

ATTAINING FUNDAMENTAL BOUNDS ON TIMING SYNCHRONIZATION

Patrick Bidigare[†]

Upamanyu Madhow[‡]

Raghu Mudumbai^{*}

Dzul Scherber[†]

[†] Raytheon BBN Technologies, Arlington, VA 22207, [pbidigar,dzul.scherber]@bbn.com

[‡] ECE Department, University of California, Santa Barbara, CA 93106, madhow@ece.ucsb.edu

^{*} ECE Department, The University of Iowa, Iowa City IA 52242, rmudumbai@engineering.uiowa.edu

ABSTRACT

In this paper, we propose an algorithm for timing synchronization that attains fundamental bounds derived by Weiss and Weinstein. These bounds state that, in addition to improving with time-bandwidth product and signal-to-noise ratio (SNR), timing accuracy also improves as the carrier frequency gets larger, if the SNR is above a threshold. Our algorithm essentially follows the logic of the Weiss-Weinstein bound, and has the following stages: coarse estimation using time domain samples, fine-grained estimation using a Newton algorithm in the frequency domain, and final refinement to within a small fraction of a carrier cycle. While the results here are of fundamental interest, we are motivated to push the limits of synchronization to enable the tight coordination required for emulating virtual antenna arrays using a collection of cooperating nodes.

Index Terms— timing synchronization, estimation, Newton method

1. INTRODUCTION

We revisit the classical problem of timing estimation based on a bandpass signal in additive white Gaussian noise (AWGN). It was shown decades ago by Weiss and Weinstein [1, 2] that, in principle, it is possible to obtain timing accuracies that improve not only with time-bandwidth product and SNR, but also with the carrier frequency, as long as the SNR is large enough. In this paper, we present an algorithm that attains the Weiss-Weinstein bounds, and provides concrete intuition regarding the structure of these bounds. Startlingly good performance is feasible: picoseconds accuracy for a carrier frequency of 1 GHz, bandwidth of 50 MHz, duration of 10 microseconds, and SNR of 10 dB. Furthermore, while such timing accuracies are orders of magnitude smaller than the inverse bandwidth, they can be attained with digital signal processing (DSP) implementations based on samples at the Nyquist rate or a small multiple thereof.

The problem of pure delay estimation, while classical in its origin, has not received much attention in recent years, perhaps because in most applications of interest (e.g., wireless communication), signals encounter multipath propaga-

tion. Especially in communications applications, the overwhelming emphasis has been on channel estimation that is “good enough” for communication, based on discrete time models at Nyquist rate or a small multiple thereof. Interestingly, we are motivated to revisit the classical problem of delay estimation because of our interest in emerging applications in cooperative communication. Specifically, we would like to be able to emulate a virtual antenna array using a collection of cooperating nodes, and to this end, would like to tightly coordinate their timing and carrier synchronization. If the links between neighbors in this cooperative cluster are near line-of-sight, then exchanges of timestamped messages can enable the nodes to synchronize their clocks, assuming symmetry in propagation and circuit delays. For the remainder of this paper, however, we do not discuss such applications, and focus instead on the fundamental problem at hand.

We note that there has been extensive research on timing estimation algorithms spanning several decades [3, 4, 5, 6, 7]. A key to the success of our timing estimation algorithm, however, is the use of a Newton algorithm in the frequency domain, an idea which we borrow from Newton-based frequency estimation algorithm from more than three decades ago [8]. Another important source of inspiration is provided by the Weiss-Weinstein bounds themselves, which guide us to simple high-SNR refinements of a baseband estimate.

