

Frequency Tracking of Periodic Signals in Noise

B.D.O. ANDERSON and P.P. PARKER

Dept of Systems Engineering, Australian National University, Canberra, A.C.T.

SUMMARY For a periodic signal measured in noise, this paper applies extended Kalman filtering to the problem of estimating the signal's frequency and the amplitudes and phases of the signal's first m harmonic components. The resultant estimator will also track the signal's frequency and its amplitudes and phases should these change over time. In this respect, it is unique among approaches to this problem. A partial theoretical analysis of the estimator appears in the paper. This analysis shows that there is some measure of decoupling in the estimator: The amplitudes are estimated as if the phase and frequency estimates were correct; the phases and frequency are estimated as if the amplitude estimates are correct. For the special case that the signal is a sinusoid and has known amplitude, the estimator becomes well known phase-locked loop. The paper also contains extensive simulations demonstrating both the tracking and the asymptotic behaviour of the estimator. The asymptotic behaviour is compared with results for another known estimators, and the relative strengths of each method examined.

1. INTRODUCTION

Consider a periodic signal which is not, in general, sinusoidal. Three sets of parameters characterise the signal: The frequency of the fundamental component; the amplitude of each harmonic component; and the phase of each harmonic component.

Suppose now that the signal is not exactly periodic but rather has frequency, amplitudes and phases that change slowly over time. Here, "slowly" means slowly compared to the period of the signal, so that over a few cycles the signal appears periodic. Suppose also that the signal is measured in the presence of additive noise. From such measurements we wish to estimate and track the time-varying frequency and the time-varying amplitudes and phase of the harmonic components.

The problem has application in many fields. For example, in sonar systems, hydrophones detect the engine noise of a vessel. Knowing the frequency, amplitudes, and phases may be sufficient for identifying the class of vessel generating the sound, hence the importance of the problem.

We now define some notation. Consider first a periodic signal $y(t)$ with zero d.c. component. A Fourier series representation is

$$y(t) = \sum_{k=1}^{\infty} r_k \sin(k\omega t + \phi_k) \quad (1.1)$$

Here, and throughout this paper, we will work in discrete time: $t = 0, 1, 2, \dots$. There is no real loss of generality in doing this. Now suppose that $y(t)$ is not exactly periodic, but has slowly time varying frequency ω , amplitudes r_k , and phases ϕ_k

$$\omega = \omega(t) \quad (1.2)$$

$$r_k = r_k(t) \quad (1.3)$$

$$\phi_k = \phi_k(t) \quad (1.4)$$

where $\omega(t)$, $r_k(t)$ and $\phi_k(t)$ are nearly constant over several cycles of $y(t)$. Now instead of $\omega t + \phi_k$ that appears in (1.1) we have

$$\theta_k(t) = \sum_{\tau=0}^t \omega(\tau) + \phi_k(t) \quad (1.5)$$

and thus

$$y(t) = \sum_{k=1}^{\infty} r_k(t) \sin \theta_k(t) \quad (1.6)$$

In this model of $y(t)$, the quantities $\omega(t)$, $r_k(t)$ and $\theta_k(t)$, are the instantaneous frequency, amplitudes, and phases. Later in this paper we will give specific models for the time variation of $\omega(t)$,

$r_k(t)$, and $\theta_k(t)$. An alternative, equivalent, parametrization for $y(t)$ replaces amplitudes and phases with sine and cosine components. That is,

$$y(t) = a_1 \sin \theta(t) + \sum_{k=2}^{\infty} [a_k(t) \sin k\theta(t) + b_k(t) \cos k\theta(t)] \quad (1.7)$$

where

$$\theta(t) = \sum_{\tau=0}^t \omega(\tau) \quad (1.8)$$

Here the parameters are $a_1(t)$, $a_2(t)$, ..., $b_2(t)$, ..., $\theta(t)$ and $\omega(t)$. Note that there is no cosinusoidal component at the fundamental frequency. This is because such a cosinusoidal component would represent a phase shift in the fundamental, which can equivalently be represented by a change in $\theta(t)$. Thus to include a $b_1(t) \cos \omega(t)$ component would be an overparametrization. Therefore, we keep to the parametrization of Equation (1.7).

