AIAA 80-0171R

Use of Cramer-Rao Bounds on Flight Data with Colored Residuals

Richard E. Maine* and Kenneth W. Iliff†
NASA Dryden Flight Research Center, Edwards, Calif.

This paper discusses the use of the Cramer-Rao bound as a means of assessing the accuracy of maximum likelihood parameter estimates obtained from dynamic flight data. Emphasizing practical considerations such as modeling error, the Cramer-Rao bound is evaluated with real and simulated data. Improved computations of the bound to correct large discrepancies caused by colored noise and modeling error are presented. This corrected Cramer-Rao bound is the best available analytical predictor of accuracy.

Nomenclature

a_y	= lateral acceleration, g
B	= band limit, Hz
C_{l_β}	= coefficient of dihedral effect
$E\{\ \}$	= expected value
FF^T	= state noise power spectral density matrix
$f(\ )$	= general function
GG^T	= measurement noise covariance matrix
$g(\ )$	= general function
J	= cost functional
M	= Fisher information matrix
N	= number of time points
n	= state noise vector
p	= roll rate, deg/s
$p(\ )$	= probability density
R	= prediction error covariance matrix
r	= yaw rate, deg/s
t	= time, s
u	= control vector
var	= variance
x	= state vector
z	= observation vector
$\hat{z}_\xi$	= prediction of z based on ξ
α	= angle of attack, deg
β	= angle of sideslip, deg
Δt	= sampling interval, s
δ_a	= aileron deflection, deg
η	= measurement noise vector
ξ	= vector of unknowns
ϕ	= bank angle, deg
∇_ξ	= gradient with respect to ξ

Superscripts

$(\hat{\ })$	= estimate
$(\ )^T$	= transpose

Introduction

PARAMETER estimation is widely used to obtain the stability and control derivatives of aircraft from dynamic flight maneuvers.[1] The estimates are used during flight test programs in order to safely expand the flight envelope, to improve the fidelity of simulators, and to analyze the effects of control system changes. Comparisons are made with theoretical and wind tunnel predictions to validate and improve the prediction and estimation techniques. Adaptive and learning control systems can also make use of the estimates.

The derivatives obtained from flight data are, by their nature, only estimates rather than exact values. In order to make effective use of these estimates, it is necessary to have some gage of their reliability, be it a statistical measure or an intuitive guess. If accurate estimates cannot be distinguished from worthless estimates, then to be safe, all estimates must be treated as worthless. Therefore, measures of the estimate accuracy are as valuable as the estimates themselves.

Because of the numerous qualifications about modeling error and noise statistics, it is necessray to validate the theoretical accuracy measures before much confidence can be placed in them. The Cramer-Rao bounds[2-6] have been advanced as the best of the theoretical measures of accuracy. Reference 2 treats this subject in detail. Comparison of the Cramer-Rao bounds and the sample standard deviation obtained from the data scatter gives a good indication of the adequacy of the assumptions made in the theoretical development.

It has long been known that such comparison shows significant discrepancies in actual flight data. This paper will examine these discrepancies and explain their cause. An approximate correction for these effects will then be suggested and evaluated. The examples center on maximum likelihood estimates of aircraft stability and control derivatives, but it should be plain that many of the principles have much wider application.

Maximum Likelihood Estimation

Maximum likelihood estimation is by far the most commonly used technique for estimating parameters from dynamic flight data. The theory of maximum likelihood estimation is extensive; thus, a great deal can be inferred about the accuracy of maximum likelihood estimates. The emphasis of this paper will be on maximum likelihood estimates; a brief review of the most basic results on maximum likelihood estimation is therefore presented. No attempt is made here to derive these results or discuss them in detail.[3,4]

The derivations in this paper all use continuous-time models with discrete observations. The basic model form is

$$\begin{aligned} \dot{x}(t) &= f(x(t),u(t),t,\xi) + Fn(t) \\ z(t_i) &= g(x(t_i),u(t_i),t_i,\xi) + G\eta(t) \end{aligned} \tag{1}$$

and the unknown vector ξ is to be estimated. The F and G matrices can also be functions of ξ, but they usually receive special treatment. The maximum likelihood estimates can be

Received Dec. 14, 1979; presented as Paper 80-0171 at the AIAA 18th Aerospace Sciences Meeting, Pasadena, Calif., Jan. 14-16, 1980; revision received June 23, 1980.

*Aerospace Engineer. Member AIAA.

†Aerospace Engineer. Associate Fellow AIAA.

obtained by minimizing the following cost functional:

$$J(\xi)=\frac{1}{2}\sum_{i=1}^{N}[\hat{z}_\xi(t_i)-z(t_i)]^T R^{-1}[\hat{z}_\xi(t_i)-z(t_i)] \quad (2)$$

where $\hat{z}_\xi$ is the Kalman filter predicted z based on the postulated ξ, and R is the prediction error covariance matrix. Many algorithms could be used to minimize J. Reference 6 describes the Newton-Balakrishnan algorithm, which has been found to work well.

Two special cases are often used. If F is 0, then $\hat{z}$ is obtained simply by integration and R is GG^T; this is referred to as an output error method. If G is 0, and g can be inverted to obtain x from z, the estimator is essentially an equation error method. The general estimator with neither F nor G zero is referred to as a prediction error method. Although most of the examples in this paper use an output error method, the principles are applicable to the equation error or prediction error forms.