2. MODEL

The transmitter sends the passband waveform

$$u_p(t) = \text{Re} \left(b(t) e^{j2\pi f_c t} \right), \quad 0 \leq t \leq T_o, \quad (1)$$

where $b(t)$ is a complex baseband waveform designed to have good autocorrelation properties, approximately timelimited to observation interval T_o , approximately bandlimited to bandwidth W , and f_c is the carrier frequency. The received passband signal is given by (up to scale)

$$y_p(t) = A \text{Re} \left(b(t - \tau) e^{j2\pi f_c (t - \tau)} \right) + n_p(t)$$

where $n_p(t)$ is passband WGN, and τ is the delay to be estimated. We emphasize that the delay is being estimated with

respect to the receiver's timebase. That is, $t = 0$ is an arbitrary time reference at the receiver, and it seeks to decide what is the best way of τ such that the received noisy signal best matches the noiseless template $\text{Re}(b(t - \tau)e^{j2\pi f_c(t - \tau)})$. It can be shown that, for typical time durations of interest and typical carrier frequency uncertainties, carrier frequency offsets between transmitter and receiver can be ignored, at least for stationary nodes.

Downconverting (recall that this is with respect to the receiver's local oscillator $e^{j2\pi f_c t}$), we obtain the complex baseband received signal

$$y(t) = b(t - \tau)e^{-j2\pi f_c \tau} + n(t) \quad (2)$$

where n is complex WGN. We will actually work with the discrete time signal obtained by sampling this at rate $1/T$ (typically a small integer multiple of the bandwidth W),

3. ALGORITHM

While we implement the algorithm in discrete time, it is useful to first sketch it in continuous time. We seek the maximum likelihood (ML) delay estimate, given by maximizing the following cost function, based on the log likelihood for the model (2):

$$J_{bp}(\tau) = \text{Re} \left(\int y(t)b^*(t - \tau)e^{j2\pi f_c \tau} dt \right) \quad (3)$$

where the subscript denotes "bandpass." The carrier term $e^{j2\pi f_c \tau}$ is what enables us to get amazing accuracy in the bandpass regime, but it is also responsible for rapid fluctuations in the cost function that makes it difficult to pick out the maximum directly. Accordingly, we adopt a multi-stage strategy that mirrors the different stages of the Weiss-Weinstein bound:

(a) First, treat the phase term $e^{-j2\pi f_c \tau}$ as a nuisance term, and perform *noncoherent* delay estimation maximizing the cost function $J_{nc}(\tau) = |\int y(t)b^*(t - \tau)dt|^2$. This is a well behaved cost function, whose maximization provides timing accuracy that, in principle, improves with bandwidth (of $b(t)$) and SNR. As we shall see, we implement this in two stages, passing through a "hypothesis testing" regime and a "fine-grained estimation" regime.

(b) Once we have a delay estimate based on $J_{nc}(\tau)$, then we pick the peak of $J_{bp}(\tau)$ nearest to it. For example, setting $\int y(t)b^*(t - \tau)dt = re^{j\theta}$ ($r \geq 0$), we can make a correction of size $\frac{\theta}{2\pi f_c}$ in the delay estimate in order to make $J_{bp}(\tau) = r \geq 0$ (we also consider a more sophisticated refinement, described later). Of course, such refinements only work when stage (a) has brought us to within a carrier cycle of the true delay, which only happens when the SNR is large enough to place us in the bandpass regime.

We have swept under the rug a key challenge in stage (a). Given samples at the Nyquist rate (or a small multiple

thereof), how do we get estimation accuracies that are arbitrarily smaller than the sampling accuracy? Rather than time domain interpolation strategies that depend in a complicated manner on the transmitted waveform, it is effective to express the cost function in the frequency domain:

$$J_{nc}(\tau) = \left| \int y(t)b^*(t - \tau)dt \right|^2 = \left| \int Y(f)B^*(f)e^{j2\pi f \tau} df \right|^2 \quad (4)$$

In the frequency domain, it is easy to compute the derivatives of the cost function with respect to τ , in a manner independent of the choice of the pulse b , since they involve differentiating a smooth complex exponential. Once we get close enough to the true delay (e.g., using a discrete time matched filter), we employ a Newton method to rapidly refine the delay estimate.