We will call the first model, that of equations (1.5) and (1.6), the polar model and the second model, that of equations (1.7) and (1.8), the rectangular model. Most of the analysis and all of the simulations in this paper will be for the polar model; however, throughout the paper, we will point to places where the two models give similar results, and to places where they give differing results.

Regardless of which model we work with, we assume that the signal $y(t)$ is corrupted by noise. Denote the noise by $n(t)$ and let it be white and gaussian. Then we assume we have measurements

$$x(t) = y(t) + n(t) \quad (1.9)$$

available. The task, restated, is to take measurements $\{z(\tau): \tau = 0, 1, \dots, t\}$ and from these measurements, estimate $r_1(t)$, ..., $r_m(t)$, $\theta_1(t)$, ..., $\theta_m(t)$, $\omega(t)$ for the polar model, or $a_1(t)$, ..., $a_m(t)$, $b_2(t)$, ..., $b_m(t)$, $\theta(t)$, $\omega(t)$ for the rectangular model. Furthermore, we wish to update these estimates as each new measurement $z(t)$ becomes available, thus tracking changes in the parameters. Note carefully that we only estimate parameters up to the m^{th} harmonic: The higher harmonics are assumed to be negligible, perhaps having been eliminated by an anti-aliasing filter. In short, we approach the problem as a $2m + 1$ parameter estimation and tracking problem.

Let us note some qualitative issues surrounding the problem. In estimating the signal's frequency, not only the fundamental, but also the higher harmonics contain information. Indeed, in the extreme case that no fundamental exists, with only the second,

third, and higher harmonics present, the fundamental obviously contains no information about the frequency $\omega(t)$. Thus, we seek an estimator that uses not only the energy in the fundamental, but also the energy in the higher harmonics to estimate the signal's frequency. Moreover, the information about the frequency contained in any harmonic will depend upon the energy in that harmonic, or, more precisely, the signal to noise ratio of that harmonic. But we do not assume a priori knowledge of the different harmonics' amplitudes; indeed, estimating these harmonic amplitudes is part of our problem. Given that a knowledge of the harmonic amplitudes helps in finding the frequency, we would expect that the estimates of the harmonic amplitudes would also help in calculating the frequency. In actual fact, the estimator we propound here estimates the frequency $\omega(t)$ using estimates of the harmonic amplitudes as though they are the true, correct amplitudes. This is an intuitively appealing result.

The converse argument also applies. That is, a knowledge of the frequency helps in calculating the amplitudes and phases - or, for the rectangular model, the sine and cosine amplitudes. Indeed, if the frequency is known exactly, then estimating the sine and cosine components in the rectangular model is a linear estimation problem. Since we do not assume knowledge of the frequency, the best we can use is an estimate of the frequency. It is a hallmark of our estimator that it uses the frequency estimate as though it is the true frequency when estimating the sine and cosine components or the phases and amplitudes.

A short survey of other frequency estimators would now be necessary. Frequency estimation is a mature problem with many contributions from many authors. Naturally we cannot include an exhaustive survey. To highlight the salient aspects of other schemes, we divide them into four classes, depending upon the assumed characteristics of the incoming signal. These classes are (1) constant frequency, sinusoidal signal (i.e. no higher harmonics); (2) constant frequency, higher harmonics present; (3) time-varying frequency, sinusoidal signal; (4) time-varying frequency, higher harmonics present. The problem we address here falls into category (4), and is, to our knowledge, the only estimator yet to cover this category. The other categories have a number of candidate estimators as we now show.

Constant Frequency, Sinusoidal Signal

The most common estimator in this category is the discrete Fourier transform (DFT). By introducing another level of algorithm that takes the output from the DFT and passes it through a hidden Markov model estimator, Streit and Barrett (1988) have overcome some of these difficulties. The hidden Markov model estimator is recursive, and works at impressively low signal to noise ratio levels. Amongst the many other frequency estimators for constant frequency sinusoidal signals, is adaptive line enhancement (ALE) see Widrow and Glover (1975).