One of the statistical properties of maximum likelihood estimates is of significant importance to this paper. The estimates are asymptotically efficient—that is, their variance approaches the limit given by the Cramer-Rao inequality as the amount of data used increases. If the data length is long enough, the limit will be a good approximation to the variance; "long enough" can usually be equated with a few periods of the lowest system natural frequency.

The Cramer-Rao inequality[2-6] is a measure of the information content of a maneuver. It gives a lower bound on the variance of the estimates; no estimator can do better than this bound. For unbiased estimates, the inequality is

$$\mathrm{var}(\hat{\xi}) \geq M^{-1} \quad (3)$$

where M is the Fisher information matrix.

$$M=\mathrm{E}\{[\nabla_\xi \log P(Z|\xi)][\nabla_\xi \log P(Z|\xi)]^T\} \quad (4)$$

This matrix inequality means that the left-hand side of Eq. (3) minus the right-hand side is positive semidefinite.

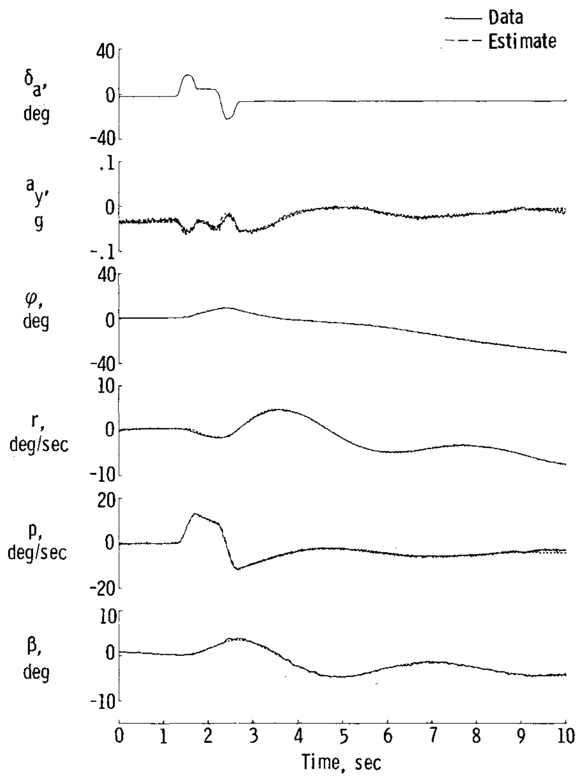

Fig. 1 Time-history match of flight data.

For the system of Eq. 1, the information matrix can be closely approximated by

$$M \cong \sum_{i=1}^{N} \nabla_\xi \hat{z}_\xi^T(t_i) R^{-1} \nabla_\xi^T \hat{z}_\xi(t_i) \quad (5)$$

This matrix is recognized as the dominant term of the Hessian matrix $\nabla_\xi^2 J(\xi)$. Its computation or a suitable approximation is required by most algorithms for minimizing the cost functional of Eq. (2). The computation of Eq. (5) is exactly that used by the Newton-Balakrishnan algorithm. When F is 0 and G is known, it is an exact expression for the information matrix. Otherwise, it is a good approximation.

Since maximum likelihood estimates are asymtotically unbiased and efficient, M^{-1} is a good approximation to the variance of $\hat{\xi}$ for long maneuvers.

Discrepancy in the Cramer-Rao Bound

Data from a PA-30 aircraft was chosen to evaluate the Cramer-Rao bounds. Eighteen maneuvers were obtained from this vehicle. Each maneuver consisted of an aileron input initiated from steady flight. The derivatives were estimated by the MMLE3 program[6] using an output error maximum likelihood method (measurement noise only). A typical time-history match is shown in Fig. 1.

Figure 2 presents the estimates of C_{l_β} and the Cramer-Rao bounds obtained from these maneuvers. The vertical scales in these plots are exaggerated in order to show the scatter, and the α scale is exaggerated in order to separate the maneuvers. No significant differences in the derivatives are expected over this small angle-of-attack range; the 18 maneuvers can be regarded as being at essentially the same flight condition.

The Cramer-Rao bounds on this plot are so small that they are difficult to see—they are roughly the same size as the symbols. The data scatter is much larger than indicated by the Cramer-Rao bounds. Quantitatively the sample standard deviation is about nine times the average Cramer-Rao bound. This discrepancy between the data scatter and the Cramer-Rao bound has long been known. It has become a common practice to multiply the Cramer-Rao bounds by a "fudge factor" of 5 to 10 (Ref. 7). The resulting values proved useful for evaluating the accuracy of estimates, but the necessity of the unexplained fudge factor detracted from the confidence. Because of this problem with the Cramer-Rao bounds, several reports[8-10] used the estimated correlations as primary indicators of accuracy in spite of the problems with the correlations mentioned in Ref. 2. The estimated correlations are based on the same theoretical foundation as the Cramer-Rao bounds and thus should be equally suspect if errors are known to exist. Other reports ignored the fudge factor and quoted overly optimistic values of the accuracy.[11]

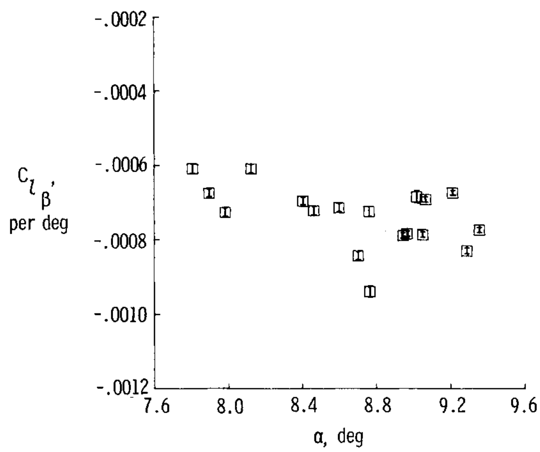

Fig. 2 Estimates from flight data.

