Proceedings of the
American Control Conference
Baltimore, Maryland • June 1994

Hybrid H_2/H_∞ Estimation for Phase-Locked Loop Filter Design *

Dan Simon & Hossny El-Sherief
TRW Systems Integration Group
Mail Stop SB2-1051
PO Box 1310
San Bernardino, CA 92402
dsimon@thor.bmd.trw.com

Abstract

A method of combining Kalman filtering and minimax filtering is proposed and demonstrated in an application to phase-locked loop design. Kalman filtering suffers from a lack of robustness to departures from the assumed noise statistics. But minimax filtering has the drawback of ignoring the engineer's (admittedly incomplete) knowledge of the noise statistics. It is shown in this paper that hybrid Kalman/minimax filtering can provide the "best of both worlds." Phase-locked loop filter design is used in this paper to demonstrate an application of hybrid estimation.

I Introduction

H_2 filtering, also known as Kalman filtering, is a well-established technology which dates back to the 1960s and has its roots in the late 1700s [14]. H_2 filtering is an estimation method which minimizes the variance of the estimation error, and assumes that the noisy inputs have known statistical properties.

Unfortunately, the assumption that the statistical properties of the noise are known limits the applicability of Kalman filters. This limitation has given rise to a recent interest in minimax estimation, also known as H_∞ filtering. The optimality measure which is used in H_∞ filtering is the magnitude of the maximum singular value of the transfer function from the noise to the estimation error. No knowledge of the noise statistics is assumed. H_∞ filtering appears to have first been introduced in 1987 [3], with roots dating back to 1981 [17].

If the H_2 approach to filtering assumes too much, the H_∞ approach assumes too little. Generally, an engineer has less knowledge about the noise than an H_2 filter requires, but more knowledge than an H_∞ filter can use. This motivates an interest in designing an estimation filter which uses the best characteristics from each type of filter. This type of cross between H_2 and H_∞ filtering can be called a *hybrid* filter.

The motivation in this paper for using a hybrid filter is digital phase-locked loop (PLL) design. PLLs are used to track the phase and frequency of the carrier component of a sinusoidal signal [9]. The development of digital PLLs began in the late 1960s, and reached a reasonable state of maturity by the early 1980s. Many different approaches have been taken in the past to PLL filter design. Perhaps the most successful approach for highly dynamic trajectories, based on comparisons in [15, p. 55], is the use of the Kalman filter.

PLLs are of particular interest to the Global Positioning System (GPS) community. GPS is a satellite-based navigation system which provides position and velocity information to any user with a GPS receiver [5]. The user position is obtained by tracking a known binary pseudo-random (PR) code transmitted by the GPS satellites, and the user velocity is obtained by tracking the sinusoidal carrier which modulates the PR code. It is clearly desirable to provide robust algorithms for the GPS receiver's PLL, or the user's velocity information may be lost.

II Optimal Filtering Fundamentals

This section reviews some of the fundamental theory of optimal filtering. First the problem is defined, and then H_2 filtering and H_∞ filtering are discussed. Finally, a method of combining these two approaches is proposed.

Consider a linear, discrete, time-invariant system given by

$$\begin{aligned} x_{k+1} &= \phi x_k + v_k \\ y_k &= H x_k + n_k. \end{aligned} \tag{1}$$

$x_k \in R^n$ is the state vector, $y_k \in R^m$ is the measurement, and v_k and n_k are noise processes. We also define an augmented noise vector as

$$\omega_k = \begin{pmatrix} v_k \\ n_k \end{pmatrix}. \tag{2}$$

It is desired to find an estimate $\hat{x}_k$ for x_k based on measurements y_i, $i \le k$. The estimator structure is assumed to be

$$\hat{x}_{k+1} = \phi \hat{x}_k + K_k (y_{k+1} - H\phi \hat{x}_k) \tag{3}$$

where K_k is a gain to be determined. Define the estimation error as

$$e_k = x_k - \hat{x}_k. \tag{4}$$

The transfer function matrix from the noise ω_k to the estimation error e_k is denoted $G_{e\omega}$. If $K_k = K$ is a constant, this transfer function is given by

$$\begin{aligned} G_{e\omega}(z) = \; & [zI - (I - KH)\phi]^{-1} \times \\ & \{(I - KH)[I_n \quad 0_{nm}] - K[0_{mn} \quad I_m]z\} \end{aligned} \tag{5}$$

where I_n is the $n \times n$ identity matrix and 0_{pq} is the $p \times q$ zero matrix.