We now discuss discrete time implementation using samples of $y(t)$ at rate $1/T$, denoted by $\{y[n] = y(nT)\}$. The corresponding samples of the baseband template waveform $b(t)$ are denoted by $\{b[n] = b(nT)\}$. We proceed in three stages.

3.1. Stage 1: Hypothesis Testing

We begin by "hypothesis testing" style coarse delay estimation, passing $\{y[n]\}$ through the discrete time matched filter $b_{mf}[n] = b^*[-n]$ and picking the peak. If the peak is at \hat{n} , then the corresponding continuous time coarse delay estimate is $\hat{\tau}_1 = \hat{n}T$.

3.2. Stage 2: Fine-Grained Baseband Estimation

We are now going to maximize the cost function (4) using the frequency domain expression for it. To get its discrete time version, let M denote the smallest power of 2 at least as large as the length of the discrete time received signal and the discrete time template $b[n] = b(nT)$, compute the DFTs and then "fftshift" them so that DC falls in the center. For these centered DFTs, denoted by $\{Y[m]\}$ for the received signal and $\{B[m]\}$ for the template, the m th frequency corresponds to $f_m = m/(MT)$, $m = -M/2, \dots, M/2 - 1$. The discrete time approximation of the noncoherent cost function can now be written as $J(\tau) = |Q(\tau)|^2$ where

$$Q(\tau) = \sum_m Y[m]B^*[m]e^{j2\pi f_m \tau} \quad (5)$$

We can now use the hypothesis testing based estimate $\hat{\tau}_1$ as a starting point for Newton iterations of the form:

$$\hat{\tau}[n] = \hat{\tau}[n - 1] - \frac{J'(\hat{\tau}[n - 1])}{J''(\hat{\tau}[n - 1])}$$

($\hat{\tau}[0] = \hat{\tau}_1$). The first and second derivatives can be easily computed in terms of the following DFT-like expressions.

$$T_0 = \sum_m Y[m]B^*[m]e^{j2\pi f_m \tau}$$

$$T_1 = \sum_m Y[m]B^*[m]f_m e^{j2\pi f_m \tau}$$

$$T_2 = \sum_m Y[m]B^*[m]f_m^2 e^{j2\pi f_m \tau}$$

The derivatives are then given by (see appendix)

$$J'(\tau) = 4\pi \text{Im}(T_0 T_1^*)$$

$$J''(\tau) = 8\pi^2 (|T_1|^2 - \text{Re}(T_0 T_2^*))$$

We have found the Newton method to converge quickly, within 2-3 iterations. We denote the estimate at the end of this stage as $\hat{\tau}_2$.

3.3. Stage 3: Bandpass Refinement

In our final stage, we go back to the ML cost function (3), and applying Parseval's identity, approximate it using DFTs as $\tilde{J}(\tau) = \text{Re}(P(\tau))$, where

$$P(\tau) = \sum_m Y[m]B^*[m]e^{j2\pi(f_m + f_c)\tau} \quad (6)$$

Comparing with (5), we realize that $|P(\tau)| = |Q(\tau)|$. Once we have maximized $|Q(\tau)|$, and hence $|P(\tau)|$, to within a small fraction of a carrier cycle in stage 2, maximizing the real part of $P(\tau)$ can be accomplished by refining the delay estimate to set the phase, or the imaginary part, of $P(\tau)$ to zero. If the refinement is small enough (small fraction of a carrier cycle), it does not substantially change $|P(\tau)| = |Q(\tau)|$, whose variation is governed by the inverse bandwidth.