Constant Frequency, Higher Harmonics Present

In a recent paper, Nehorai and Porat (1986) extended the ALE method to cope with higher harmonics. Their method uses the information in all of the first m harmonics in finding the frequency. The second part of Nehorai and Porat's algorithm is to estimate the amplitudes of sinusoidal and cosinusoidal components. Barrett and McMahon (1987) weight the different harmonics according to their amplitude estimates when calculating the frequency estimates. In this respect, Barrett and McMahon improve upon Nehorai and Porat.

Time-Varying Frequency, Sinusoidal Signal

To this point, we have considered estimators for constant frequency signals; yet there does exist an estimator for time-varying frequency well known both through its theory and its application: It is the phase locked loop (PLL). Now the PLL is known to be an approximation of the Extended Kalman Filter (EKF), applied to the frequency estimation. See Anderson and Moore (1979) for a treatment of the EKF. Snyder (1969) gives the equivalence between the PLL and EKF in continuous time,

Kelly and Gupta (1972) give the discrete time case. But the EKF is a much more general tool than just the PLL, so it would seem possible to take the EKF as applied to the frequency estimation problem - which reduces to the PLL - and to generalize this EKF, allowing for higher harmonics with time varying, unknown amplitudes. It is just this path we follow in this paper.

2. THE ESTIMATOR

As already stated, we apply an extended Kalman filter to estimating frequency, phases, and amplitudes. To do this, we need a state space model for the quasi periodic signal defined for the polar case Equations (1.5) and (1.6) or alternatively for the rectangular case in Equations (1.7) and (1.8).

Taking the polar case first, and using the notation of (1.5) and (1.6) the state space model is

$$x(t+1) = Fx(t) + v(t) \quad (2.1)$$

$$z(t) = y(t) + n(t) \quad (2.2)$$

$$= h(x(t)) + n(t) \quad (2.3)$$

where

$$x(t) = [r_1(t), r_2(t), \dots, r_m(t), \omega(t), \theta_1(t), \theta_2(t), \dots, \theta_m(t)]^T \quad (2.4)$$

$$F = (f_{ij}) \text{ with } f_{ii} = \forall i, f_{m+2, m+1} = 1 \\ f_{m+3, m+1} = 2, \dots, f_{2m+1, m+1} = m \\ \text{otherwise } f_{ij} = 0 \quad (2.5)$$

$$h(x(t)) = \sum_{k=1}^m r_k(t) \sin \theta_k(t) \quad (2.6)$$

and $v(t)$ is white gaussian noise, with zero mean and variance

$$E[v(t)v(t)^T] = Q \quad (2.7)$$

while $n(t)$ is white, gaussian noise, of zero mean and

$$E[n(t)v(t)] = 0 \quad (2.8)$$

$$E[n^2(t)] = R \quad (2.9)$$

To this signal model, we apply an extended Kalman filter, which takes the measurement $z(t)$ and produces estimates $\hat{x}(t|t)$ or $\hat{x}(t|t-1)$ of $x(t)$. Here $\hat{x}(t|t)$ denotes an estimate of $x(t)$ given measurements $z(\tau)$ up to and including time t ; $\hat{x}(t|t-1)$ is a one step prediction of $x(t)$, that is an estimate of $x(t)$ given $z(\tau)$ up to time $t-1$. The extended Kalman filter, Anderson and Moore (1979) p. 195 is

$$\hat{x}(t|t) = \hat{x}(t|t-1) + L(t)[z(t) - h(\hat{x}(t|t-1))] \quad (2.10)$$

$$\hat{x}(t+1|t) = F \hat{x}(t|t) \quad (2.11)$$

$$L(t) = \Sigma(t) H(t) (H(t)^T \Sigma(t) H(t) + R)^{-1} \quad (2.12)$$

$$\Sigma(t+1) = F(\Sigma(t) - \Sigma(t) H(t) [H(t)^T \Sigma(t) H(t) + R]^{-1} H(t)^T \Sigma(t)) F^T + Q \quad (2.13)$$

where

$$H(t) = \frac{\partial h(\hat{x}(t|t-1))}{\partial \hat{x}(t|t-1)} \quad (2.14)$$

$$= [\sin \hat{\theta}_1(t|t-1) \dots \sin \hat{\theta}_m(t|t-1) \hat{O}^T r_1(t|t-1) \cos \hat{\theta}_1(t|t-1) \dots] \quad (2.15)$$