*This work was supported by Independent Research and Development Project 93268305, sponsored by the Guidance, Systems and Controls Department of the TRW Systems Integration Group.

In H_2 filtering, also known as Kalman filtering, it is assumed that the noise processes v_k and n_k are zero-mean. The gain K_k is computed according to the formulas

$$\begin{aligned} K_k &= P_k(-)H^T(HP_k(-)H^T+R_k)^{-1} \\ P_k(+) &= (I-K_kH)P_k(-) \\ P_{k+1}(-) &= \lambda\phi P_k(+)\phi^T+Q_k \end{aligned} \qquad (6)$$

where $R_k = E(n_k n_k^T)$, $Q_k = E(v_k v_k^T)$, and λ is a forgetting factor. If $\lambda = 1$, then the Kalman filter is the affine filter which minimizes $E(e_k S_k e_k^T)$ for any positive semidefinite weight matrix sequence $\{S_k\}$. This is commonly expressed by stating that the Kalman filter is the linear minimum variance estimator [1, chap. 5]. If w_k is white and wide-sense stationary with a power spectral density of $S_w(\omega) = S$, then the Kalman filter is also the affine filter which minimizes the S-weighted 2-norm of G_{ew} [7],[8, sec. 6.5],[10, sec. 10.4]. The S-weighted 2-norm of a discrete transfer function matrix is given by

$$\begin{aligned} \|G\|_S^2 &= \frac{1}{2\pi}\int_{-\pi}^{+\pi} \text{trace}\,[G(e^{j\omega})SG^*(e^{j\omega})]\,d\omega \\ &= \frac{1}{2\pi}\int_{-\pi}^{+\pi} \sum_{i=1}^{n} \lambda_i[G(e^{j\omega})SG^*(e^{j\omega})]\,d\omega \end{aligned} \qquad (7)$$

where G^* is the Hermitian transpose of G. If the system (1) is completely observable and completely controllable, the gain K_k will reach a unique steady state, denoted by $K^{(2)}$. In order to save computational expense, the steady-state gain (which can be computed off-line) is often used in real-time systems. The resulting filter is identical to the Wiener filter [2, sec. 4.7]. If $\lambda > 1$, then greater emphasis is given to more recent data [1, sec. 6.2]. This results in greater stability and improved performance in many practical cases.

The fact that the Kalman filter is the linear minimum variance estimator is a powerful and attractive result. But several facts may indicate against the use of a Kalman filter [11]:

1. The Kalman filter minimizes $E(e_k e_k^T)$, while the user may be more interested in minimizing the worst-case error.
2. The Kalman filter assumes that $E(n_k n_k^T)$ and $E(v_k v_k^T)$ are known.
3. The Kalman filter assumes that $E(n_k)$ and $E(v_k)$ are known.

These considerations have led to the statement of the H_∞ filtering problem. Several H_∞ problem formulations have been presented in the literature [11]. The problem which is considered in this paper can be posed as follows [16]. Given the system in (1) and the estimator structure in (3), find a gain K_k such that $\|G_{ew}\|_\infty < \gamma$, where (as before) G_{ew} is the transfer function matrix from the noise w_k to the estimation error e_k, and $\|G\|_\infty$ is the magnitude of the largest singular value of G (over all frequencies).

$$\|G\|_\infty = \sup_{\omega\in[-\pi,\pi]} \lambda_{max}[G(e^{j\omega})G^*(e^{j\omega})]. \qquad (8)$$