1. **One-shot adjustment.** In this method, we perturb the delay to set the phase of $P(\tau)$ to zero, resulting in the following delay estimate:

$$\hat{\tau}_3 = \hat{\tau}_2 - \frac{\arg(P(\hat{\tau}_2))}{2\pi f_c}$$

2. **Newton-based adjustment.** A more sophisticated refinement is to drive $G(\tau) = \text{Im}(P(\tau))$ to zero using another Newton stage, starting from $\tau = \hat{\tau}_2$. This corresponds to the iterations

$$\hat{\tau}[n] = \hat{\tau}[n-1] - \frac{G(\tau[n-1])}{G'(\tau[n-1])}$$

$$= \hat{\tau}[n-1] - \frac{\text{Im}(S_0)}{2\pi \text{Re}(S_1)} \quad (7)$$

where

$$S_0 = \sum_m Y[m]B^*[m]e^{j2\pi(f_m + f_c)\tau}$$

$$S_1 = \sum_m (f_m + f_c)Y[m]B^*[m]e^{j2\pi(f_m + f_c)\tau} \quad (8)$$

are analogous to the terms T_0 , T_1 defined earlier for baseband estimation in stage 2, except for having a carrier term. See appendix.

4. SIMULATION RESULTS




Fig. 1. Attaining the Weiss-Weinstein Bound

Figure 4 shows the Weiss-Weinstein bound [1, 2] and the simulated performance of our algorithm for a typical set of parameters. There is a nice interpretation for the various regions of the bound. For very low SNR, the delay estimate is garbage. As the SNR increases, we obtain accuracies close to the inverse bandwidth; we call this the “hypothesis testing regime,” since this can be viewed as testing a set of discrete delays spaced by the sampling interval, and picking the best. As the SNR increases further, we can get accuracies better than the inverse bandwidth by refining the hypothesis testing based estimate; we term this the “baseband estimation regime,” since the delay estimate here does not depend on the carrier frequency. Finally, once the SNR is large enough that the baseband estimate is within a carrier cycle of the truth, we can get to within an even smaller fraction of a carrier cycle by refining further; we can term this the “bandpass regime,” where the accuracy improves with carrier frequency. As seen from the plot, our algorithm essentially works through these regimes in a succession of stages. The plot also shows that the “one-shot” phase refinement algorithm is more robust than the Newton-based method near the Weiss-Weinstein threshold, but the latter gets closer to the CRLB at high SNRs, which is given by

$$\text{var}(\hat{\tau}) \geq (8\pi^2 \rho \langle f^2 \rangle)^{-1},$$

where ρ is the integrated SNR and

$$\langle f^2 \rangle = \frac{\sum_m (f_c + f_m)^2 |B[m]|^2}{\sum_m |B[m]|^2} \approx f_c^2$$

is the mean square frequency of the transmitted passband waveform. The CRLB explicitly shows that sub-carrier

period delay accuracy can be obtained. We also observe that both methods “plateau” at high SNRs: the one-shot refinement beyond 40 dB, and the Newton based refinement beyond 60 dB. While the plateau may be of little practical significance (the error is better than 10^{-4} carrier cycles for the one-shot adjustment and better than 10^{-6} carrier cycles for the Newton-based adjustment), it raises an interesting theoretical question as to how to attain the CRLB at arbitrarily high SNRs. A possible approach is to employ a composite cost function at the last stage, which simultaneously accounts for both the magnitude and phase of $P(\tau)$. We leave this as an interesting topic for future work.

5. CONCLUSIONS

We have presented a low-complexity delay estimation algorithm which approaches fundamental bounds. When the SNR is high enough, the accuracy can be a very small fraction of a carrier cycle, and improves with carrier frequency as well as with time-bandwidth product and SNR. We have carried out extensive simulations (not reported here due to lack of space) showing that the results are insensitive to specific waveform choices. Among the waveforms considered: up-chirp, up-down chirp, BPSK modulation using a raised cosine pulse, and MSK. In recent over-the-air experiments at BBN Raytheon, this algorithm was successfully used for fine-grained delay estimation.

The algorithm is easily extended to colored noise by suitably whitening the cost functions. Such extensions have been tested successfully using measured data on noise and interference in the television band, which has a highly colored power spectral density, together with synthetically generated signals.