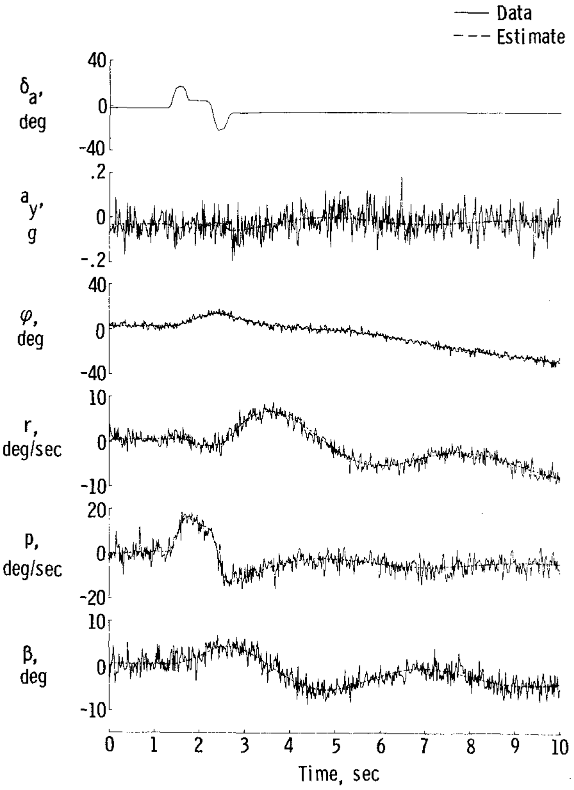

Fig. 3 Time-history match of simulated data.

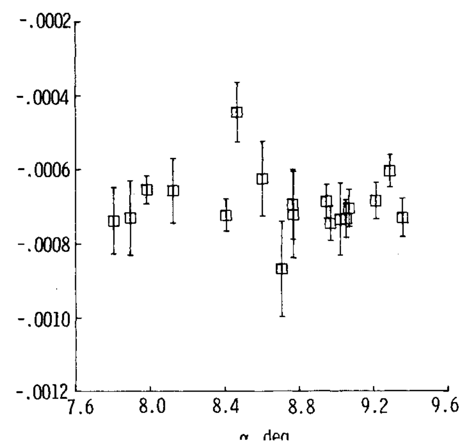

Fig. 4 Estimates from simulated data.

Discrepancies larger than the quoted accuracy were then attributed to various effects without sufficient data points to establish whether the observed differences were significant or lay within the scatter band.

The evaluation of accuracy measures on actual flight data is complicated by the impossibility of establishing true values for comparison and by the inevitable presence of unmodeled effects. Although tests on actual flight data are necessary for final validation, simulated data provide a more controlled environment which may aid preliminary work. The above experiment was therefore repeated with simulated data.

In order to mimic the flight data experiment as closely as feasible, the control inputs measured from the flight data were used to create the simulated data. Eighteen simulated maneuvers were created using the same flight conditions as the 18 actual manuevers. The same model was used for the simulation as for the estimation. The same true values of the nondimensional derivatives were used for all 18 maneuvers. A pseudorandom noise generator was used to add simulated white Gaussian measurement noise to the responses. The measurement noise power for each signal was proportional to the average residual power observed on the flight data for that signal. The proportionality constant was adjusted to obtain the same magnitude of scatter in the estimates from the simulated data as was observed in the flight data. The derivatives and Cramer-Rao bounds were estimated from this simulated data using the same program as for the flight data. A typical time-history match is shown in Fig. 3.

The estimates of C_{l_β} from the simulated data are shown in Fig. 4, plotted to the same scale as Fig. 2. The true value is -0.0007.

The scatter on this plot is about the same as for the flight data because the simulated noise power was adjusted to achieve this scatter. The Cramer-Rao bounds differ drastically on the flight and simulated data. The Cramer-Rao bounds on the simulated data are about 10 times that of the flight data and agree well with the observed scatter. For the 11 derivatives estimated (not including bias terms), the ratios of the sample standard derivations to the Cramer-Rao bounds range from 0.68 to 1.32. This is excellent agreement for a sample of only 18 maneuvers.

The Cramer-Rao bound thus agrees very well with the scatter on simulated data, but disagrees drastically on flight data. The excellent performance of the bound on simulated data destroys many of the suggestions that could be advanced to explain the discrepancy. First, and most obvious, is the possibility of an error in the formulation or computer programming. Such an error would have shown up in the simulated as well as the flight results, since the same computer program was used. Second is the fact that the Cramer-Rao bound is only a lower bound. The maximum likelihood estimator is proven to be asymptotically efficient; that is, the Cramer-Rao inequality approaches an equality as time approaches infinity. For finite time, the equality does not hold. Intuition, which can be backed up by analysis in this case, suggests that a few periods of the natural frequency should be enough that the asymptotic result is closely approached. The fact that the simulated data agreed so well with the scatter verifies that the time was long enough to make the equality a very good approximation.