It can be shown that if this problem has a solution for a given γ, then it can be solved by a constant gain, denoted by $K^{(\infty)}$ [16]. Amazingly enough, as $\gamma \to \infty$, the solution of the H_∞ problem is identical to the steady-state Kalman filter when $R_k = Q_k = I$. The H_∞ filtering solution for a specified γ is given by

$$\begin{aligned} K^{(\infty)} &= (I+P/\gamma^2)^{-1}PH^T \\ P^{-1} &= M^{-1} - I/\gamma^2 + H^TH \\ M &= \phi P\phi^T + I. \end{aligned} \qquad (9)$$

One method to solved these equations is given in [16]. Alternatively, (9) can be solved iteratively. Note from the problem statement that as γ gets larger, the problem is "easier" to solve. If γ is too small, the problem will not have a solution, and X_1 will be singular.

II-A Hybrid H_2/H_∞ Filtering

H_2 and H_∞ filters both have pros and cons. The H_2 filter assumes that the noise statistics are known. The H_∞ filter does not make this assumption, but further assumes that absolutely nothing is known about the noise characteristics.

Suppose that although the noise statistics are not perfectly known, the user does have a rough idea of these statistics. Also suppose that a user desires to minimise some combination of the H_2 and H_∞ objective functions. What could be done? Perhaps a hybrid H_2/H_∞ filter could be used.

Several approaches to hybrid filtering have been proposed in the literature [7]. In this paper we propose a heuristic hybrid filtering approach. We will simply use a weighted combination of the steady-state H_2 and H_∞ gains in the estimator. That is,

$$K = d\,K^{(2)} + (1-d)K^{(\infty)} \qquad (10)$$

where $[0,1] \ni d =$ the relative weight given to H_2 performance. $K^{(2)}$ is given by the steady-state solution of (6) and $K^{(\infty)}$ is the solution of (9). The key design parameter in the hybrid filter is the weight d. This weight must be chosen so as to ensure stability. A convex combination of two stable estimators is not necessarily stable, as is shown in Section IV. So the first criterion for the choice of d is stability. The second criterion is the relative weight given by the user to H_2 performance versus H_∞ performance. This relative weight can be determined on the basis of the engineer's confidence in the *a priori* statistics.

III Application to Phase-Locked Loop Design

Consider the problem of tracking a sinusoidal signal with an unknown, time-varying phase $\theta(t)$:

$$s(t) = A\cos\theta(t). \qquad (11)$$

This signal is corrupted by noise. The device used to track such a signal is called a phase-locked loop. PLLs are of particular interest in Global Positioning System receivers. A GPS satellite transmits a sinusoidal signal modulated by a known pseudo-random binary code. After the PR code is removed from the signal, the receiver has access to the sinusoid. Since the sinusoid is transmitted at a known frequency, the frequency which the receiver tracks can be used to compute the doppler between the user and the satellite. The satellite orbit is known fairly accurately, so the doppler frequency can be used to obtain the user's velocity. A GPS receiver can therefore be used as a navigational instrument in place of more expensive and complex inertial instruments. The receiver architecture considered in this paper is shown in Figure 1.

Note from Figure 1 that the output of the arctan phase discriminator is modulo 2π. That is, the phase discriminator does