Important topics for future work include performance evaluation and extensions of our algorithm to near-LoS environments in which there is some multipath, and its applications at the system level for the cooperative communication applications that motivated it.

Appendix: Derivatives for Newton method

Baseband estimation: Recall that $J(\tau) = |Q(\tau)|^2 = Q(\tau)Q^*(\tau)$, where $Q(\tau)$ is given by (5). It is easy to see that

$$Q'(\tau) = j2\pi T_1, \quad Q''(\tau) = -4\pi^2 T_2$$

where T_0, T_1, T_2 are as defined in Section 3.2. Since conjugation is interchangeable with differentiation with respect to a real variable, we obtain the following expressions for the derivatives of the noncoherent cost function:

$$\begin{aligned} J'(\tau) &= Q(\tau)Q'^*(\tau) + Q'(\tau)Q^*(\tau) \\ &= 2\text{Re}(Q(\tau)Q'^*(\tau)) = 4\pi\text{Im}(T_0T_1^*) \end{aligned}$$

$$\begin{aligned} J''(\tau) &= 2Q'(\tau)Q'^*(\tau) + Q(\tau)Q''^*(\tau) + Q''(\tau)Q^*(\tau) \\ &= 2|Q'(\tau)|^2 + 2\text{Re}(Q(\tau)Q''^*(\tau)) \\ &= 8\pi^2(|T_1|^2 - \text{Re}(T_0T_2^*)) \end{aligned}$$

Bandpass refinement: We wish to drive

$$G(\tau) = \text{Im}(P(\tau)) = \frac{P(\tau) - P^*(\tau)}{2j}$$

to zero. Note that $P(\tau) = S_0$, so that $G(\tau) = \text{Im}(S_0)$. Its derivative is given by

$$G'(\tau) = \frac{P'(\tau) - P'^*(\tau)}{2j}$$

From (6), we see that $P'(\tau) = j2\pi S_1$ and $P'^*(\tau) = -j2\pi S_1^*$. We therefore obtain after straightforward manipulations that $G'(\tau) = 2\pi\text{Re}(S_1)$.

6. REFERENCES

- [1] A. Weiss and E. Weinstein, “Fundamental limitations in passive time delay estimation—part i: Narrow-band systems,” *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 31, no. 2, pp. 472 – 486, apr 1983.
- [2] E. Weinstein and A. Weiss, “Fundamental limitations in passive time-delay estimation—part ii: Wide-band systems,” *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 32, no. 5, pp. 1064 – 1078, oct 1984.
- [3] M.-A. Pallas and G. Jourdain, “Active high resolution time delay estimation for large bt signals,” *Signal Processing, IEEE Transactions on*, vol. 39, no. 4, pp. 781 – 788, apr 1991.
- [4] H.C. So, P.C. Ching, and Y.T. Chan, “A new algorithm for explicit adaptation of time delay,” *Signal Processing, IEEE Transactions on*, vol. 42, no. 7, pp. 1816 – 1820, jul 1994.
- [5] Jian Li and Renbiao Wu, “An efficient algorithm for time delay estimation,” *Signal Processing, IEEE Transactions on*, vol. 46, no. 8, pp. 2231 – 2235, aug 1998.
- [6] Jr. Marple, S.L., “Estimating group delay and phase delay via discrete-time cross-correlation,” *Signal Processing, IEEE Transactions on*, vol. 47, no. 9, pp. 2604 – 2607, sep 1999.
- [7] Say Song Goh, T.N.T. Goodman, and F. Shang, “Joint estimation of time delay and doppler shift for band-limited signals,” *Signal Processing, IEEE Transactions on*, vol. 58, no. 9, pp. 4583 – 4594, sept. 2010.
- [8] T. Abatzoglou, “A fast maximum likelihood algorithm for frequency estimation of a sinusoid based on newton’s method,” *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 33, no. 1, pp. 77 – 89, feb 1985.