Frequency Estimation in the Presence of Cycle Slips: Filter Banks and Error Bounds for Phase Unwrapping

K. Kastella,* R. Mudumbai,** and T. Stevens*

*SRI International
Ann Arbor, MI
[Keith.Kastella, Troy.Stevens]@SRI.com

**Electrical & Computer Engineering
The University of Iowa, Iowa City, IA
RMudumbai@engineering.uiowa.edu

***Abstract*—We consider a setting in which a receiver uses a sequence of short, narrowband training burst signals from a transmitter to jointly estimate the time delay and frequency offset of its local clock with respect to the transmitter. A key challenge in this estimation problem is in handling cycle-slips arising from ambiguities in phase unwrapping when (a) the repetition rate of the training signal is small compared to the frequency offset, and (b) the bandwidth of the training signal is small relative to the carrier frequency. We propose a novel Bayesian filter-bank approach to handling these ambiguities. We present numerical simulations to show the effectiveness of this approach and compare our results with the fundamental posterior Cramer-Rao lower bound. The filter achieves the bound for signals between about 5 and 35 dB SNR, showing that it is optimal in this regime.**

Keywords-frequency estimation, cycle slips, filter banks, phase unwrapping, nonlinear filtering, posterior Cramer-Rao bound

I. INTRODUCTION

We consider the problem of a receiver jointly estimating the time delay and frequency offset of its clock with respect to a transmitter. While this is a classic estimation problem, we focus on one aspect of the problem that has not been systematically studied: estimation algorithms that can handle phase unwrapping ambiguities on delay and frequency estimates.

The setup for our estimation problem is illustrated in Fig. 1, which shows a system where a transmitter regularly sends a pulsed training signal of duration T_p. The receiver's clock, in general, is driven by a different local oscillator than the one the transmitter uses, and has timing and frequency offsets with respect to the transmitter's clock. The problem is to estimate these offsets. We consider a nonlinear filtering approach where the receiver takes an initial estimate and updates it every time it receives a training signal. The clock offsets vary in time due to the stochastic dynamics of the clock drifts, and/or Doppler effects due to relative motion between the transmitter and the receiver.

There is substantial literature on frequency and delay estimation, which are classic problems in estimation theory. Delay estimation using narrowband training signals was studied by Weiss and Weinstein [1], who derived bounds tighter than the CRLB that explicitly account for the effect of carrier phase ambiguities.

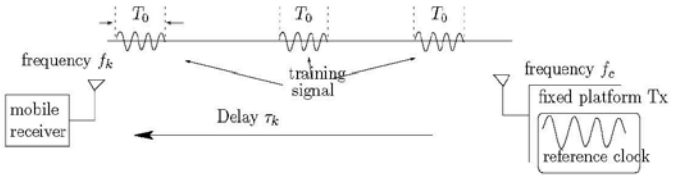

Figure 1. Signaling set-up for delay/frequency tracking problem.

To see the effect of phase ambiguities on the delay estimation problem, consider a narrowband training signal $s(t) = Re[b(t)e^{j2\pi f_c t}]$ sent by the transmitter. A noisy, delayed version of this signal $r(t) = Re\left[b(t-\tau)e^{j2\pi f_c(t-\tau)}\right]$ is received at the receiver. Now consider the same signal received with a slightly different delay $\tau' = \tau + 1/f_c$. This new signal $r'(t) = Re\left[b(t-\tau')e^{j2\pi f_c(t-\tau')}\right] = Re\left[b(t-\tau-1/f_c)e^{j2\pi f_c(t-\tau)}\right]$ has the same carrier phase as $r(t)$ and differs from $r(t)$ only in the envelope $b(t-\tau-1/f_c)$. However, by definition, the envelope narrowband signal $b(t)$ varies very little over a carrier period i.e. $b(t-\tau-1/f_c) \approx b(t-\tau)$. In other words, a delay of $\tau' = \tau + 1/f_c$ is virtually indistinguishable from a delay of τ.