and the filter is initialized by

$$\hat{x}(0|0-1) = E[x(0)] = \bar{x}_0 \quad (2.16)$$

$$\Sigma(0) = E[(x(0) - \bar{x}_0)(x(0) - \bar{x}_0)^T] \quad (2.17)$$

$x(0)$ having been assumed a random variable. When quantities in (2.16) or (2.17) are unknown, then the algorithm must be initialized with some reasonable values of $\hat{x}(0|0-1)$ and $\Sigma(0)$. Thus, to implement our estimator requires initialization, plus the update at each time instant of the four equations (2.10), (2.11), (2.12) and (2.15). This compares well with more than twenty equations updated at every time instant in Nehorai and Porat's algorithm.

The application to the rectangular parametrization is almost the same as for the polar case. Equations (2.1) - (2.3) remain unchanged, while Equations (2.4) - (2.6) become, respectively

$$x(t) = [a_1(t), a_2(t), \dots, a_m(t), b_2(t), b_3(t), \dots, b_m(t), \omega(t), \theta(t)] \quad (2.18)$$

$$F = (f_{ij}) \text{ with } f_{ii} = 1 \forall i, f_{2m+1, 2m} = 1 \\ \text{otherwise } f_{ij} = 0 \quad (2.19)$$

$$h(x(t)) = a_1 \sin \theta(t) + \sum_{k=2}^m [a_k(t) \sin k\theta(t) + b_k \cos k\theta(t)] \quad (2.20)$$

where the notation is consistent with that in Equations (1.7) - (1.8). This defines the state space signal model for the rectangular case. The estimator for the rectangular case needs no change to the estimator for the polar case, save that Equation (2.15) becomes

$$H(t) = [\sin \hat{\theta}(t|t-1) \dots \cos m \hat{\theta}(t|t-1) 0 \hat{\gamma}(t|t-1)]^T \quad (2.21)$$

where

$$\hat{\gamma}(t|t-1) = \sum_{k=1}^m k \hat{a}_k(t|t-1) \cos k \hat{\theta}(t|t-1) - \sum_{k=2}^m k \hat{b}_k(t|t-1) \sin k \hat{\theta}(t|t-1) \quad (2.22)$$

As the model stands, it has $2m + 1$ states. The price of increasing the number of harmonics m is to increase the state size and thus the computational burden and complexity.

Whether the rectangular model and estimator is better or the polar model and estimator is better is a moot point. Since the models are non-linear, it is quite likely that the extended Kalman filter does perform better for one model compared to the other. However, the non-linearity would make any theoretical resolution very difficult. Perhaps simulations would resolve the matter; we have not performed such simulations here. However, it is common knowledge that for radar tracking systems, extended Kalman filters that use rectangular spatial coordinates outperform those that use polar spatial coordinates. It is also true that the rectangular model is more linear. However, if one actually wants to find the amplitudes and phases, it is easier to estimate them directly with the polar estimator rather than use the rectangular estimator and then make a rectangular-polar conversion. Therefore, each model has its advantages.

In both models there are some potential pitfalls of implementation, largely caused by the non-linear models. The first of these is that $\hat{\omega}$ could lock onto a fraction or multiple of the true frequency ω . The difficulty is inherent in the problem and could equally occur in Nehorai and Porat's estimator. Essentially it is a problem caused by poorly initializing the state estimate, in particular the frequency estimate. It is well known that a phase locked loop will not capture a signal if its centre frequency - that is initial frequency estimate - is too distant from the signal's actual frequency. Therefore, it is no surprise that good initialization is also important in our more general problem.

Another potential difficulty is cycle slipping. That is, θ_k and $\hat{\theta}_k$ for some k differing by a multiple of 2π . In simulations, such cycle slipping was only observed at low SNR levels, and then only in the transient phase.

A more subtle form of cycle slipping can occur only in the polar model. In this case \hat{r}_k could become negative for some k thus tracking $-r_k$ and $\hat{\theta}_k$ would then differ from θ_k by π . The remedy is easy. If \hat{r}_k is negative, change its sign, and add π to $\hat{\theta}_k$. This type of cycle slipping was also observed in simulations.

3. THEORETICAL UNDERSTANDING OF THE ESTIMATOR

Space limitations prevent us from setting out the results in detail. We shall summarize them here.