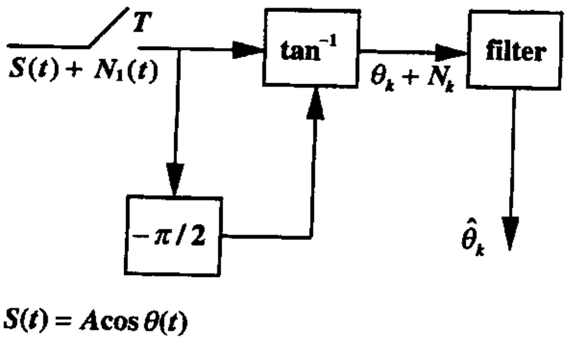

Figure 1: Phase-locked loop architecture

not know the difference between θ radians and $\theta + 2\pi$ radians. If the phase estimation error suddenly goes from zero to some multiple of 2π, it is said that a *cycle slip* has occurred. So it is more important in a PLL to prevent cycle slips than it is to maintain a small phase error. If the PLL maintains lock on the phase, the PLL contribution to a GPS receiver's velocity error is small compared to other sources of velocity error [12, 13]. For instance, the velocity error due to a typical 4° RMS PLL tracking error may be on the order of 0.01 feet/second. But the velocity error due to all other sources may be on the order of 0.10 feet/second. If a cycle slip occurs, then the velocity error due to the PLL tracking error momentarily jumps to 0.90 feet/second. So undetected cycle slips can be catastrophic. In some cases, the noise is so high or the phase dynamics are so severe that the estimation error begins growing without bound. In this case it is said that *loss of lock* has occurred, and the user loses all velocity information from the GPS receiver. Therefore, for a GPS receiver, it is primarily loss of lock and secondarily cycle slips which are of greatest concern (rather than phase error).

Optimal filtering can be used in PLL design by a method similar to that used in [15]. We create a state vector from successive derivatives of the incoming phase.

$$x_k = \begin{pmatrix} \theta_k & \omega_k & \alpha_k & \beta_k \end{pmatrix}^T \tag{12}$$

where

$$\begin{aligned} \theta_k &= \theta(t_k) \qquad \omega_k = \theta'(t_k) \\ \alpha_k &= \theta''(t_k) \qquad \beta_k = \theta'''(t_k). \end{aligned} \tag{13}$$

We therefore obtain the state transition equations [15]

$$\beta_{k+1} = \beta_k + \int_{kT}^{(k+1)T} \theta^{(4)}(t)\,dt \tag{14}$$

$$\alpha_{k+1} \approx \alpha_k + T\beta_k + \int_{kT}^{(k+1)T} t\theta^{(4)}(t)\,dt \tag{15}$$

$$\omega_{k+1} \approx \omega_k + T\alpha_k + T^2\beta_k/2 + \int_{kT}^{(k+1)T} \frac{t^2}{2}\theta^{(4)}(t)\,dt \tag{16}$$

$$\theta_{k+1} \approx \theta_k + T\omega_k + T^2\alpha_k/2 + T^3\beta_k/6 + \int_{kT}^{(k+1)T} \frac{t^3}{6}\theta^{(4)}(t)\,dt \tag{17}$$

where the approximations are valid for a small sample period T. This gives rise to the system description

$$\begin{aligned} x_{k+1} &= \theta_k = \phi x_k + v_k \\ y_k &= H x_k + n_k \end{aligned} \tag{18}$$

where the system matrices ϕ and H are apparent from (12) – (17). As a further approximation, we can model the fourth derivative of the phase as a white noise process with variance N.

$$E[\theta^{(4)}(t)\theta^{(4)}(\tau)] = N\delta(t-\tau). \tag{19}$$

Similarly, we can model the continuous-time phase measurement noise as a white noise process with variance N_0, so the variance of the sampled version is N_0/T.

$$\begin{aligned} E[n(t)n(\tau)] &= N_0\delta(t-\tau) \\ E[n(t_k)n(t_j)] &= \frac{N_0}{T}\delta_{k-j}. \end{aligned} \tag{20}$$

These considerations lead us to the assumed noise statistics

$$E(n_k n_j) = R\delta_{k-j} = (N_0/T)\delta_{k-j} \tag{21}$$

$$\begin{aligned} E(v_k v_j^T) &= Q\delta_{k-j} \\ &= NT \begin{pmatrix} T^6/252 & T^5/72 & T^4/30 & T^3/24 \\ T^5/72 & T^4/20 & T^3/8 & T^2/6 \\ T^4/30 & T^3/8 & T^2/3 & T/2 \\ T^3/24 & T^2/6 & T/2 & 1 \end{pmatrix} \delta_{k-j} \end{aligned} \tag{22}$$

The filter structure used to obtain a state estimate is given by

$$\hat{x}_{k+1} = \phi\hat{x}_k + K(y_{k+1} - H\phi\hat{x}_k). \tag{23}$$

A constant gain K will be used due to real-time computational constraints. Note that the estimate $\hat{\theta}_k$ (the first component of $\hat{x}_k$) is actually an estimate of the phase modulo 2π. This estimate can be placed in the proper phase cycle by using the frequency estimate $\hat{\omega}_k$.

