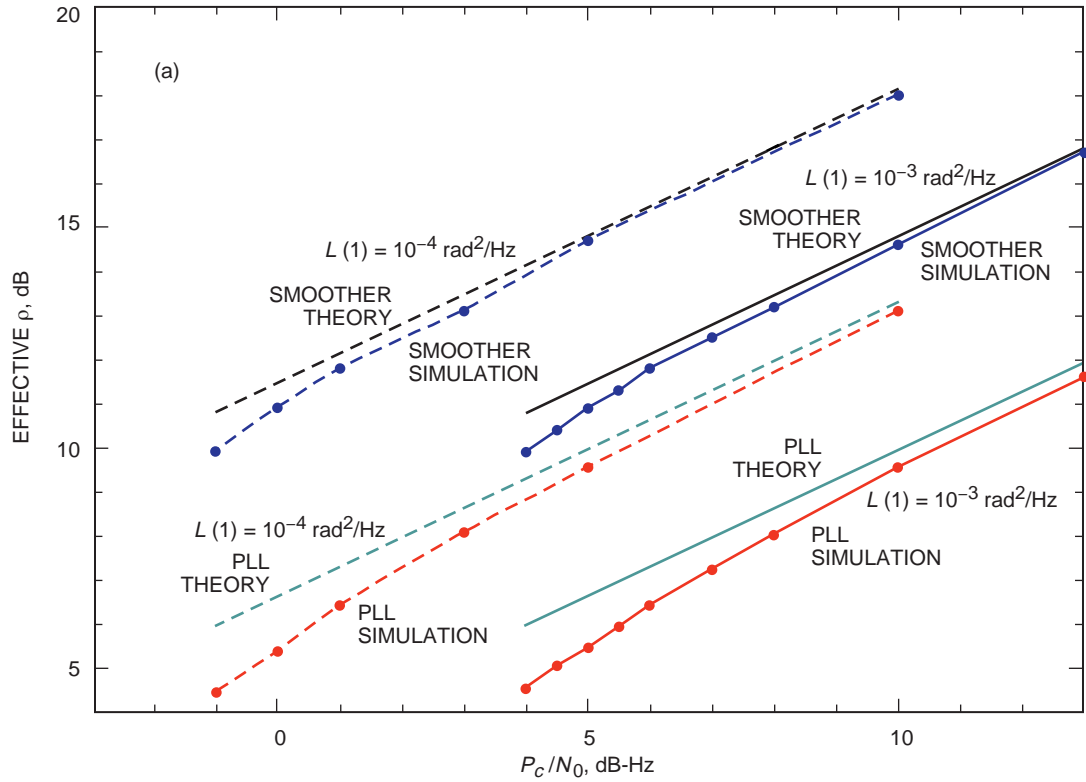


Fig. 2. Effective ρ versus P_c/N_0 for causal and noncausal tracking using optimum bandwidths, using a type II digital phase-locked loop (DPLL) for causal tracking and frequency estimates, $\alpha = 3$: (a) $L(1) = 10^{-3}$ and $10^{-4} \text{ rad}^2/\text{Hz}$ and (b) $L(1) = 10^{-3}$ and $10^{-2} \text{ rad}^2/\text{Hz}$.

Table 2(a). Performance of a noncausal filter for $L(1) = 1 \times 10^{-4} \text{ rad}^2/\text{Hz}$, $\alpha = 3$, aided by a PLL running in parallel.

P_c/N_0 , dB-Hz	N samples	$B_{L \text{ opt,}}$ Hz	ρ_{eff} predicted, dB	ρ_{eff} simulated, dB	$\Delta\rho$, dB
10.0	647	0.1030	18.10	18.00	0.10
5.00	951	0.0701	14.80	14.60	0.20
3.00	1111	0.0600	13.40	13.10	0.30
1.00	1293	0.0516	12.10	11.75	0.35
0.00	1397	0.0477	11.40	10.90	0.50
-1.00	1509	0.0442	10.80	9.93	0.87

Table 2(b). Performance of a causal filter for $L(1) = 1 \times 10^{-4} \text{ rad}^2/\text{Hz}$. $\alpha = 3$.