(a) As already noted, the estimator has two parts, one estimating frequency and phase, and using amplitude estimates as if they were exact, the other estimating amplitudes, using frequency and phase estimates as if they were exact

(b) the amplitude estimation uses a linear Kalman filter

(c) the frequency and phase estimation uses one summer

(integration in continuous time) for each phase and one for the frequency. The end result looks like an interconnection of second order phase-locked loops, all sharing one summer, and if there is only a fundamental, a conventional PLL is obtained.

4. SIMULATIONS

All the simulations presented here are for the polar model. The first set of simulations demonstrate asymptotic behaviour; the second set of simulations shows transient and tracking behaviour. Because the estimator we have propounded is highly non-linear - a necessary consequence of the problem - it is impossible to analyse it completely. Therefore, the simulations are an important part of evaluating the estimator's performance.

To make this evaluation easier, the asymptotic simulations are exactly the same as those performed by Nehorai and Porat (1986). We present both our results and those of Nehorai and Porat. The input signal to the estimator was, as in Nehorai and Porat (1986),

$$z(t) = \sum_{k=1}^5 r_k \sin 2\pi 0.08kt + n(t) \quad (4.1)$$

That is, the number of harmonics $m=5$, while $n(t)$ was unit variance, zero mean white gaussian noise. The signal's frequency, $\omega = 2\pi 0.08$, and the amplitudes r_k were constant over time. The amplitudes were

$$r_k = r_1/k \quad k = 2, 3, 4, 5 \quad (4.2)$$

where r_1 was chosen such that the signal to noise ratio, given by

$$\text{SNR(dB)} = 10 \log_{10} \left(\sum_{k=1}^5 r_k^2 / 2 \right) \quad (4.3)$$

was at 0dB, 8dB or 16dB depending on the trial. For each of the SNR levels of 0, 8 and 16dB, one hundred independent Monte Carlo experiments were run, each over 200 times samples. At the final time instant, the error in the state estimate was calculated, then the statistics over the hundred Monte Carlo runs were computed. Exactly the same experiments were repeated for 500 time samples on each Monte Carlo run.

When the algorithm was initialized at the start of each Monte Carlo run, the following initial conditions were used.

$$\hat{r}_1(0|1) = (3/4 \times E\{z(t)^2\})^{1/2} \quad (4.4)$$

$$\hat{r}_2(0|1) = \hat{r}_3(0|1) = \hat{r}_4(0|1) = \hat{r}_5(0|1) = 0 \quad (4.5)$$

The value of $\hat{r}_1(0|1)$ was chosen in (4.4) as a way of easily estimating the fundamental component's strength from only the statistics of the incoming signal. Setting the estimates of higher harmonics' amplitudes to zero in (4.5) reflects a lack of knowledge about the signal's harmonic structure as may well be encountered in practice.

Additional initial conditions were as follows

$$\hat{\theta}_k(0|1) = 0 \quad k = 1, \dots, 5 \quad (4.6)$$

and

$$\hat{\omega}(0|1) = 2\pi 0.05 \quad (4.7)$$

Equations (4.4) - (4.7) describe the initial, one step ahead, state estimate. The initial Σ matrix was

$$\Sigma(0) = \text{diagonal} \{ \sigma_1, \sigma_1/4, \sigma_1/9, \frac{\sigma_1}{16}, \frac{\sigma_1}{25}, (0.06\pi)^2, 0.02, \dots, 0.02 \} \quad (4.8)$$

where

$$\sigma_1 = (\hat{r}_1(0|1))^2/4 \quad (4.9)$$

Additionally we chose $R = 1$, and

$$Q = \text{diagonal} \{ q_1, q_1, \dots, q_1, q_2, q_1, \dots, q_1 \} \quad (4.10)$$

where $q_1 = 1 \times 10^{-4}$ and $q_2 = 1 \times 10^{-8}$, the latter being the element of Q corresponding to the frequency ω . The values of $\Sigma(0)$, Q and R were chosen by trial and error, and may or may not be close to the optimal values.