The preceding attempts to explain the discrepancy between the Cramer-Rao bounds and the scatter implicitly assumed that the scatter is a reasonable measure of the accuracy. For the simulated data, where the true values are known to be constant, this statement is almost a definition of accuracy. Other possibilities must be considered for the flight data. Recall that the simulated data noise level was chosen to make the scatter match that of the flight data, and the resulting Cramer-Rao bounds of the simulated data were larger than those of the flight data. As is evident by comparing Figs. 1 and 3, the noise power in this simulated data is much larger than that of the flight data residuals. The difference in the Cramer-Rao bounds arises directly from the difference in the noise power. If the simulated data noise power were lowered to be the same as that of the flight data residuals, the simulated and flight Cramer-Rao bounds would be the same. Of course the scatter of the simulated data would be much less than the flight data.

This suggests the possibility that, instead of the Cramer-Rao bounds being too small, the scatter is too large in the flight data to properly represent the accuracy. It might be true, for instance, that the individual estimates are as accurate as indicated by the Cramer-Rao bounds and that the scatter reflects actual changes in the coefficients. In principal, this would explain the discrepancy. However, there is no physical reason to suspect such large variations in the aerodynamic derivatives at essentially the same flight condition. Furthermore, no ascertainable pattern can be detected in the scatter that relates to any flight condition parameter. Although only a single example is shown here, the same discrepancy is noted on every class of vehicle tested, including small, general-aviation aircraft,[7] airliners,[12] military

fighters,[13] large supersonic aircraft,[14] and unconventional vehicles.[15,16] The universality of the discrepancy argues strongly against the possibility of apparently random changes in the actual derivatives.

A related possibility is that the scatter in the estimates results from unmodeled errors that would not be reflected in the Cramer-Rao bounds. An obvious example of such a problem would be an error in the measurement of the flight condition. For instance, if the dynamic pressure measurement were inaccurate, the Cramer-Rao bound and the estimates of the dimensional derivatives would not be affected. The nondimensional derivatives, however, would have larger errors than otherwise predicted. The occurrence of the same discrepancy on many different aircraft and data systems argues strongly against this case. Most of the data systems used are believed to be accurate enough to eliminate problems of errors in the flight condition.

None of the suggestions advanced above has proven a satisfactory explanation of the discrepancy observed in the flight data. For several years the necessity for the fudge factor on flight data was left essentially without explanation. It was argued that modeling errors existed which invalidated the Cramer-Rao bound. Such modeling errors, of course, were not present in the simulated data. Although this argument is virtually irrefutable, it does little to explain the problem. The subject of what types of modeling error might exist that would have such effects was not addressed. This argument, amounting to a dismissal of the problem, does not give any basis for confidence in the use of the fudge factor. Although some authors found the values to be of empirical use when the fudge factor was applied, others rejected the Cramer-Rao bound as invalid.

Explanation of the Discrepancy

In Ref. 17, the authors advanced the first satisfactory explanation of the discrepancy in the Cramer-Rao bound. The discrepancy was traced to the theoretical assumptions about the independence of the noise samples.

The exact discrete-time theory of estimation in the presence of colored noise is trivial when the spectral shape of the noise is known. The application is also easy in the frequency domain, but considerations such as nonlinearities and time variation severely limit the application of frequency domain estimation. The time domain application of the exact theory of estimation with colored noise is overly cumbersome.

The effort is not justified by the small benefits expected. Furthermore, the noise spectral characteristics are seldom precisely known, and incorrect specification could actually make the estimates worse. Estimation of the spectral characteristics is possible in principle, but adds further unacceptable computational complications. The theory of estimation with colored noise has therefore been largely relegated to textbook examples.

Abandoning the exact approach, we recognize that approximating the effects of colored noise can still provide more useful results than ignoring the question. For a first approximation, assume that the noise is band limited white with band limit B. Then, rather than derive a maximum likelihood estimator for this system, analyze the performance of the estimator based on white noise when the actual noise is colored in this manner. The results of this analysis[18] agree well with intuitive expectations. As long as the noise bandwidth is much larger than the system bandwidth, there is negligible effect on the estimates. Stated loosely, the estimation errors are caused by the noise near the natural system frequencies; the estimator will always mistake some percentage of this noise as actual system response and vice versa. Noise far above the system bandwidth is readily identified as noise rather than as system response. A good estimator should be little influenced by such high-frequency noise.

This result is readily generalized. It is obvious that the exact shape of the noise rolloff was of little consequence in the preceding analysis. The analysis could be repeated with different types of rolloff characteristics and the same results would be obtained. In short, the high-frequency characteristics of the noise do not materially affect the estimates.

This conclusion is a welcome validation of the practice of using the maximum likelihood estimator based on independent noise samples, even though the actual residuals are known to be significantly correlated. The quantitative interpretation of "high" frequency is somewhat difficult, and skepticism is prudent when the noise bandwidth nears the system bandwidth, as often occurs. Nonetheless, this theory provides a much stronger base than totally ignoring the question.