P_c/N_0 , dB-Hz	$B_{L \text{ opt,}}$ Hz	ρ_{eff} predicted, dB	ρ_{eff} simulated, dB	$\Delta\rho$, dB
10.0	0.313	13.30	13.11	0.19
5.00	0.213	9.94	9.55	0.39
3.00	0.183	8.61	8.07	0.54
1.00	0.157	7.28	6.43	0.85
0.00	0.145	6.61	5.37	1.24
-1.00	0.135	5.94	4.45	1.49

Table 2(c). Performance of a noncausal filter for $L(1) = 1 \times 10^{-3} \text{ rad}^2/\text{Hz}$, $\alpha = 3$, aided by a PLL running in parallel.

P_c/N_0 , dB-Hz	N samples	$B_{L \text{ opt,}}$ Hz	ρ_{eff} predicted, dB	ρ_{eff} simulated, dB	$\Delta\rho$, dB
20.0	139	0.480	21.40	21.31	0.086
15.0	205	0.325	18.10	18.01	0.09
13.0	239	0.279	16.80	16.69	0.11
10.0	299	0.223	14.80	14.64	0.16
8.00	349	0.191	13.40	13.16	0.24
7.00	377	0.177	12.80	12.54	0.26
6.00	409	0.163	12.10	11.75	0.35
5.00	441	0.151	11.40	10.87	0.53
4.00	477	0.140	10.80	9.93	0.87

Table 2(d). Performance of a causal filter for $L(1) = 1 \times 10^{-3} \text{ rad}^2/\text{Hz}$, $\alpha = 3$.

P_c/N_0 , dB-Hz	$B_{L \text{ opt}}$, Hz	ρ_{eff} predicted, dB	ρ_{eff} simulated, dB	$\Delta\rho$, dB
20.0	1.460	16.60	16.40	0.198
15.0	0.992	13.30	13.05	0.251
13.0	0.851	11.90	11.60	0.301
10.0	0.676	9.94	9.55	0.390
8.00	0.580	8.61	8.02	0.593
7.00	0.537	7.94	7.23	0.707
6.00	0.497	7.28	6.44	0.845
5.00	0.460	6.61	5.46	1.15
4.00	0.426	5.94	4.53	1.41

Table 2(e). Performance of a noncausal filter for $L(1) = 1 \times 10^{-2} \text{ rad}^2/\text{Hz}$, $\alpha = 3$, aided by a PLL running in parallel.

P_c/N_0 , dB-Hz	N samples	$B_{L \text{ opt}}$, Hz	ρ_{eff} predicted, dB	ρ_{eff} simulated, dB	$\Delta\rho$, dB
20.0	65	1.030	18.20	18.00	0.200
15.0	95	0.702	14.80	14.55	0.253
13.0	111	0.601	13.50	13.21	0.288
11.0	131	0.509	12.20	11.74	0.459
10.0	141	0.473	11.50	10.90	0.599
9.00	151	0.442	10.80	9.96	0.836

Table 2(f). Performance of a causal filter for $L(1) = 1 \times 10^{-2} \text{ rad}^2/\text{Hz}$, $\alpha = 3$.

P_c/N_0 , dB-Hz	$B_{L \text{ opt}}$, Hz	ρ_{eff} predicted, dB	ρ_{eff} simulated, dB	$\Delta\rho$, dB
20.0	3.09	13.3	12.86	0.443
15.0	2.11	10.0	9.44	0.557
13.0	1.81	8.67	7.96	0.715
11.0	1.55	7.33	6.34	0.991
10.0	1.44	6.67	5.43	1.24
9.00	1.33	6.00	4.51	1.49

in telemetry systems, this is the approximate threshold for useful telemetry. (Note that the actual loop SNR required to achieve this effective loop SNR is 1.76-dB higher because of the contribution of phase process noise to rms phase error.) The simulation results for the PLL achieve an effective ρ of 9.55 dB, or approximately 0.40-dB lower than the theoretical curve, which is fairly consistent with other PLL results. For the smoother, the effective ρ of the linear model is 14.80 dB, or 4.86-dB higher than for the PLL, according to the ratio of the optimum bandwidth parameters b_0 in Table 1. The simulation results for the smoother achieve an effective ρ of 14.64 dB, which is 5.1-dB better than the PLL.

