KALMAN FILTERING WITH A TWO-STATE CLOCK MODEL

Fran B. Varnum, Defense Mapping Agency

ABSTRACT

In this paper, a formulation of the operative equations that might be employed to filter measurements taken of an atomic clock are developed in an intuitive way. Emphasis is upon making clear, in a practical example, concepts and matrix components that might otherwise appear somewhat abstract to the uninitiated in a formal derivation.

Particular attention is given to the concept and calculation of the process noise required to process timing data in a Kalman filter. The author has tried to present the material in such a manner that a person with little or no experience could begin experimenting with the use of a Kalman filter to process real or simulated timing data.

INTRODUCTION

This paper assumes a rudimentary understanding of the Kalman filter concept and algorithm. It's purpose is to extend that understanding through the intuitive development of a practical example. The example chosen is a 2-state (phase and frequency) filter to process time difference data between a subject clock and a reference clock. The reference is assumed perfect and the subject clock is assumed to be perturbed by white noise on it's frequency. How to include a white noise source on the frequency drift (random walk FM) is an extension that will be developed but not included in the full example.

The behavior of a cesium clock system is, for the most part, deterministic. However, random fluctuations in the on-going physical processes give rise to some measure of unpredictable behavior. It will be the function of the Kalman filter to make optimum estimates of the deterministic parameters (phase and frequency states) given measurements that are a function not only of these parameters but also of the perturbing noise source(s). This task is ideally suited for a Kalman filter.

One practical reason for estimating deterministic parameters is to use those parameters to predict a clock's behavior. For example, this is the purpose of estimating satellite clock states at the Master Control Station of the Global Positioning System (GPS).

Prediction of a clock's behavior will generally always be in error for two reasons. First, even though given the best possible values for the deterministic parameters describing past behavior, one could not predict subsequent random motion. Second, the original estimates of the deterministic parameters must necessarily have been made in the presence of random behavior and are, themselves, subject to error. The function of the Kalman filter is to make an "optimum" estimate of the deterministic parameters based on whatever data is available in the belief that these will provide the best possible prediction under the circumstances.

Included, additionally, as part of these circumstances is the user-specified "model" which describes the dynamical relationships between the selected filter states together with a specification of the noise sources that are assumed to exist within the clock. It is very important that the model adequately represent the system under consideration in order to obtain the best possible result. The model is usually expressed very compactly in the form of a vector matrix differential equation. In the following section, the standard clock model in state-space, vector-matrix notation will be examined.

CLOCK MODEL

$$\dot{X} = A\,X + B\,N \qquad (1)$$

Equation 1 is the matrix differential equation form of the clock model. X is the matrix vector of the states chosen to describe the deterministic behavior of the clock. In our case, as previously stated, we have chosen phase and frequency. Call these states X1 and X2 respectively. N is a matrix vector of the noise sources we are identifying as the cause of random behavior. Call these N1 and N2.

Equation (1) can expanded as:

$$\begin{vmatrix} \dot{X}1 \\ \dot{X}2 \end{vmatrix} = \begin{vmatrix} 0 & 1 \\ 0 & 0 \end{vmatrix} \begin{vmatrix} X1 \\ X2 \end{vmatrix} + \begin{vmatrix} 1 & 0 \\ 0 & 0 \end{vmatrix} \begin{vmatrix} N1 \\ N2 \end{vmatrix} \qquad (2)$$

N is a matrix vector consisting of zero mean, white, normally distributed noise sources. In our example, white FM is being modeled, so the elements in the B matrix are set appropriately to reflect this model.

If random walk frequency were to be modeled, then N2 should be incorporated into the model (random walk frequency is generated from white noise on frequency drift) by setting the zero in the lower right-hand corner of the B matrix to a one.