Since the estimates are essentially unaffected by high-frequency noise, it immediately follows that all of the statistical characteristics of the estimates are equally unaffected. Thus, in particular, the same expression for the Cramer-Rao bound should still be valid. This would seem to refute the thesis that the discrepancy in the Cramer-Rao bound is related to the noise spectrum. In fact, the relationship is so elementary that it has been overlooked.

In stating that the estimates are essentially unaffected by high-frequency noise, we are comparing noise spectra that are the same at low frequencies. The high-frequency spectra, and thus the total noise power, will differ. The low-frequency spectral density, rather than the total power, is the important statistic. All of the programs in use are written in terms of the total noise power (or, equivalently, the noise variance). Programs based on the continuous-time theory use the spectral density, but the spectral density is estimated in practice by dividing the discrete total power estimate by the Nyquist frequency.

Let us consider the effects of using the total power instead of the spectral density. Imagine a system with total noise power R. If the noise samples are independent, the classic analysis is valid and the Cramer-Rao bounds should be correct. The noise power spectral density of this system is $2R\Delta t$, since the noise spectrum is flat out to the Nyquist frequency $1/2\Delta t$. Now, imagine a second system with the same total power but with a one-sided noise bandwidth of B. Figure 5 shows the noise spectra of these two systems.

The Cramer-Rao inequality for both systems, expressed in the usual way, is

$$\operatorname{var}(\hat{\xi}) \geq \left\{ \sum_{i=1}^{N} [\nabla_{\xi}\hat{z}(t_i)^T] R^{-1} [\nabla_{\xi}\hat{z}(t_i)^T]^T \right\}^{-1} \tag{6}$$

where $\nabla_{\xi}\hat{z}^T$ does not depend on the noise statistics. Since the total power R is the same for both systems, the same values are computed for Cramer-Rao bounds. This is the computation used by current programs. Expressed in terms of the power spectral density GG^T, the Cramer-Rao inequality

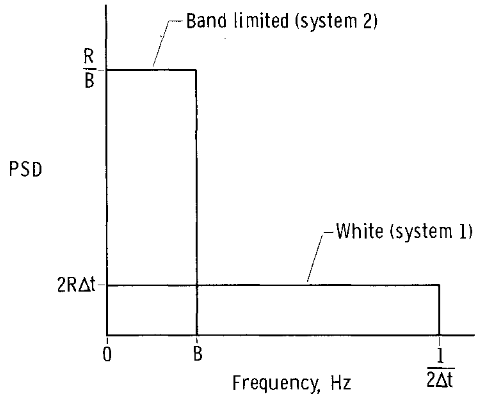

Fig. 5 Assumed noise spectra.

should be

$$\mathrm{var}(\hat{\xi}) \geq \left\{ \sum_{i=1}^{N} [\nabla_\xi \hat{z}(t_i)^T] \left(\frac{1}{2\Delta t} GG^T \right)^{-1} [\nabla_\xi \hat{z}(t_i)^T]^T \right\}^{-1} \tag{7}$$

For the first system, $(1/2\Delta t)GG^T = R$, so this computation is the same as the previous one. For the second system, however, $(1/2\Delta t)GG^T = (1/2B\Delta t)R$, thus

$$\mathrm{var}(\hat{\xi}) \geq \frac{1}{2B\Delta t} \left\{ \sum_{i=1}^{N} [\nabla_\xi \hat{z}(t_i)^T] R^{-1} [\nabla_\xi \hat{z}(t_i)^T]^T \right\}^{-1} \tag{8}$$

The Cramer-Rao bound computation based on the total power was, therefore, too small by a factor of $1/2B\Delta t$ on the variance (or $\sqrt{1/2B\Delta t}$ on the standard deviation). The Cramer-Rao bound computation based on the power is assuming that the power is evenly spread over the Nyquist range as in the first case. Thus, for Eq. (6), B is assumed to be $1/2\Delta t$. If this assumption is incorrect, the resulting variance computation will be proportionally incorrect.

These results explain excellently the discrepancies previously observed. The noise samples on the simulated data were independent, and thus the corresponding Cramer-Rao bound computations were valid. Modern flight test instrumentation is accurate enough that the largest component of the residual error in flight data is from modeling error rather than true measurement error. The "measurement noise" statistics must include all such unmodeled effects. (The philosophical question of precisely what can be included in "measurement noise" is not addressed in detail in this paper.[5] Suffice it to say that since our theory contains only the linear system model and measurement noise, anything that is not included in the system model must be part of the "measurement noise." Otherwise our theory is denying its existence; such solipsism seems unwise.) The "measurement noise" in real flight test data therefore tends to be quite colored. Figure 6 shows power spectral density plots of the residuals from the flight test data used in Fig. 2. A precise breakpoint is not obvious, but a value in the neighborhood of 1 Hz seems reasonable. The Nyquist frequency for this 50 sample/s data is 25 Hz, so the computed Cramer-Rao bounds should be increased by a factor of about $\sqrt{25/1} = 5$. The sample standard deviations were about nine times the Cramer-Rao bounds before this correction. The agreement is now quite reasonable considering the vaguely defined value of the break frequency. The noise break frequency is close enough to the system frequencies that the approximations in the theory are subject to question, but the experimental results hold up well.