A recently introduced maximum likelihood estimation (MLE) Newton-search algorithm [2] approaches the Weiss-Weinstein performance bounds. The MLE can fail at modest SNR, since it tracks only a single likelihood maximum. The contribution of this work is that it develops a Bayesian nonlinear filtering scheme that can track several likelihood local maxima. The particle filter is less likely than the MLE to be captured by suboptimal local maxima, improving performance. The resulting state estimate is evaluated as the conditional expectation over the entire posterior density. We evaluate the posterior Cramer-Rao bound (PCRB) and show that the filter achieves the bound for a range of SNR values. No unbiased estimator can achieve better estimation performance than the particle filter in this regime.

This paper is organized as follows. Section II formulates the problem, Section III details the particle filter solution, and Section IV derives the posterior Cramer-Rao bound. Numerical results are given in Section V. Section VI concludes with a discussion of combined delay and frequency ambiguities due to phase unwrapping, and suggests how this

work can be extended to incorporate slow-time jitter to address the problem of frequency ambiguities.

II. Problem Formulation

Our objective is to track the phase offset at a clock on a mobile platform relative to the clock on a fixed reference. The reference clock is assumed perfect and the phase offset is generated by the combined effects of the mobile's clock drift and motion. At discrete times $t_k = kT$, $k = 1, \dots, N_k$, the reference transmits the passband pulse of duration $T_0 \ll T$ with envelope $b(t)$ and carrier frequency f_c :

$$u_k^p(t) = Re[b(t)e^{j2\pi f_c t}],\ 0 \le t \le T_0. \tag{1}$$

We follow radar literature practice and refer to t_k as the slow-time index and let t denote the fast-time within the pulse relative to t_k. The signal is received at the mobile and demodulated to yield the complex baseband waveform

$$y_k(t) = Ab(t - \tau_k)\exp(j(-2\pi f_c\tau_k - 2\pi\nu_k t) + n_k(t). \tag{2}$$

Denoting clock phase on the mobile at t_k by θ_k, the total delay due to propagation and clock drift is $\tau_k = \theta_k/2\pi f_c$. The instantaneous frequency of the mobile is $f_k = f_c + \nu_k = f_c + \dot{\theta}_k/2\pi$. The power spectral density of the baseband noise $n_k(t)$ is N_0. The phase between the carrier and the envelope b is 0. The received amplitude $A = 1$ is assumed known. τ and ν are slowly varying and can be treated as constant within the pulse.

The phase is tracked by a discrete-time filter with slow-time sampling interval T. The phase is modeled using a two-state model of the oscillator dynamics,

$$x_{k+1} = Fx_k + w_{k+1}, \tag{3}$$

where $x_k = \left(\theta_k, \dot{\theta}_k\right)^{\mathrm{T}}$. The state transition matrix F and the real-valued 0-mean 2-vector white Gaussian process noise obey

$$F = \begin{pmatrix} 1 & T \\ 0 & 1 \end{pmatrix} \quad w_k \sim N(0, Q) \tag{4}$$

and

$$Q = q_c^2\begin{pmatrix} T & 0 \\ 0 & 0 \end{pmatrix} + q_m^2\begin{pmatrix} T^3/3 & T^2/2 \\ T^2/2 & T \end{pmatrix}. \tag{5}$$

The leading contribution of the clock drift is Brownian motion with strength q_c, while the platform motion contributes Wiener process noise of strength q_m.