The results of the simulations appear in Tables I - III. In each of these tables, the numbers in parentheses are the results of Nehorai and Porat (1986), while the numbers without parentheses are our results. In Table I, the third column, labelled "outliers", is an appellation given by Nehorai and Porat to those

Monte Carlo experiments where the error in the estimated frequency was greater than an arbitrary threshold of 0.003π radian/sample. Since this error threshold is only 1.875% of the true frequency, it is rather a tight definition of an outlier; but, in the interests of comparison with Nehorai and Porat's results we have kept the same definition. For the same reason, we excluded the outlier trials when compiling the statistics of the frequency, amplitude, and phase estimates.

Examining Table I shows that while the number of outliers we incurred was approximately the same as Nehorai and Porat at SNR of 8dB and 16dB, at the lower SNR of 0dB we performed much worse. Indeed, at 0dB our method appears to work very poorly, scoring 39 outliers versus 7 for Nehorai and Porat for a number of time samples $N=200$, and 34 outliers versus 6 for $N=500$.

**TABLE I
STATISTICAL RESULTS FOR THE
FUNDAMENTAL FREQUENCY ESTIMATE**

N (samples)	SNR (dB)	Outliers	ω	
			Bias $\times 10^{-6}$	Standard Dev $\times 10^{-5}$
200	0	39 (7)	79.6 (-98.7)	40 (42.3)
	8	4(3)	11.8 (26.0)	10 (14.3)
	16	0 (0)	2.59(1.20)	10 (5.05)
500	0	34 (6)	4.76 (22.6)	20 (14.4)
	8	4 (3)	15.9(0.4)	10 (2.03)
	16	0 (0)	6.52(0.99)	10 (0.78)

**TABLE II
STATISTICAL RESULTS FOR THE AMPLITUDE
ESTIMATE**

r_1				
N (samples)	SNR (dB)	True	Bias	Standard Deviation
200	0	1.17	-0.12(-0.08)	0.19(0.15)
	8	2.94	-0.039(-0.13)	0.17(0.32)
	16	7.38	-0.007(-0.06)	0.11(0.22)
500	0	1.17	-0.053(-0.04)	0.11(0.10)
	8	2.94	-0.005(-0.02)	0.009(0.08)
	16	7.38	-0.01(-0.02)	0.07(0.12)

**TABLE III
STATISTICAL RESULTS FOR THE PHASE
ESTIMATES**

θ_1			
N (samples)	SNR (dB)	Bias $\times 10^{-1}$	Standard Deviation $\times 10^{-1}$
200	0	-4.47(0.68)	3.38(4.48)
	8	-4.97(-0.00)	0.55(1.25)
	16	-5.07(-0.08)	0.32(0.56)
500	0	-5.01(0.06)	1.49(2.42)
	8	-4.95(0.01)	0.54(0.51)
	16	-5.01(0.03)	0.33(0.25)

We now consider the tracking behaviour of our algorithm, something that other algorithms cannot emulate. Whereas the asymptotic results we presented were tabular and quantitative, the tracking results are graphical and quantitative. Figures 4.1 - 4.3 show the results of a single trial. In this case the signal was initially the same as that in the asymptotic results, with SNR = 8dB. However, we added some process noise $v(t)$ - see Equation (2.1) - into the signal model, causing the frequency, amplitudes, and phases to evolve with time. The process noise had variance

$$E[v(t)v(t)^T] = \text{diagonal } \{a_1, a_1, a_1, a_1, a_1, a_1, a_2, a_1, a_1, a_1, a_1, a_1\}$$

where $a_1 = 1 \times 10^{-3}$, $a_2 = 3 \times 10^{-7}$. Note that the state with variance of a_2 corresponds to the frequency ω . The initialization of the estimator was the same as for the asymptotic trials.

Figure 4.1 shows how the true frequency and its estimate evolve with time. Note that although the frequency estimate starts at a very inaccurate initial value, it locks onto the true value and tracks it closely after time $t = 50$. Figure 4.2 is an expanded scale version of Figure 4.1, showing the tracking behaviour from $t = 60$ to $t = 200$. Throughout this time, the frequency error never exceeds 0.6% of the actual frequency.

Figure 4.3 shows the estimated and true amplitudes of the fundamental frequency component r_1 . Tracking is achieved after $t = 60$. The amplitudes of harmonics and their estimates exhibited a similar tracking response and are not shown here. For the phases, there is quite a vigorous transient response, after which the error wanders fairly closely around zero radians.