If the noise processes n_k and v_k are zero-mean, Q_k and R_k are known, and the user wants to minimize the variance of the phase estimation error, then the steady-state Kalman filter gain can be used. If, on the other hand, the noise processes may or may not be zero-mean, Q_k and R_k are *not* known, and the user wants to minimize the worst-case effect of the noise on the phase estimation error, then the H_∞ filter gain can be used. If the user has some idea of the noise statistics (but doesn't know them exactly) and wants to minimize some combination of the H_2 and H_∞ objective functions, then the hybrid H_2/H_∞ filter discussed in Section II-A can be used.

IV Simulation Results

The hybrid H_2/H_∞ filter discussed in this paper was simulated for a GPS receiver used for missile navigation. The behavior of the H_2/H_∞ PLL was investigated by examining its ability to track the phase between the missile and one GPS satellite for the first 60 seconds of boost (i.e., during Stage I burn). The filter rate was fixed at 50 Hz. The satellite-to-missile range acceleration is depicted in Figure 2. The relationship between the phase θ and the range ρ is given by

$$\rho(t) = \frac{c\theta(t)}{2\pi f} \tag{24}$$

where c is the speed of light and f is the frequency of the transmitted sinusoid. We concentrate in this paper on tracking

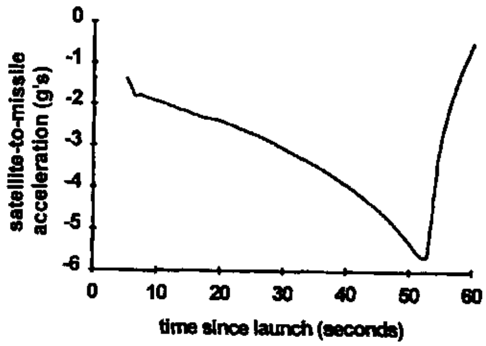

Figure 2: Missile-to-satellite range acceleration

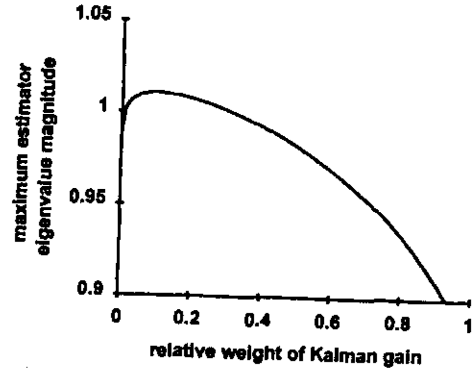

Figure 3: The stability of the hybrid filter

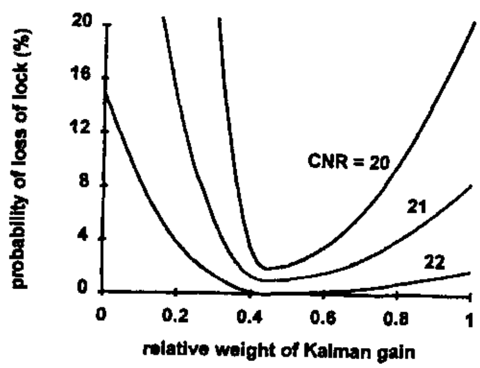

Figure 4: Probability of loss of lock as a function of relative weight of Kalman gain

the GPS L1 carrier at a frequency of 1.575 GHz. We assume without loss of generality (see Figure 1) that the magnitude of the carrier is unity.