Performing the indicated matrix multiplication in equation (2), a pair of differential equations are obtained:

$$\dot{X}1 = X2 + N1 \qquad (3)$$

$$\dot{X}2 = 0 \qquad (4)$$

These equations are an identical but a more explicit statement of the clock model of equation (1). Equation (3) states that the time derivative of state 1 equals state 2 plus a white noise term. Remembering that states 1 and 2 are phase and frequency respectively and that only white noise FM is being considered, the reader should have no difficulty accepting this equation. Equation (4) states that the time derivative of the second state is zero, or, in other words, the frequency is being modeled as a constant. The fact that a real cesium exhibits a random walk in frequency is being ignored. If it were to be considered, then equation (4) would be:

$$\dot{X}2 = N2$$

Because phase, frequency, and frequency drift are integrally related and perfectly analogous to position, velocity, and acceleration, the basic clock model is identical in form to that used to model many dynamic systems in state space.

The next step, in preparation for implementing our Kalman filter, is to note the solution to the clock model. What is required is the solution to a stochastic differential equation and in our case one with constant coefficients (the A matrix).

SOLUTION TO THE CLOCK MODEL

The solution to the deterministic part of matrix equation (1) is:

$$e^{At} = \phi(t) = \sum_{n=0}^{\infty} \frac{1}{n!} t^n A^n \qquad n = 0,1,\ldots$$

Since $A = \begin{vmatrix} 0 & 1 \\ 0 & 0 \end{vmatrix}$

$$\phi(t) = \begin{vmatrix} 1 & t \\ 0 & 1 \end{vmatrix} \qquad \text{(n = 0, 1 are the only nonzero terms)}$$

$\phi(t)$ is the "transition matrix" which plays an important role in the Kalman filter algorithm. As the solution to equation (1), it describes the evolution of the modeled states with time and is a very aptly named quantity. The value of the states at any time, t , is:

$$X(t) = \phi(\delta t) \ X(t-\delta t) \qquad (5)$$

where δt is some interval of time prior to t. For the sake of notational simplicity let this earlier time be (T0). In expanded form (5) is:

$$\begin{vmatrix} X1(t) \\ X2(t) \end{vmatrix} = \begin{vmatrix} 1 & \delta t \\ 0 & 1 \end{vmatrix} \begin{vmatrix} X1(0) \\ X2(0) \end{vmatrix}$$

Notice that X(TO) has been simplified to X(0).

Performing the indicated multiplication gives:

$$X1(t) = X1(0) + X2\ \delta t$$

$$X2(t) = X2(0)$$

Expressing our state variables X1 and X2 (phase and frequency) as T and $\dot{T}$

$$T(t) = T(0) + \dot{T}(0)\delta t \qquad (6)$$

$$\dot{T}(t) = \dot{T}(0) \qquad (7)$$

Intuitively we see that these equations are correct. Equation (6) states that the phase at any time (t) is equal to the phase at an earlier time plus an accumulation due to the constant frequency offset. Equation (7) states that the frequency is constant.

The next concern, in finding the solution to equation (1), is to consider it's nondeterministic term. The solution to this part is a difficult integral. However the covariance of that integral is the quantity of interest (1). It is, itself, another integral, but one easily evaluated. Before examining it, a short discussion of it's physical significance is in order.

The covariance matrix of the solution to the random part of the clock model specifies the uncertainty in the clock's output due to the white noise sources incorporated in the model. In other words, it is a statistical measure of the inability of the deterministic states to completely model the clock's behavior. This is precisely the quantity needed to "Q" the Kalman filter; or more properly stated, it specifies the amount of "process noise" to be incorporated for each filter state.

The solution to the random part of the clock model is:

$$\int_{t-\delta t}^{t} \phi(u)\ BN\ du$$

with covariance:

$$\int_{t-\delta t}^{t} \phi(u)\ \mathrm{cov}\ (BN)\ \phi(u)'\ du \qquad (9)$$

Evaluation of this integral (Eq. 9) will specify the "Q's" required for each state in the filter. Because of its importance, the evaluation will be outlined for a more general model incorporating three states (frequency drift is added) and a random forcing function (white noise) associated with each state. The appropriate choice of zeros in the A and B matrices will tailor these results for a particular application.

EVALUATION OF CLOCK STATE NOISE COVARIANCE MATRIX

$$\text{Write } B\ N \text{ (Eq 1)} = \begin{vmatrix} S1 & 0 & 0 \\ 0 & S2 & 0 \\ 0 & 0 & S3 \end{vmatrix} \begin{vmatrix} N1 \\ N2 \\ N3 \end{vmatrix}$$

where S1, etc. are the standard deviations of the respective white noises in the model. In the previous development, the B matrix contained 1's which implied unit standard deviations.