Significant reduction in radio loss: To illustrate the significance of the smoother performance, let us consider the radio loss, or loss in effective telemetry SNR, for a ($k = 15, r = 1/4$) convolutional code operating at a bit-error rate of 10^{-3} [3, Fig. 19]. Under the above conditions of SNR and phase noise, the PLL is barely operating at the minimum acceptable performance, with a mean-square (MS) phase error of 0.1 rad^2 , and would have a very high radio loss—in excess of 3 dB. The smoother would achieve an MS phase error of approximately 0.03 rad^2 and would have a very acceptable radio loss of approximately 0.2 to 0.3 dB. Thus, the smoother improves the telemetry performance by 2.7 dB!

Smoother threshold: The minimum threshold for the smoother is another factor of interest. Still considering the case of $L(1) = 10^{-3} \text{ rad}^2/\text{Hz}$, we see the smoother simulation achieves an effective ρ of 9.93 dB at a P_c/N_0 of 4.0 dB-Hz. The performance is approximately 0.9-dB below the linear theory. Considering this to be the threshold is consistent with the 10-dB effective ρ that is used for PLL telemetry receivers. It is very interesting that the smoother operates at all under these conditions, because the underlying PLL is slipping cycles very frequently. From the curve for PLL simulation results in Fig. 2(a), we see that a P_c/N_0 of 10.8 dB-Hz is required to achieve an effective ρ of 10 dB with a PLL. The smoother threshold of $P_c/N_0 = 4 \text{ dB-Hz}$ to achieve a 10-dB effective ρ is approximately 6.8-dB lower than the PLL threshold of $P_c/N_0 = 10.8 \text{ dB-Hz}$.

2. Results for Other Values of Phase Noise. The results for other values of phase-noise spectral density, shown in Figs. 2(a) and 2(b), scale as expected according to the linear theory. In the linear region, increasing phase-noise density by 10 dB decreases effective ρ by 3.30 dB, when P_c/N_0 is held constant. Increasing phase-noise density by 10 dB increases the P_c/N_0 required to achieve a given effective ρ by 5 dB; specifically, the threshold is increased by 5 dB. Table 3 summarizes the results for each phase-noise density at a point near threshold and includes the noise bandwidths of the PLLs and the smoothers.

B. Results for FFT-based Frequency Estimation

Using an FFT-based scheme to estimate the frequency of the tracked signal would seem to enable operation at lower SNR values because the FFT does not suffer from cycle slips. However, in our

Table 3. Comparison of the causal and noncausal receiver performances, for $L(1) = 1 \times 10^{-3} \text{ rad}^2/\text{Hz}$, $\alpha = 3$.

Input parameters		Causal receiver, type II DPLL		Noncausal receiver with triangular weights, aided by PLL-derived frequency estimates		SNR improvement
P_c/N_0 , dB-Hz	$L(1)$, rad^2/Hz	Optimum B_L , Hz	Theory/simulation effective ρ , dB	Optimum B_L , Hz	Theory/simulation effective ρ , dB	Theory/simulation $\Delta\rho$, dB
0.0	1×10^{-4}	0.15	6.6/ 5.4	0.05	11.4/ 10.9	4.8/ 5.5
5.0	1×10^{-3}	0.46	6.6/ 5.4	0.15	11.4/ 10.9	4.8/ 5.5
10.0	1×10^{-2}	1.44	6.67/ 5.4	0.47	11.5/ 10.9	4.8/ 5.5

simulations, this scheme resulted in poorer SNR performance than the PLL-based scheme. Table 4 shows the relative SNR performance when using PLL- and FFT-derived frequency estimates. FFT computations were based on the same samples that the smoother used in its observation window. The samples were scaled by weights with a triangular windowing shape (the same as for the smoother); then the samples were padded with enough zeroes to fill an FFT array of 2^{13} complex samples in order to improve frequency resolution. There exist methods to improve the accuracy of FFT frequency estimates, but these were not pursued in our simulations.