We found, in agreement with [15], that $\lambda \approx 1.055$ resulted in the best performance for the Kalman filter. Typical carrier-to-noise ratios (CNRs) for GPS are around 30 to 40 dB-Hz [4]. But if atmospheric conditions are severe or jamming is present, the CNR could drop into the 20's and biases could be introduced into the measurement noise. The CNR is related to the variance of the measurement noise (R) by

$$\text{CNR} = \frac{1}{2TR} \tag{25}$$

where $T = 0.02$ seconds is the filter rate [6, p. 282]. Our interest in this paper is tracking the GPS L1 carrier in spite of the fact that the CNR is significantly different than expected and the noise is not zero-mean. It is assumed in this section that there is a constant, unknown phase measurement bias of 1 radian. So the *true* measurement equation is

$$y_k = \theta_k + n_k + 1 \tag{26}$$

where n_k is zero-mean noise, but the filters are designed according to the incorrect measurement equation

$$y_k = \theta_k + n_k. \tag{27}$$

The simulated noise n_k was generated with a Laplacian (exponential) density, which has heavier "tails" than a Gaussian density.

The steady-state Kalman filter gain for CNR = 20 dB-Hz and $\lambda = 1.055$ was found to be

$$K^{(2)} = [\ 0.4926 \quad 7.6403 \quad 65.9854 \quad 274.1215\ ]^T. \tag{28}$$

The H_∞ estimation problem was found numerically to be solvable for $\gamma > 1$, so $\gamma = 1.01$ was used in (9) to compute the H_∞ gain

$$K^{(\infty)} = [\ 0.9127 \quad 1.8732 \quad 1.7398 \quad 0.6833\ ]^T. \tag{29}$$

As discussed in Section II-A, the user needs to choose a value of the Kalman weight gain d such that the hybrid filter is stable. Figure 3 shows the magnitude of the largest eigenvalue of the hybrid estimator as a function of d. It is seen that the estimator is unstable for $0.01 < d < 0.31$. This shows that d must be chosen greater than 0.31 for satisfactory estimator performance. How much greater? For the pure H_∞ filter, $|\lambda_{max}| = 0.988$. We see from Figure 3 that as d increases, $|\lambda_{max}|$ also increases at first, then begins decreasing. $|\lambda_{max}|$ drops back down to 0.988 at $d \approx 0.45$. So in this application, a rule of thumb for hybrid H_2/H_∞ filter design is to choose $d \geq 0.45$. This ensures that the hybrid filter is at least as stable as the pure H_∞ filter.

Recall that our primary interest is in maintaining lock during the mission. With this in mind, the probability of loss of lock was obtained experimentally for various values of the Kalman gain weight d in (10). Recall further that $d = 0$ corresponds to a pure H_∞ filter, while $d = 1$ corresponds to a pure H_2 filter. Probability of loss of lock was obtained by conducting 100 Monte Carlo samples for each data point. This probability is shown in Figure 4 as a function of d for three values of CNR. It is seen that the use of hybrid H_2/H_∞ ($d < 1$) filtering results in a noticable improvement in phase lock over pure H_2 or pure H_∞ filtering. Furthermore, the advantage becomes more significant as the CNR decreases. For example, a pure Kalman filter with a CNR of 20 dB-Hz has a 20% chance of losing lock. But a hybrid H_2/H_∞ filter with a weight d around 0.4 or 0.5 has only a 2% chance of losing lock.

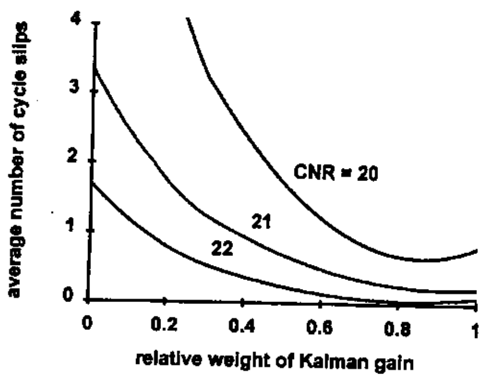

Figure 5: Average number of cycle slips when the PLL maintains lock

When the Kalman filter *does* maintain lock, it performs better than the hybrid filter. This is seen in Figure 5, which shows the average number of cycle slips as a function of d for various values of CNR. The numbers in Figure 5 are derived only from those Monte Carlo samples which did not lose lock. So if the Kalman filter does not lose lock, chances are it will have fewer cycle slips than the hybrid filter. But the advantage is not significant. For example, at a CNR of 20 dB-Hz, the Kalman filter (if it maintains lock) slips an average of one cycle, while a hybrid filter with d around 0.4 or 0.5 slips an average of two cycles.