By definition: $\quad COV\ BN = E\ (BN\ (BN)') = B\ E(NN')\ B'$

where E is the expectation and the prime indicates the matrix transpose.

Assuming the noise sources to be uncorrelated:

$$E(NN') = \delta \begin{vmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{vmatrix}$$

where δ indicates the Dirac delta function.

So:

$$COV\ BN = \begin{vmatrix} V1 & 0 & 0 \\ 0 & V2 & 0 \\ 0 & 0 & V3 \end{vmatrix}$$

where V1, etc., are the variances of the noise sources.

Now ϕ COV (BN) ϕ' is determined by performing the indicated multiplication. For the three state model, the transition matrix $\phi(\delta t)$ is:

$$\begin{vmatrix} 1 & t & \frac{\delta t^2}{2} \\ 0 & 1 & \delta t \\ 0 & 0 & 1 \end{vmatrix}$$

The result is:

$$\phi(\delta t)\ COV\ BN\ \phi(\delta t)' = M1 + M2\delta t + M3\delta t^2 + M4\delta t^3 + M5\delta t^4 \qquad (10)$$

The matrices M1 through M5 are:

$$M1 = \begin{vmatrix} V1 & 0 & 0 \\ 0 & V2 & 0 \\ 0 & 0 & V3 \end{vmatrix}$$

$$M2 = \begin{vmatrix} 0 & \frac{V1}{2} & 0 \\ \frac{V1}{2} & 0 & \frac{V2}{2} \\ 0 & \frac{V2}{2} & 0 \end{vmatrix} \qquad M3 = \begin{vmatrix} 0 & 0 & \frac{V1}{6} \\ 0 & \frac{V1}{3} & 0 \\ \frac{V1}{6} & 0 & \frac{V2}{3} \end{vmatrix}$$

$$M4 = \begin{vmatrix} 0 & 0 & 0 \\ 0 & 0 & \frac{V1}{8} \\ 0 & \frac{V1}{8} & 0 \end{vmatrix} \qquad M5 = \begin{vmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \frac{V1}{20} \end{vmatrix}$$

Finally, equation (10) is integrated over the interval δt, from ($t-\delta t$ to t) to obtain the $Q(\delta t)$ matrix required to optimally tune the Kalman filter. The result is:

$$Q(\delta t) = M1\delta t + M2(\delta t)^2 + M3(\delta t)^3 + M4(\delta t)^4 + M5(\delta t)^5$$

This same process, applied to the simpler model chosen for our example, gives:

$$Q(\delta t) = \begin{vmatrix} V1\delta t & 0 \\ 0 & 0 \end{vmatrix}$$

By inspection one sees that non-zero process noise (Q) is applied to the phase state only and that its magnitude is proportional to time: specifically, the time between measurement updates to the filter. In other words, while the last measurement will have presumably decreased the uncertainty in the phase state over what it had been previously, the uncertainty should start growing again until the next measurement. Some sense of the dynamics of a Kalman filter can be had at this point by imagining that due to a loss in some measurements, a longer than normal time elapses between the latest and the previous measurement. Because of the relatively larger value for "Q"(δt increased) and hence growth in uncertainty in the state, the Kalman gain will be correspondingly larger than otherwise for the state. This, in turn, causes the delayed measurement to have a larger impact on the estimate of this state than under the postulated "normal conditions".

It is this process that gives rise to the idea of a kalman filter "memory", and the fact that it is effected by the filter process noise. Generally larger values of process noise tend to weight the latest measurements more heavily.

The converse is that without the incorporation of adequate process noise, too little weight is given to the latest measurements. Suffice it to say, the choice of "Q's" is critical and can mean the difference between optimum and worthless results in a given situation. It is for this reason that some attention has been given to the concept of filter "Q'ing" in this discussion.