As a verification of the theoretical results of this section, a new set of simulated data was created. The noise for this data was created by passing the pseudorandom independent noise through a fifth-order Chebychev filter with a break frequency of 1 Hz. This filter has a sharp break at 1 Hz as shown by Fig. 7, which is a power spectral density of one of the resulting measurement noise signals. A typical time-history match to this data is shown in Fig. 8. Note that this data exhibits deterministic appearing characteristics such as phase shift and flattened peaks more typical of flight data than the simulations shown in Fig. 3. The Cramer-Rao bounds and estimates of C_{l_β} are shown in Fig. 9. We have succeeded in duplicating the discrepancy with simulated data by using band-limited noise. The band limit is well defined here, and when the Cramer-Rao bounds are corrected for the colored noise, they agree excellently with the scatter.

These results support the conclusion that the discrepancy in the Cramer-Rao bound has now been adequately explained by the presence of colored noise. An important consideration is that in order to accurately reflect flight data scatter, the noise statistics used in computing the Cramer-Rao bound must represent the entire residual error including modeling error

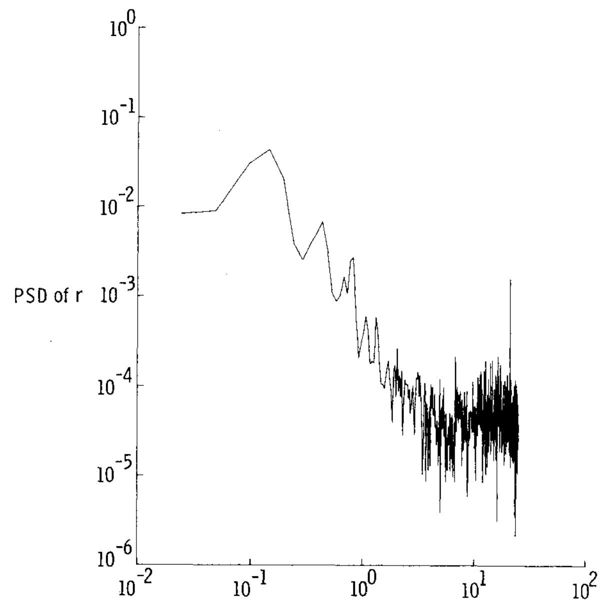

Fig. 6 Power spectral density of yaw rate residuals from flight data.

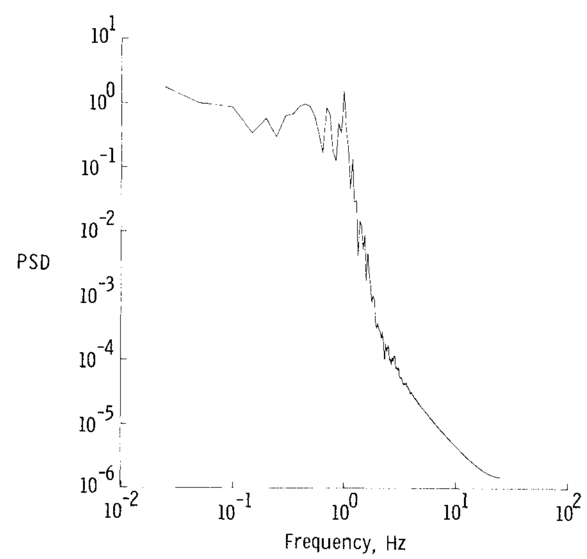

Fig. 7 Power spectral density of simulated colored noise.

contributions (which may be much larger than the actual instrumentation error). Therefore, studies made solely on the basis of instrumentation characteristics[19-21] are likely to be extremely overoptimistic.

Suggested Implementation

The previous section explained the reasons for the discrepancies observed in the Cramer-Rao bounds. It remains to discuss practical implementation of a corrected computation of the bound. Three approaches are considered.

The first approach is to continue the use of fudge factors. The Cramer-Rao bounds, computed ignoring noise coloring and then multiplied by a fudge factor of 5 to 10, have proven useful in practice. The objection to this approach has not been to its utility, but rather to its ad hoc nature and total lack of theoretical justification. Now that the theory has provided an understanding of the need for an extra term, it is not unreasonable to use a value of this term based on past experience rather than analytically computing it for each maneuver. It could be looked upon as an empirically determined spectral adjustment factor instead of as a mysterious fudge factor. The factor has been observed to be relatively

constant over large classes of cases, adding justification to this approach. The advantages are the simplicity and the fact that no changes to current programs are required. The disadvantage is that the engineer must watch for changes in the vehicle or in the analysis which might significantly affect the spectral characteristics of the residuals and thus the factor used. If such changes occur, or if discrepancies are noted, it may be necessary to adjust the factor used. The approach is subject to criticisms of arbitrariness, but when considered as a tool to aid the engineer's evaluation, instead of as an absolute value of accuracy, it has been and can continue to be useful.