The results show clearly that tracking is possible in our algorithm. After an initial transient phase, the estimated parameters settle down and track the true parameters. This is a unique quality of our algorithm. Other algorithms, specifically that of Nehorai and Porat, and that of Barrett and McMahon, cannot simultaneously track frequency, phase and amplitude variations even though such variations do occur in real world applications.

5. CONCLUSIONS

This paper has presented a new frequency estimation algorithm with a number of unique properties. The algorithm is based on the extended Kalman filter with its attendant ideas of mean square optimal estimates and recursive calculation. It is the only algorithm known to us that can track simultaneously amplitudes, phase, and frequency variations for a signal which is quasi-periodic rather than merely sinusoidal.

A number of possible extensions can be contemplated, e.g. to cope with Dopplers shift, and to marry this algorithm, with a hidden Markov model algorithm to achieve good low SNR performance.

A further extension is to fixed lag smoothing. By allowing a certain lag in producing the frequency, amplitude, and phase estimates, the accuracy of estimates should improve. Kumar (1987) has done some significant work upon smoothers applied to phase estimation of a sinusoidal signal.

A final extension is to the case when more than one signal is present, such as in sonar when two or more ships can be detected. Here the work of Sundresh, Cassara, and Schachter (1977) and Cassara, Schachter, and Simowitz (1980) may be of use. These authors looked at the problem of simultaneously demonstrating two different FM signals. Their estimator consisted of two inter-connected phase-locked loops. For our problem, the same structure could be used, replacing the two phase locked loops with two of our extended Kalman filter estimators.

These are but a few representatives of the many extensions that exist for this problem, and show that the scope of the problem reaches far beyond that which a single paper can contain.

REFERENCES

- ANDERSON, B.D.O. and MOORE, J.B. (1979). *Optimal Filtering*, Englewood Cliffs, New Jersey: Prentice Hall.
- BARRETT, R.F. and MCMAHON, D.R.A. (August 1987). "ML Estimation of the Fundamental Frequency of a Harmonic Series", Proc. ISSPA 87, pp 333-336, Brisbane, Australia.
- CASSARA, F.A., SCHACHTER, H. and SIMOWITZ, G.H. (June 1980). "Acquisition Behaviour of the Cross-Coupled Phase-Locked Loop FM Demodulator", IEEE Trans. Communications, Vol. COM-28, pp 897-904.

KELLY C.N. and GUPTA, S.C. (July 1972). "Discrete Time Demodulation of Continuous-Time Signals", IEEE Trans. Information Theory, Vol. IT-18, pp 488-493.

KUMAR, R. (1987). Optimum Filters and Smoothers Design for Carrier Phase and Frequency Tracking, NASA JPL Publication 87-10, Jet Propulsion Laboratory, Pasadena, California.

NEHORAI, A. and PORAT, B. (October 1986). "Adaptive Comb Filtering for Harmonic Signal Enhancement", IEEE Trans. Acoustics, Speech and Signal Processing, Vol. SSP-34, pp 1124-1138.

SNYDER, D.L. (1969). The State Space Approach to Continuous Estimation with Application to Analog Communication Theory, Cambridge, Mass: MIT Press.

STREIT, R.L. and BARRETT, R.F. "Frequency Line Tracking Using Hidden Markov Models", submitted to IEEE Trans. Acoustics, Speech and Signal Processing.

SUNDRESH, T.S., CASSARA, F.A. and SCHACHTER, H. (December 1980). "Maximum a Posteriori Estimator for Suppression of Interchannel Interference in FM Receivers", IEEE Trans. Communications, Vol. COM-25, pp 1480-1485.

WIDROW, B. and GLOVER, J.R., Jr. et al. (December 1975). "Adaptive Noise Cancelling: Principles and Applications", Proc. IEEE, Vol. 63, pp 1692-1716.




Figure 4.3 Amplitude of fundamental component: True Value and estimate. Tracking example.




Figure 4.1 True frequency and frequency estimate: Tracking example.




Figure 4.2 True frequency and frequency estimate: Expanded Scale. Tracking example.