THE KALMAN FILTER ALGORITHM

The state variables (X) to be estimated have already been established as the phase and frequency (X1 and X2). Let it further be specified at this point that these are "epoch states". That is to say, a phase and frequency will be estimated for a specified time in the past such that when propagated to current time by the transition matrix, one obtains the corresponding "current-time" states. It can be shown that epoch and current-time state formulations result in identical estimates of the current-time states (3). If there is an advantage to estimating epoch states it is the relative ease with which predictions can be made relative to a fixed vis a variable epoch.

Use of epoch states introduces a conceptual subtlety that initially caused this author some difficulty in formulating the filter algorithm under discussion. It has to do with the state-transition matrix which has already been introduced and described validly as the model for the time-evolution of states. But note that epoch states, themselves, are generally constant by definition. Continuous reestimation is required due to the stochastic behavior of the clock, but in the absence of data the transition matrix for the epoch states is unity. There are, in other words, two transition matrices: the first describes the time evolution of states and was used to calculate the process noise. The second propagates epoch states between measurements. More importantly, it propagates the epoch state covariance matrix between measurements-a process called the "time update" and the first step in the Kalman cycle.

TIME UPDATE

This process propagates the epoch state covariance from its computed value immediately following the last measurement update to the time of the next measurement. Call these covariances P(k+) and P(k+1-) respectively. In this notation k+ means immediately after the last measurement update, k+1- means at the time of but before processing the next measurement. The time update consists of computing P(k+1-) given P(k+) from the last Kalman cycle.

$$P(k+1-) = P(k+) + \phi^{-1} \quad Q \quad \phi'^{-1} \qquad (11)$$

The inverse of the transition matrix means that the process noise added to account for the growth in uncertainty during the interval between measurements is propagated back to epoch. In our example pre and post multiplication by the inverse and the inverse of the transpose of the transition matrix leaves the Q matrix unchanged. If there had been a noise term associated with the frequency rate term (random walk frequency) this would not be the case. Note that the transition matrix normally encountered with the propagation of the P matrix itself is missing. This is, of course, because it is unity as discussed above.

The three remaining steps in the Kalman cycle are to:

compute the gain
make the new estimate (measurement update)
compute the new state covariance

Once this is complete, the cycle starts again with the time update as described in the previous paragraph.

Before writing the three equations that describe these steps only two new matrices need to be introduced. They are the measurement noise covariance matrix (R) and the measurement matrix (H).

The measurement noise matrix for our example is simply a 1 x 1 with its single element equal to the variance of the measurement error.

The measurement matrix (H) is the matrix that defines the relationship between the measurement and the states being estimated. In our example, the phase difference between a subject clock and a reference clock is being measured. The measurement is the state X1 propagated to current time (the time of the measurement). The transition matrix provides the relationship, namely, equation (6) which in matrix form is written:

$$T(t) = \begin{vmatrix} 1 & \Delta t \end{vmatrix} \quad \begin{vmatrix} T(0) \\ \dot{T}(0) \end{vmatrix}$$

Where Δt is the elapsed time since epoch.

Letting $T(t) = Z - v$, where Z is the measurement incorporating a measurement error (v) and defining $\begin{vmatrix} 1 & \Delta t \end{vmatrix} = H$

$$Z = H\ X - v \qquad (12)$$

which is the standard form for the measurement equation in the Kalman algorithm.

Notice that in order for equation (12) to be correct X must be the true clock states and not the estimated states. We expect some error in the estimate or else there would be no need to update it. The quantity $Z-H\hat{X}$ (where the hat indicates the estimate) plays a major role in the measurement update process. This residual carries the information relating to the difference between the actual measurement and measurement predicted from existing estimates.

To complete the Kalman cycle the operative equations are:

$$K = P(k+1-)\ H'[H\ P(k+1-)\ H' + R]^{-1} \qquad \text{(compute gain)}$$

$$\hat{X}(k+1) = \hat{X}(k) + K \quad Z - H\hat{X}(k) \qquad \text{(update estimate)}$$

$$P(k+1+) = [I - K\ H]\ P(k+1-) \qquad \text{(update covariance)}$$

These equations are easily implemented on a small computer. The largest matrix is a 2 X 2 and the entire algorithm is reducible to a few algebraic manipulations. Performance of the filter under a variety of conditions can be empirically evaluated. Simulated data can be created by integrating the output of a Gaussian random number generator where the numbers are interpreted as white noise on frequency. The standard deviation of the number generator should be chosen to appropriately represent the noise characteristics of the clock being simulated. This follows from the sample time chosen and the Allen variance that characterizes the clock. While the model chosen for this tutorial is over-simplified for a real situation, it should serve quite well as a first step in developing insight and experience with the use of Kalman filters to process timing data.