V Conclusion

A hybrid H_2/H_∞ filtering approach has been proposed and applied to phase-locked loop design. This hybrid approach not only takes advantage of the noise statistics knowledge which is inherent in H_2 filter design, but also takes advantage of the robustness of H_∞ filtering. It is seen from the simulation data that, in general, the hybrid filter provides a large advantage over the pure Kalman filter and the pure H_∞ filter. This advantage is particularly noticable at low CNRs, even when the Kalman filter designer has perfect knowledge of the true CNR. The Kalman filter is not robust to departures from the assumed noise statistics. But the H_∞ filter does not take advantage of the designer's (albeit incomplete) knowledge of the noise properties. hybrid filtering is an approach which combines the best of both worlds – or at least avoids the worst of both worlds. It is thus recommended that the hybrid H_2/H_∞ filter proposed in this paper be given serious consideration for PLL design in particular, and for state estimation in general.

References

[1] B. Anderson and J. Moore, *Optimal Filtering*, Englewood Cliffs, NJ: 1979.

[2] A. Gelb (editor), *Applied Optimal Estimation*, Cambridge, MA: The MIT Press, 1984.

[3] M. Grimble, "H_∞ Design of Optimal Linear Filters," *International Symposium on Mathematical Theory of Networks and Systems*, Phoenix, AZ, 1987.

[4] W. Hurd, J. Statman, and V. Vilnrotter, "High Dynamic GPS Receiver Using Maximum Likelihood Estimation and Frequency Tracking," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 23, pp. 425-436, July 1987.

[5] P. Janicsek and S. Gilbert (editors), *Global Positioning System Papers, Vols. 1-3*, Washington, DC: The Institute of Navigation, 1980, 1984, 1986.

[6] M. Jeruchim, P. Balaban, and K. Shanmugan, *Simulation of Communication Systems*, New York: Plenum Press, 1992.

[7] P. Khargonekar and M. Rotea, "Mixed H_2/H_∞ Filtering," *IEEE Conference on Decision and Control*, Tucson, AZ, pp. 2299-2304, Dec. 1992.

[8] H. Kwakernaak and R. Sivan, *Linear Optimal Control Systems*, New York: John Wiley & Sons, 1972.

[9] W. Lindsey and C. Chie (editors), *Phase-Locked Loops*, New York: IEEE Press, 1986.

[10] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, New York: McGraw-Hill Book Company, 1984.

[11] U. Shaked and Y. Theodor, "H_∞-Optimal Estimation: A Tutorial," *IEEE Conference on Decision and Control*, Tucson, AZ, pp. 2278-2286, Dec. 1992.

[12] D. Simon and H. El-Sherief, "GPS Receiver Design for Integrated Navigation Systems," *Institute of Navigation 49th Annual Meeting*, Cambridge, MA, pp. 437-445, June 1993.

[13] D. Simon and H. El-Sherief, "Design of Global Positioning System Receivers for Integrated Inertial Navigation Systems," *IEEE Conference on Decision and Control*, San Antonio, TX, pp. 1476-1477, 1993.

[14] H. Sorenson (editor), *Kalman Filtering: Theory and Applications*, New York: IEEE Press, 1985.

[15] V. Vilnrotter, S. Hinedi, and R. Kumar, "A Comparision of Frequency Estimation Techniques for High-Dynamic Trajectories," *JPL Publication 88-21*, Pasadena, CA: Jet Propulsion Laboratory, Sep. 1988.

[16] I. Yaesh and U. Shaked, "A Transfer Function Approach to the Problems of Discrete-Time Systems: H_∞-Optimal Linear Control and Filtering," *IEEE Transactions on Automatic Control*, vol. 36, pp. 1264-1271, Nov. 1991.

[17] G. Zames, "Feedback Optimal Sensitivity: Model Preference Transformation, Multiplicative Seminorms and Approximate Inverses," *IEEE Transactions on Automatic Control*, vol. 26, pp. 301-320, 1981.