On the mobile, the received signal is sampled in fast-time with period t_s. The sampled signal is $y_k[l] = y_k(lt_s)$, $l = 1, \dots, N_s$. The covariance of the complex noise samples $n_k[l]$ is $\sigma_n^2 = N_0/t_s$. Denote the sampled replica signal with delay $\tau' = \theta_k/(2\pi f_c)$ and frequency offset $\nu' = \dot{\theta}_k/(2\pi)$ by $b(\tau', \nu')[l] = b(t_s l - \tau')\exp(j(-2\pi f_c\tau' - 2\pi\nu' t_s l)$.

$$\left[\{x_k^i, m_k^i\}_{i=1}^{N_p}\right] = SIR\left[\{x_{k-1}^i, m_{k-1}^i\}_{i=1}^{N_p}, y_k(t)\right]$$

FOR $i = 1{:}N_p$

  Draw $x_k^i \sim p(x|x_{k-1}^i)$

  Calculate $m_k^i = p(y_k|x_k^i)m_{k-1}^i$

END FOR

Calculate total weight $t = \sum_{i=1}^{N_p} m_k^i$

FOR $i = 1{:}N_p$

  Normalize: $w_k^i := t^{-1}w_k^i$

END FOR

Calculate $\hat{x}_k = \sum_{i=1}^{N_p} w_k^i x_k^i$,

  $P_k = \sum_{i=1}^{N_p} w_k^i\left(x_k^i - \hat{x}_k\right)\left(x_k^i - \hat{x}_k\right)^{\mathrm{T}}$,

  and $N_{eff} = \left(\sum_{i=1}^{N_p}\left(w_k^i\right)^2\right)^{-1}$

IF $N_{eff} < N_{req}$

  Resample

END IF

Figure 2. SIRPF Algorithm

Up to an additive constant, the log-likelihood ratio for the observed data $y_k[l]$ given x_k is

$$\log p(y_k|x_k) = \frac{1}{\sigma_n^2}\mathrm{Re}\sum_{l=1}^{N_s} y_k[l]b^*(\tau', \nu')[l] \tag{6}$$

III. Particle Filtering

Particle filtering [3] is a Monte Carlo approximation to the Bayes-optimal recursive filter. For this study, we have implemented the Sampling Importance Resampling Particle Filter (SIRPF) algorithm detailed in [4]. At time step k the algorithm stores samples of the oscillator state (particles) x_k^i, and normalized weights m_k^i, $i = 1, \dots, N_p$ used to approximate the posterior density

$$p(x|y_k, \dots, y_0) \approx \sum_{i=1}^{N_p} m_k^i\delta\left(x - x_k^i\right). \tag{7}$$

The samples and weights are updated according to Fig. 2. In SIRPF, (3) is used to draw time-updated samples $x_k^i \sim p(x|x_{k-1}^i)$. The time-updated samples are used in (6) to update the weights as $m_k^i = p\left(y_k|x_k^i\right)m_{k-1}^i$, which are then normalized and used to evaluate the estimated state $\hat{x}_k$ and error covariance P_k. If this recursion is followed over time without interruption, the system degenerates and most of the weights become small, while a single weight tends towards unity. This pathological behavior is avoided by computing the

effective sample size $N_{eff} = 1/\sum_{i=1}^{N_p}(w_k^i)^2$ and resampling with replacement when $N_{eff} < N_{req}$. In principle, the required number of effective particles is a parameter that could be optimized, but in this work, $N_{rec} = N_p/2$.

IV. Posterior Cramer Rao Bound

Let $\hat{x}_k$ be an estimator of the system state x_k. The PCRB on the estimation error is [5]

$$E[(\hat{x}_k - x_k)(\hat{x}_k - x_k)^{\mathrm{T}}] \geq (J_k)^{-1}, \tag{8}$$

where the Fisher information matrix J_k obeys the recursion

$$J_{k+1} = D_k^{22} - (D^{12})^{\mathrm{T}}(J_k + D^{11})^{-1}D^{12}, \tag{9}$$

with $D^{11} = F^{\mathrm{T}}Q^{-1}F$, $D^{12} = -F^TQ^{-1}$,

$$D_k^{22} = Q_k^{-1} - \frac{\partial^2 \Lambda_k}{\partial x_k^2}, \tag{10}$$

and $\Lambda(x_k) = E[\log p(y_k|x_k)]$. The expected log-likelihood is evaluated by approximating (6) as an integral, yielding

$$\Lambda(x_k) = \frac{2}{N_0}\chi_b(\tau' - \tau, \nu' - \nu) \times \cos(2\pi f_c(\tau - \tau') - 2\pi(\nu - \nu')(\tau + \tau')/2), \tag{11}$$

where the ambiguity function for waveform b is

$$\chi_b(\tau, \nu) = \int b(t + \tau/2)b^*(t - \tau/2)\exp(-j2\pi\nu t)dt. \tag{12}$$