The second approach is to examine (manually or automatically) the actual spectrum of the residuals. The break frequency can be evaluated, or the spectral density can be directly used. This method has the advantage of providing the most information. The spectral characteristics of each signal can be adjusted separately, instead of using a single factor for all of the signals. The entire spectral shape can be examined for peculiarities such as resonant modes. The disadvantages are twofold. First, this is the most complex approach. A Fourier transform routine must be included in the analysis program if the adjustment is to be automatic; appropriate plotting routines will also be desired. It is always a good practice to examine at least a few sample residual power spectral density plots anyway, but it keeps things simple to create the power spectral density plots in a separate program. The second disadvantage is that the value to use for the spectral density or the break frequency is not usually obvious from the plot. Figure 6 is a typical example of this problem. The spectrum does not exhibit an obvious flat area followed by a well-defined break. Picking a specific value from the plot can be as much of an art as picking a value for the spectral adjustment factor from experience.

The third approach is a compromise between the first two. It obtains information from the actual residuals, but keeps the programming relatively simple. This approach uses the total power of the low-pass filtered residuals. A simple single-pole filter is used with a break frequency two or three times the system natural frequency. The power of the filtered residuals is then divided by the filter break frequency to give the average power spectral density at frequencies near and below the system natural frequencies. This method does not provide the complete spectral information of a power spectral density plot, but it does pull out a reasonable estimate of the single value we are after with very little work. The bulk of the implementation is filtering the residuals, which is done in the same time loop that computes them. The method is somewhat approximate and requires picking a value for the filter break frequency. The approximations are such that a factor of about 2 typically remains to be multiplied.

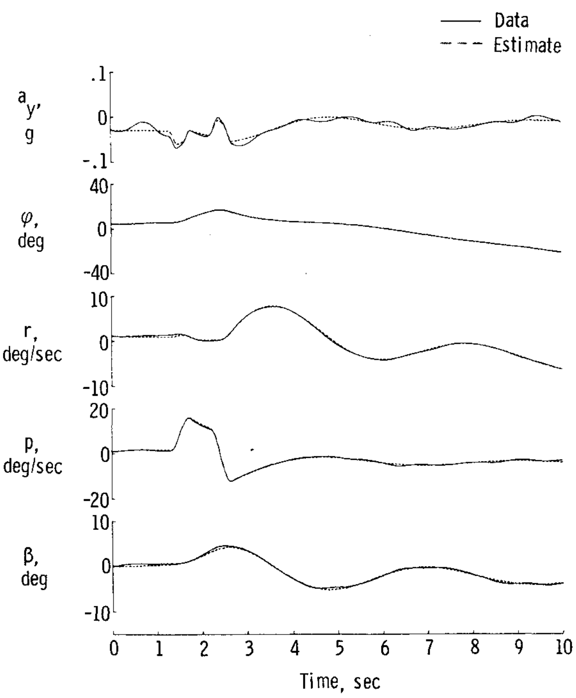

Fig. 8 Time-history match of simulated data with filtered noise.

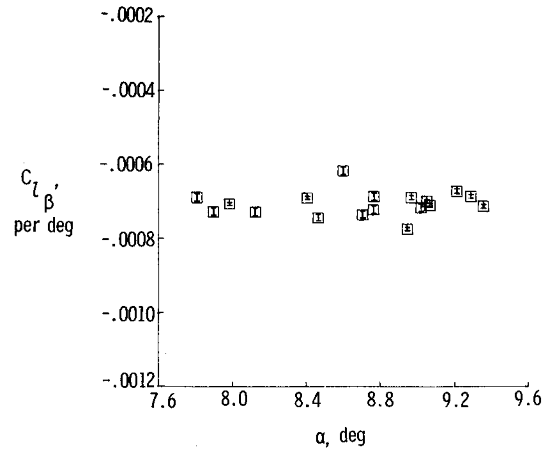

Fig. 9 Estimates from simulated data with colored noise.

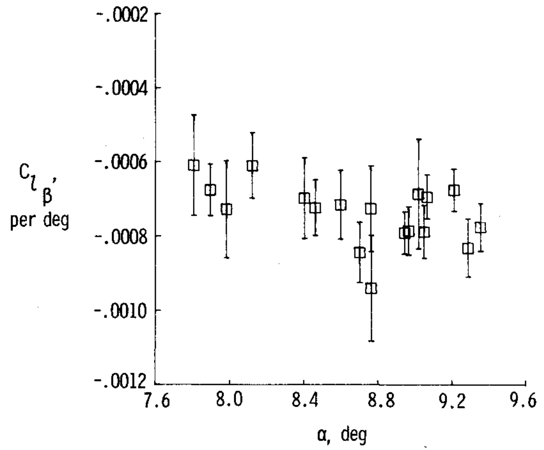

Fig. 10 Estimates from flight data—bounds are adjusted for colored noise.

The method might therefore seem to be little improvement over the first approach, since an empirically determined factor remains (although the factor is now smaller). The advantage is that this approach will notice and approximately account for large changes in the spectral characteristics. If the spectral characteristics of all of the maneuvers are similar, the results from this approach and the first one are equally applicable.

Figure 10 shows the flight data from Fig. 2 with Cramer-Rao bounds computed from the filtered residuals. A break frequency of ½ Hz was used for the filters. The correction factors computed from the filtered residual power were from 3 to 6, and a remaining empirical factor of 2 was used. The magnitudes of the Cramer-Rao bounds on this plot are reasonable and give a good visual indication of the estimate accuracy. Although this plot shows no outstandingly good or poor maneuvers, there is a noticeable tendency for the estimates near the center of the scatter band to have smaller Cramer-Rao bounds than the outliers. This approach is seen to result in a useful estimate of the accuracy.