Table 4. Comparison of the smoother’s SNR performance with frequency estimates derived from a PLL, an FFT, and the smoother’s last $N/2$ phase values, for $L(1) = 1 \times 10^{-3} \text{ rad}^2/\text{Hz}$, $\alpha = 3$.

P_c/N_0 , dB-Hz	ρ_{eff} ideal, dB	ρ_{eff} with PLL, dB	ρ_{eff} using FFT, dB	FFT penalty in ρ , dB	ρ_{eff} using last $N/2$, dB	$N/2$ penalty in ρ , dB
6.00	12.1	11.7	10.4	1.3	10.8	0.9
5.00	11.4	10.7	8.3	2.4	9.97	0.7
4.00	10.8	9.93	6.3	3.63	8.47	1.46

C. Results for Frequency Estimation by Phase Differencing

Frequency estimates based on the smoother’s last two phase estimates were not as good as the ones obtained from PLL frequency values. When the smoother is linearized by PLL-derived frequency estimates, the two-point estimates are reasonably good, although a little noisier than the other frequency-estimating schemes. However, when the smoother is controlled by the two-point smoother-derived frequency estimates, the performance drops noticeably. The performance of this scheme needs further investigation. Increasing the number of smoother phase-estimated values from the last two to the last $N/2$ improves the quality of the frequency estimates and, hence, the performance of the smoother, but still this improvement falls short of what one gets with the PLL-based algorithm. Table 4 shows the relative SNR performance when using PLLs and the last $N/2$ smoother-points-derived frequency estimates.

D. Results for FLL-Based Frequency Estimation

No useful results for our study were obtained from simulations with a cross-product frequency-locked loop, which is considered a “classic” among similar automatic frequency control devices [4]. The main limitation of this loop, as well as of all similar FLL devices, is its inherent squaring loss, which requires a relatively high P_c/N_0 value in order to stay locked on a signal with rapidly changing frequency. In the region of interest to us, which was for P_c/N_0 values below 15 dB-Hz, the FLL would simply not track the fast changing frequency of a signal corrupted by phase noise. This may have been because the various parameters of the FLL, including predetection bandwidth and loop bandwidth, were not optimized. This optimization was outside of the scope of this work.

In summary, the simulation has shown that frequency estimates obtained from linear fits to NCO phase values are reasonably good even when the PLL is experiencing occasional cycle slips, which become quite frequent below 5.5 dB of effective loop SNR.

V. Iterative Smoothing

As suggested by Dr. Peter Kinman, the smoothing operation can be performed multiple times over the same data set to improve the smoother’s phase estimates and SNR performance. One way of implementing

this iterative scheme is first to run the smoother with any of the desired frequency compensation schemes while saving to a file the smoother's phase estimates. Then, on the next smoother run, these saved phase values are used to get more accurate frequency estimates. Care should be taken to align properly in time past and current phase estimates. Because the smoother's effective SNR always is better than that of a PLL for the same signal conditions, the frequency estimates computed from the smoother's phase values also will be better than the ones derived from the PLL's phase values (or any other scheme that we tested). Hence, the smoother's performance improves on the second run. The smoothing process can be repeated any number of times until no further performance improvement is achieved. Figure 3 depicts the smoother's effective SNR after several iterative passes at $P_c/N_0 = 4.5$ dB-Hz and $L(1) = 10^{-3}$ rad²/Hz.

The curve in Fig. 3 was obtained averaging statistics over 6×10^6 samples per point with PLL-based frequency estimates on the first pass. With the iterative smoothing approach, the most noticeable improvement in ρ is achieved on the second pass, with asymptotically diminishing improvements on each subsequent pass. The improvement in SNR on the second pass is about 0.27 dB. At high SNR values, say at $P_c/N_0 > 7$ dB-Hz, when the degradation is smaller than 0.3 dB, repeating the smoothing process might not be worthwhile.