In the following simulation example, we use the Gaussian chirp waveform with pulse length T_p, chirp rate α, $b_{GC}(t) = \exp\left(-\pi(t/T_p)^2 + j\pi\alpha t^2\right)$, and ambiguity function

$$\chi_{GC}(\tau, \nu) = \frac{T_p}{\sqrt{2}}\exp\left[-\frac{\pi}{2}\left\{\left(\frac{\tau}{T_p}\right)^2 + \left(T_p(\nu - \alpha\tau)\right)^2\right\}\right]. \tag{13}$$

For b_{GC}, the energy per pulse is $E_p = T_p/\sqrt{2}$ and the output SNR is $\eta = 2E_p/N_0$. The log-likelihood derivatives in (10) required to evaluate the Fisher information (9) are

$$\frac{\partial^2 \Lambda_k}{\partial \theta^2} = -\frac{\eta}{(2\pi f_c)^2}\left[(2\pi f_c)^2 + \frac{\pi}{{T_p}^2} + \pi\alpha^2{T_p}^2\right], \tag{14}$$

$$\frac{\partial^2 \Lambda_k}{\partial \dot{\theta}^2} = -\frac{\eta}{(2\pi f_c)^2}\left[\pi(f_cT_p)^2 + \theta^2\right], \tag{15}$$

and

$$\frac{\partial^2 \Lambda_k}{\partial \theta \partial \dot{\theta}} = \frac{\eta}{2\pi f_c}\left[\frac{\alpha{T_p}^2}{2} + \theta\right]. \tag{16}$$

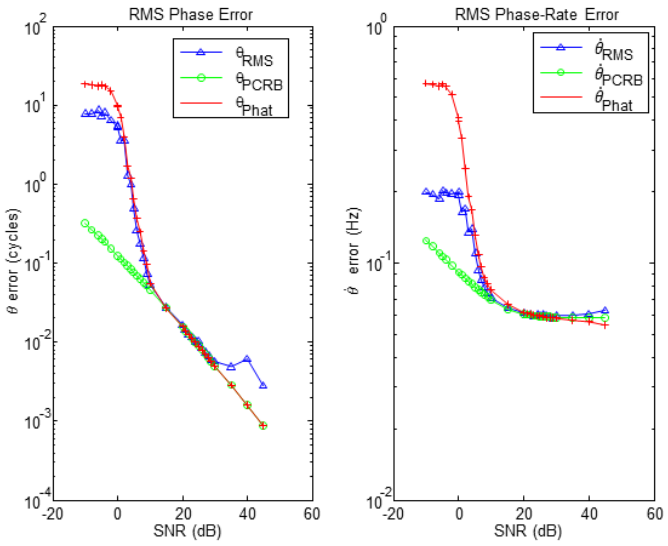

Figure 3. Phase and phase rate error results (blue diamond) compared to the CRLB (green circle) and filter-generated (red cross) error estimate.