The two-state Kalman filter described in this paper was implemented and a simulated set of 500 data points processed with three different values of process noise. The magnitudes of the process noise were chosen to represent under, over, and optimum "Q'ing" respectively. In all other respects, including the input data, the runs were identical. The input data samples, (representing simulated phase difference measurements) were generated by integrating (summing) over a set of normally distributed random numbers. The resulting filter source data is shown plotted in Figure 3. The "true" states, in this instance, are zero for both phase and frequency since the only source of variation in the data is the random component. The filter states were initialized with values of 20ns and + 1 part in 10 E 13 respectively. The covariance matrix was initialized with the square of these values.

The resulting phase estimates for the three cases cited are shown plotted in Figure 1. Clearly, increasing the "Q" increases the variation in the phase estimate. In accordance with the development of the "Q" matrix for our simplified noise model, process noise is added to the phase state only, and it follows that this should happen. Perhaps it is not quite as obvious that just the reverse should happen to the frequency estimate even though it is not being "Q'ed". It is because the filter insures that the time history of the estimated "state pair" is such that it will always closely predict the measurement data (via the measurement matrix). Increasing the range over which the phase estimate is allowed (or forced) to vary, will result in a decrease in the variation required of the companion state. The plots of the frequency state estimates in Figure 2 confirm this. In fact, too much process noise for the phase state prevented the frequency estimate from changing at all. The important general observation is that the choice of "Q" for even a single state can effect the estimates of other states.