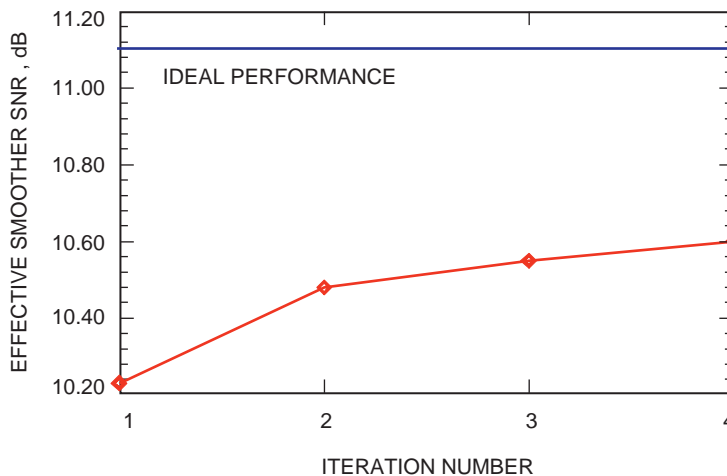


Fig. 3. Effective smoother SNR versus iteration number ($P_c/N_0 = 4.5$ dB; $L(1) = 1.0 \times 10^{-3}$ rad²/Hz; sampling rate = 100 samples/s; and batch size = 6×10^6 samples).

VI. Conclusion

It has been shown that a smoothing filter has significant advantages over a phase-locked loop and is practical to implement. Specifically, it has been shown that

- (1) The threshold signal SNR, P_c/N_0 , for phase-coherent tracking useful for telemetry systems is approximately 6.8-dB lower when a smoother is used in conjunction with a PLL than when a PLL alone is used.
- (2) Linear theory, which predicts that a smoother can achieve a given phase-estimation error with 7.5-dB lower signal SNR than a PLL, is valid when the SNR is 2 dB or more above threshold.
- (3) For fixed conditions of signal SNR and phase noise, the smoother achieves an effective loop SNR, ρ , of 5-dB higher than a PLL.

- (4) When a PLL operates at the minimum threshold for useful telemetry, the smoother improves the telemetry bit SNR (reduces radio loss) by 2.7 dB with (15,1/4) convolutionally encoded data.
- (5) At low signal SNR values, the smoothing process can be repeated several times in order to improve performance.

References

- [1] W. J. Hurd, "Optimum and Practical Noncausal Smoothing Filters for Estimating Carrier Phase With Phase Process Noise," *The Telecommunications and Data Acquisition Progress Report 42-128, October–December 1996*, Jet Propulsion Laboratory, Pasadena, California, pp. 1–7, February 15, 1997.
http://tmo.jpl.nasa.gov/tmo/progress_report/42-128/128H.pdf
- [2] V. A. Vilnrotter, W. J. Hurd, and D. H. Brown, "Optimized Tracking of RF Carriers With Phase Noise, Including Pioneer 10 Results," *The Telecommunications and Data Acquisition Progress Report 42-91, July–September 1987*, Jet Propulsion Laboratory, Pasadena, California, pp. 141–157, November 15, 1987.
http://tmo.jpl.nasa.gov/tmo/progress_report/42-91/91O.PDF
- [3] *DSN/Flight Project Interface Design Handbook*, JPL Document 810-5, Rev. D, vol. 1, "TLM-21: DSN Telemetry System, Block V Receiver," Jet Propulsion Laboratory, Pasadena, California, December 1, 1996.
<http://deepspace1.jpl.nasa.gov/dsndocs/810-5/tl21/tl21.pdf>
- [4] F. D. Natali, "AFC Tracking Algorithms," *IEEE Transactions of Communications*, vol. COM-32, no. 8, pp. 935–947, August 1984.