V. Results

Fig. 3 shows how the filter performance varies with output SNR, $\eta = 2E_p/N_0$. In this simulation we have used $N_p = 1000$ particles. The signal parameters are filter period $T = 1$ s, center frequency $f_c =$ 200 MHz, pulse width $T_p =$ 200 μs, bandwidth $B = \alpha/T_p =$ 20 MHz, and sampling period $t_s =$ 50 ns. The clock process noise parameter is $q_c = 2 \cdot 10^{-10}\sqrt{s}$, which corresponds to an Allan deviation [6] $\sigma_y(\tau) = q_c/\sqrt{\tau}$ comparable to that of existing chip scale atomic clocks [7]. The motion process noise was selected as representative of a tethered aerostat application and is $q_m = 0.1$ m/s$^{-3/2}$. For each SNR value, 100 Monte Carlo trials were performed. Each trial ran for 100 s of simulated time. For each trial, the filter was initialized with truth and run for 50 time steps (50 s) to avoid initialization transients. Error data was then collected over the last 50 time steps. The exhibited error data (θ_{RMS} and $\dot{\theta}_{RMS}$) are the RMS error averaged over trials for each SNR value. The PCRB results (θ_{PCRB} and $\dot{\theta}_{PCRB}$) were obtained analogously, averaging over the second 50 time steps of each trial (note that the PCRB depends on the realized trajectory through the θ-dependance of the Fisher information). The plots also show the average estimated error P_k (θ_P and $\dot{\theta}_P$) computed by the filter. At high SNR (above about 5 dB), the three error measurements are in good agreement. At lower SNR, the filter does not achieve the PCRB, but shows the characteristic jump in error away from the CRB in the Weiss-Weinstein region.

VI. Frequency Estimation Ambiguities

The focus up to this point has been on the effect of phase ambiguity on delay estimation, but similar ambiguities also arise for frequency estimation when the interval between training signals becomes large. The frequency estimation problem with infrequent retransmissions of short duration training signals is the mathematical dual of delay estimation using narrowband signals.

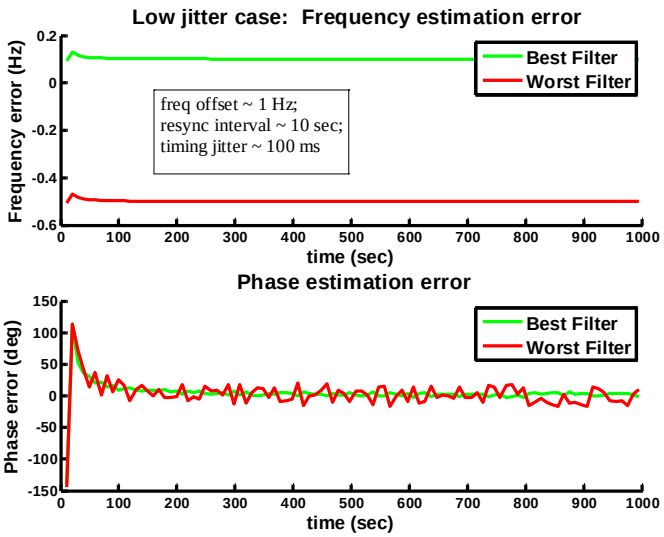

Figure 4. Frequency estimation performance with low jitter. The best filter retains a frequency bias in this low-jitter case.

This duality is explained as follows. Consider a simple estimation problem where a transmitter sends a training signal $s(t) = Re[b(t)e^{j2\pi f_c t}]$ and the receiver receives this signal with a frequency offset $y(t) = Re\left[b(t)e^{j2\pi(f_c+\Delta f)t}\right]$. The receiver uses this signal to construct an estimate $\widehat{\Delta f}$ of the frequency offset Δf. Conceptually, this is equivalent to estimating the slope of the phase angle of $y(t)b^*(t)$ as it varies with time. Now consider a simple delay estimation problem where a transmitter sends the same training signal $s(t) = Re[b(t)e^{j2\pi f_c t}]$, which is received with a delay $y(t) = Re[b(t-\tau)e^{j2\pi f_c(t-\tau)}]$. Taking Fourier Transforms, we have $Y(f) = B(f+f_c)e^{-j2\pi f\tau}, f \geq 0$, and we can think of the delay estimation problem as estimating the slope of the phase angle of $Y(f)B^*(f+f_c)$.