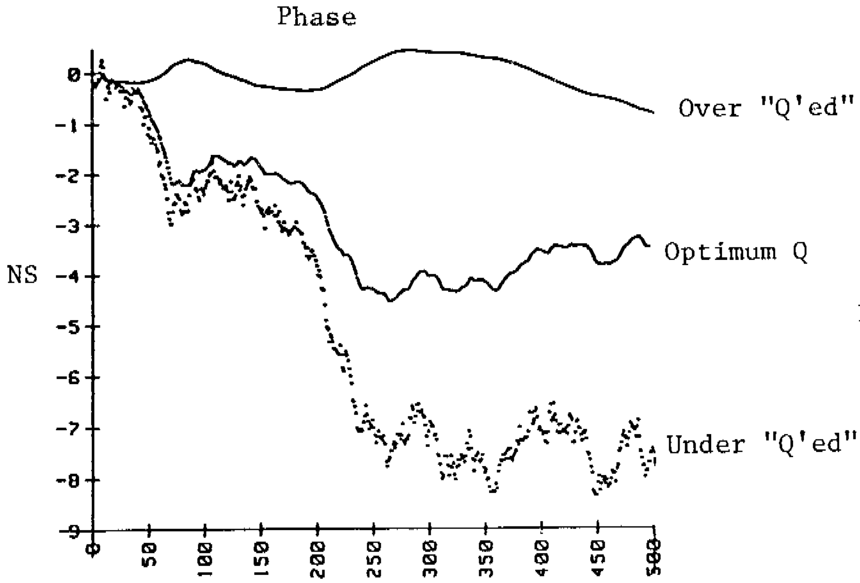

Figure 1

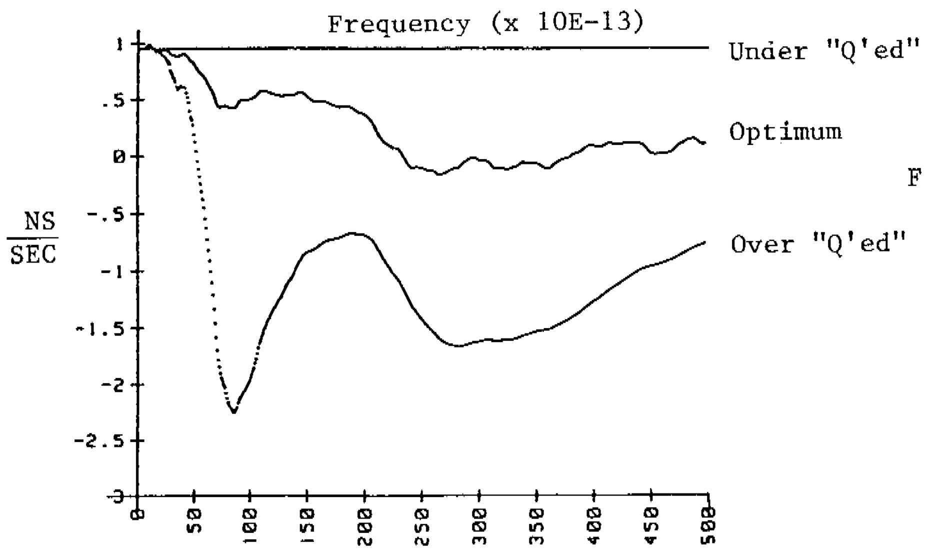

Figure 2

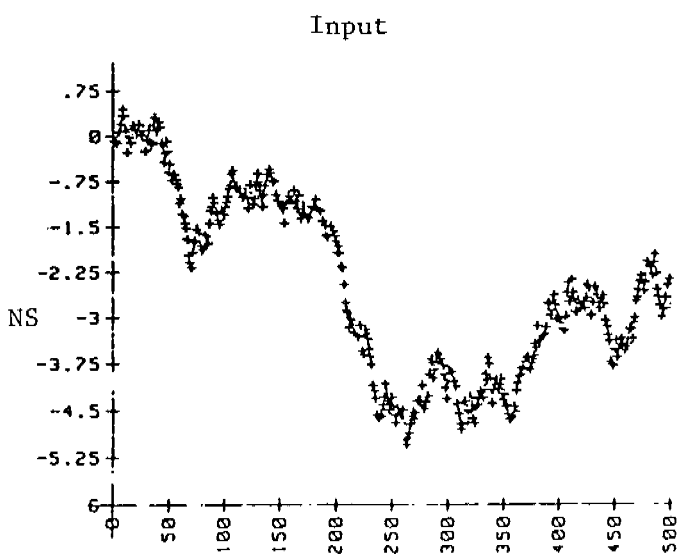

Figure 3

REFERENCES

1. Maybeck, P.S. "Stochastic Models, Estimation, and Control", Vol.1 Academic Press, 1979

2. Chaffee, J. "Derivation of the Process Noise Matrix for Satellite Clocks", unpublished memorandum (OAO Corp), March 1982.

3. Chaffee, J. "Exact Relationship between Epoch and Current Time Filters", unpublished memorandum (OAO Corp), October 1982.

4. Sorenson, H.W., Stubberud, A.R. "Linear Estimation Theory", Ch 1 "Theory and Applications of Kalman Filtering", distributed by The National Technical Information Service (NTIS), U.S. Department of Commerce, AGARDograph No. 139 1970.

5. Barham P.M., Humphries, D.E. "Further Comments on the Derivation of Kalman Filters", same publication as above.

6. Barnes, J.A., Jones, R.H., Tryon, P.V., Allen, D.W. "Stochastic Models for Atomic Clocks", NASA Conference Puplication 2265 (PTTI) 1982.

QUESTIONS AND ANSWERS

DR. WEISS:

UnderQ'ing is assuming a sigma too large or too small?

MR. VARNUM:

UnderQ'ing means that Q'ing is too small. I have under Q'd the phase state.

DR. WEISS:

So you are assuming that the sigma is smaller than it actually is for the variance of the process.

MR. VARNUM:

That's right. With the result that if you look in the filter, the uncertainty in that estimate is collapsing, and the gain is going down and less and less can be done to vary that state and it will tend to converge on a constant value, which you really don't want to happen for the phase if you have modeled the frequency as a constant. But underQ'd generally means you have put too little process noise in and overQ'd too much.