Appendix

Description of the Simulation Algorithm

Let P_c/N_0 be the ratio of carrier power-to-thermal noise spectral density of the tracked signal, Δt be the sampling time, and $L(f) = (1)/f^\alpha$ be the two-sided spectral density of the phase noise ($\alpha = 3$ for flicker noise and $\alpha = 8/3$ for scintillation noise). During the computer simulations, phase-noise samples θ_i with the desired spectral density are read from a file. These samples phase modulate the complex signal, $\mathbf{s}_i = e^{j\theta_i}$, which, for convenience, has been normalized to have $P_c = 1$. Thermal noise samples, \mathbf{n}_i , being zero-mean complex Gaussian random variables, with variance for each component $\sigma_n^2 = N_0/2P_c\Delta t$, are added to the signal. The resulting complex signal-plus-noise samples, $\mathbf{x}_i = \mathbf{s}_i + \mathbf{n}_i$, are fed into the first-in, first-out (FIFO)-type smoother array of size N .

I. Smoother Parameters

Using results given in [1], the smoother bandwidth is set to its optimum value, namely,

$$B_{sm} = b_{sm} \left(2L(1) \frac{P_c}{N_0} \right)^{1/\alpha} \quad (\text{A-1})$$

where

$$b_{sm} = \begin{cases} 0.816 & \text{for triangular weights window, } \alpha = 3 \\ 0.805 & \text{for triangular weights window, } \alpha = 8/3 \\ 0.900 & \text{for rectangular weights window, } \alpha = 3 \\ 0.895 & \text{for rectangular weights window, } \alpha = 8/3 \end{cases} \quad (\text{A-2})$$

This bandwidth determines the smoother window width, T_{sm} , which in turn determines N from the relation

$$T_{sm} = N\Delta t = \begin{cases} 2/(3B_{sm}) & \text{for triangular weighting window} \\ 1/(2B_{sm}) & \text{for rectangular weighting window} \end{cases} \quad (\text{A-3})$$

The predicted contribution to phase-error variance due to thermal noise will be $\sigma_{sm_th}^2 = B_{sm}/(P_c/N_0)$ and, due to phase noise, will be $\sigma_{sm_ph}^2 = \sigma_{sm_th}^2/(\alpha - 1)$. The total variance will be simply the sum of the two, namely, $\sigma_{sm_tot}^2 = \sigma_{sm_th}^2 + \sigma_{sm_ph}^2$.

II. PLL Parameters

A PLL precedes the smoother to track the unmodeled phase components of the tracked signal. The NCO instantaneous phase values, θ_{nco_i} , which are stored in a FIFO-type array, are needed to estimate $\hat{\theta}_{-i}$ and $\hat{\theta}_{-i}$ to linearize the phase in the smoothing interval; see Fig. 1. Because the tracking by a PLL distorts the spectrum of the tracked signal, it is necessary to restore the original signal spectrum so that the smoother operates under optimum conditions. This is done during the correlation process by introducing the term $e^{j\theta_{nco_i}}$ into the reference signal, Eq. (A-7). This in effect cancels the PLL factor $e^{-j\theta_{nco_i}}$.

For optimum performance, the PLL bandwidth is computed as suggested in [1], namely,

$$B_{pll} = b_{pll} \left(2S(1) \frac{P_c}{N_0} \right)^{1/\alpha} \quad (\text{A-4})$$

where

$$b_{pll} = \begin{cases} 2.489 & \text{for type-II loop, } \alpha = 3 \\ 2.234 & \text{for type-II loop, } \alpha = 8/3 \\ 2.752 & \text{for type-III loop, } \alpha = 3 \\ 2.490 & \text{for type-III loop, } \alpha = 8/3 \end{cases} \quad (\text{A-5})$$

The PLL variance due to thermal noise is $\sigma_{pll_th}^2 = B_{pll}/(P_c/N_0)$ and to phase noise is $\sigma_{pll_ph}^2 = \sigma_{pll_th}^2/(\alpha - 1)$. The total variance is again the sum of the two, namely, $\sigma_{pll_tot}^2 = \sigma_{pll_th}^2 + \sigma_{pll_ph}^2$.