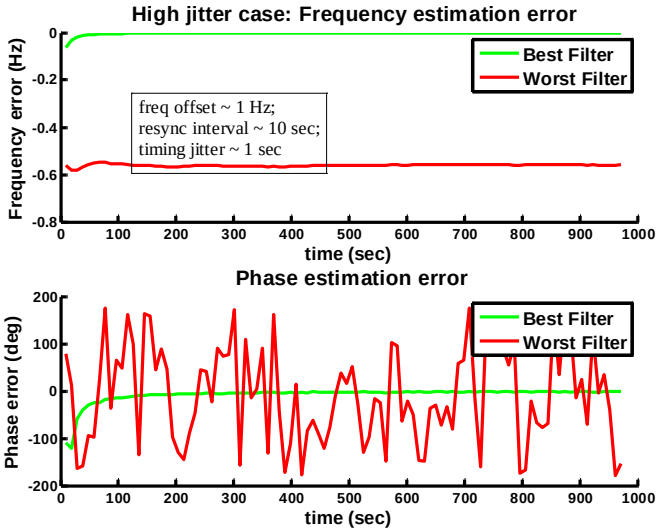

Figure 5. Frequency estimation performance with higher jitter. The best filter converges to an unbiased estimate in this high-jitter case.

Just as narrowband training signals lead to ambiguities of integer multiples of $1/f_c$ for delay estimates, short duration training signals repeated with period T lead to frequency estimate ambiguities of integer multiples of $1/T$. Increasing the bandwidth of the training signals helps resolve the $1/f_c$ delay ambiguities, and similarly, adding jitter (non-uniform slow-time sampling) to the resynchronization interval T helps resolve the $1/T$ frequency ambiguities. In principle, given enough jitter, it is possible to obtain frequency estimates to an arbitrary degree of accuracy even with infrequent training signal transmissions.

The potential utility of timing jitter for resolving frequency ambiguities is illustrated in Figs. 4 and 5, which show the estimation performance of a Kalman filter bank. The figures show frequency and phase estimation errors of two Kalman filters. These filters represent the best and worst rms phase estimates out of a bank of Kalman filters, the different filters in the bank corresponding to different initializations of the frequency estimate. In Fig. 4, where the timing jitter is only 1% of the resynchronization interval, it can be seen that the "best filter" (as measured by rms phase estimation error) is not the filter with zero steady-state frequency error. In contrast, Fig. 5 shows that when the timing jitter is increased to 10% of the resynchronization interval, the "best filter" does converge to the correct frequency estimate.

REFERENCES

[1] A. Weiss and E. Weinstein, "Fundamental limitations in passive time delay estimation–part i: Narrow-band systems," IEEE Transactions on Acoustics, Speech, and Signal Processing, vol. 31, no. 2, pp. 472–486, April 1983.

[2] P. Bidigare, U. Madhow, R. Mudumbai, and D. Scherber, "Attaining fundamental bounds on timing synchronization," IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), 2012 (in press).

[3] N. Gordon, D. Salmond, and A. Smith, "Novel approach to nonlinear non-Gaussian Bayesian state estimation," IEEE Proc. F, vol. 140, no. 2, pp. 107–113, 1993.

[4] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear / non-Gaussian Bayesian tracking," IEEE Trans. Sig. Proc., vol. 50, no. 2, pp. 174–188, February 2002.

[5] P. Tichavsky, C. H. Muarvchik, and A. Nehorai, "Posterior Cramer-Rao bounds for discrete-time nonlinear filtering," IEEE Trans. Sig. Proc., vol. 46, no. 5, pp. 1386–1396, May 1998.

[6] L. Galleani, "A tutorial on the two-state model of the atomic clock noise," Metrologia, vol. 45, pp. S175–S182, 2008.

[7] Symmetricom SA.45s CSAC Data Sheet, www.symmetricom.com.