Conclusions

The Cramer-Rao bound is significantly affected by the colored measurement noise and modeling error present in actual flight data. Computations of the bound that ignore

these effects are often in error by as much as a factor of 10. Approximate corrections for these effects are easily computed and result in reasonable values for the Cramer-Rao bounds. The Cramer-Rao bounds, corrected in this manner, are the best-known theoretical measures of accuracy. They have proven to be extremely useful in application to real flight data.

Modeling error must be considered as a component of the noise statistics. Theoretical studies based only on instrumentation noise characteristics produce extremely overoptimistic results and can lead to incorrect conclusions.

Engineering judgment remains the cornerstone of gaging the accuracy of the estimates. The empirical and theoretical measures should all be taken as tools to aid the engineer's judgment rather than as absolute incontrovertible values. The judgment function combines information from all of the available tools. It must evaluate the applicability of the tools and consider the factors such as modeling error that were ignored in the development of the tools.

References

[1] Iliff, K.W., Maine, R.E., and Montgomery, T.D., "Important Factors in the Maximum Likelihood Analysis of Flight Test Maneuvers," NASA TP 1459, 1979.

[2] Maine, R.E. and Iliff, K.W., "Estimation of the Accuracy of Dynamic Flight-Determined Coefficients," AIAA Paper 80-0171, AIAA 18th Aerospace Sciences Meeting, Pasadena, Calif., Jan. 1980.

[3] Goodwin, G.C. and Payne, R.L., *Dynamic System Identification: Experiment Design and Data Analysis,* Academic Press, New York, 1977, pp. 88-103.

[4] Balakrishnan, A.V., *Communication Theory,* McGraw Hill Book Co. Inc., New York, 1968, Chap. 3.

[5] Cramer, H., *Mathematical Methods of Statistics,* Princeton University Press, Princeton, New Jersey, 1946, pp. 473-524.

[6] Maine, R.E. and Iliff, K.W., "User's Manual for MMLE3, A General FORTRAN Program for Maximum Likelihood Parameter Estimation," NASA TP-1563, 1980.

[7] Iliff, K.W. and Maine, R.E., "Practical Aspects of Using a Maximum Likelihood Estimation Method to Extract Stability and Control Derivatives From Flight Data," NASA TN D-8209, 1976.

[8] Wells, W.R. and Ramachandran, S., "Flight Test Design for Efficient Extraction of Aircraft Parameters," *Proceedings of AIAA 3rd Atmospheric Flight Mechanics Conference,* Arlington, Texas, June 1976, pp. 101-107.

[9] Parrish, R.V. and Steinmetz, G. G., "Lateral Stability and Control Derivatives of a Jet Fighter Airplane Extracted From Flight Test Data by Utilizing Maximum Likelihood Estimation," NASA TN D-6905, 1972.

[10] Suit, W.T., "Aerodynamic Parameters of the Navion Airplane Extracted From Flight Data," NASA TN D-6643, 1972.

[11] Williams, J.L. and Suit, W.T., "Extraction From Flight Data of Lateral Aerodynamic Coefficients for F-8 Aircraft With Supercritical Wing," NASA TN D-7749, 1974.

[12] Tanner, R.R. and Montgomery, T.D., "Stability and Control Derivative Estimates Obtained From Flight Data for the Beech 99 Aircraft," NASA TM-72863, 1979.

[13] Iliff, K.W., Maine, R.E., and Steers, S.T., "Flight-Determined Stability and Control Coefficients of the F-111A Airplane," NASA TM-72851, 1978.

[14] Powers, B.G., "Phugoid Characteristics of a YF-12 Airplane With Variable-Geometry Inlets Obtained in Flight Tests at a Mach Number of 2.9," NASA TP-1107, 1977.

[15] Sim, A.G., "A Correlation Between Flight-Determined Derivatives and Wind-Tunnel Data for the X-24B Research Aircraft," NASA SX-3371, 1976.

[16] Maine, R.E., "Aerodynamic Derivatives for an Oblique Wing Aircraft Estimated From Flight Data by Using a Maximum Likelihood Technique," NASA TP-1336, 1978.

[17] Iliff, K.W. and Maine, R.E., "Further Observations on Maximum Likelihood Estimates of Stability and Control Characteristics Obtained From Flight Data," AIAA Paper 77-1133, AIAA Atmospheric Flight Mechanics Conference, Hollywood, Fla., Aug. 1977.

[18] Balakrishnan, A.V. and Maine, R.E., "Improvements in Aircraft Extraction Programs," NASA CR-145090, 1975.

[19] Hodge, W.F. and Bryant, W.H., "Monte Carlo Analysis of Inaccuracies in Estimated Aircraft Parameters Caused by Unmodelled Flight Instrumentation Errors," NASA TN D-7712, 1975.

[20] Sorensen, J.A., "Analysis of Instrumentation Error Effects on the Identification Accuracy of Aircrafft Parameters," NASA CR-112121, 1972.

[21] Gupta, N.K. and Hall, W.E. Jr., "Design and Evaluation of Sensor Systems for State and Parameter Estimation," *Journal of Guidance and Control,* Vol. 1, Nov.-Dec., 1978, pp. 397-403.