The frequency of the tracked signal, θ_i , at the instant $t_{i-(N+1)/2}$ was estimated from the slope of the least-squares linear fits to the N PLL NCO phase values, $\theta_{nco,i} \cdots \theta_{nco,1-N}$ (the same N as for the smoother).

III. Smoothing Algorithm

The main steps of the smoothing algorithm are (1) correlation, (2) delta phase estimate, (3) phase-estimate update, (4) phase-error and signal-frequency computation, and (5) model phase update. Mathematically, these steps are as follows.

A. Correlation

At time t_i , the complex signal-plus-noise sample $\mathbf{x}_i = (\mathbf{s}_i + \mathbf{n}_i)$ is fed into the FIFO-type array $\{\mathbf{x}\}$. The array is correlated with the reference signal, \mathbf{r} , to give

$$\mathbf{Z}_{i-(N+1)/2} = \sum_{k=1}^{k=N} \mathbf{x}_{i-k} \mathbf{r}_{i-k} w_k \quad (\text{A-6})$$

The variable k designates the relative position in the array and corresponds to a delay of $k\Delta T$ s. The purpose of the reference

$$\mathbf{r}_{i-k} = e^{-j(-\theta_{nco,j-k} + \theta_{m,i-(N+1)/2} + \hat{\theta}_{i-(N+1)/2-1} \Delta T((N+1)/2-k))} \quad (\text{A-7})$$

is to restore the spectrum of the tracked signal by removing the NCO phase value, and to “linearize” the signal inside the smoothing interval in order to improve the estimation of θ_i at time $t_{i-(N+1)/2}$, i.e., in the middle of the array. Here $\theta_{m,i-(N+1)/2}$ is the “model” or predicted phase, $\hat{\theta}_{i-(N+1)/2-1}$ is the estimated frequency for time $t_{i-(N+1)/2-1}$, and w_k is the triangular (or some other) weighting function.

B. Delta Phase Estimation

The smoother computes the phase difference between the tracked and the reference signal by performing the four-quadrant arc-tangent operation on the complex variable \mathbf{Z} defined in Eq. (A-6):

$$\Delta \hat{\theta}_{i-(N+1)/2} = \arg [\mathbf{Z}_{i-(N+1)/2}] \quad (\text{A-8})$$

C. Phase-Estimate Update

The total phase estimate is obtained by correcting the model phase estimate with the value of the residual phase obtained from the correlation process, namely,

$$\hat{\theta}_{i-(N+1)/2} = \hat{\theta}_{m,i-(N+1)/2} + \Delta\hat{\theta}_{i-(N+1)/2} \quad (\text{A-9})$$

D. Phase-Error and Frequency Computation

The phase error, defined as the difference between the estimated and the actual phase values, is

$$\phi_{i-(N+1)/2} = \theta_{i-(N+1)/2} - \hat{\theta}_{i-(N+1)/2} \quad (\text{A-10})$$

As stated earlier, the instantaneous values of $\hat{\theta}$, redefined here as $\hat{\theta}_{i-(N+1)/2}$, were computed from least-squares fits to the NCO phase values of the PLL.

E. Phase-Model Update

Finally, the modeled or predicted phase for the next iteration is computed using the relation

$$\theta_{m,i-(N+1)/2+1} = \hat{\theta}_{i-(N+1)/2} + \hat{\dot{\theta}}_{i-(N+1)/2}\Delta T \quad (\text{A-11})$$

The operations described in Subsections III.A through III.E are repeated for each smoothing phase update instant.

The statistics of the phase error, ϕ_i , are evaluated at the end of the run, and the effective SNR of the smoother, defined as

$$\rho_{sm} = 10 \log_{10} \left[\frac{1}{\sigma_{\phi}^2} \right] \quad (\text{A-12})$$

is computed.

Tables 2(a) through 2(f) and Figs. 2(a) and 2(b) summarize the theoretical and the simulation results.