Application of Parameter Estimation to Highly Unstable Aircraft

Richard E. Maine* and James E. Murray†

NASA Ames Research Center, Edwards, California

This paper deals with the application of parameter estimation to highly unstable aircraft. It includes a discussion of the special problems in applying the output error method to such aircraft and demonstrates that the filter error method eliminates these problems. This paper shows that the maximum likelihood estimator with no process noise does not reduce to the output error method when the system is unstable. It also proposes and demonstrates an ad hoc method that is similar in form to the filter error method but applicable to nonlinear problems. Flight data from the X–29 forward–swept–wing demonstrator are used to illustrate the problems and methods discussed.

Nomenclature

A,B,C,D	= system matrices
C_{m_α}	= coefficient of pitching moment due to angle of attack
$C_{m_{\delta c}}$	= coefficient of pitching moment due to canard
C_{m_q}	= coefficient of pitching moment due to pitch rate
C_{N_α}	= coefficient of normal force due to angle of attact
FF^*	= process noise covariance matrix
$f(.)$	= state derivative function
GG^*	= measurement noise covariance matrix
$g(.)$	= response function
$J(.)$	= cost function
K	= Kalman gain matrix
MMLE3	= modified maximum likelihood estimator program, version 3
N	= number of time points
n	= process noise vector
pEst	= parameter estimation program
Q	= covariance matrix of predicted state
R	= innovation covariance matrix
t	= time
u	= input vector
x	= state vector
z	= response vector
η	= measurement noise vector
ξ	= vector of unknown parameters
Φ	= state transition matrix
Ψ	= input transition matrix

Superscripts

$(\tilde{n})$	= estimated or predicted
$(\acute{n})$	= corrected
$(n)^*$	= complex conjugate transpose

Introduction

THE output error method of parameter estimation[1-6] is the most widely used means of estimating aircraft stability and control derivatives from flight data. There is a large body of experience in applying this method to real flight data.[2-5] Like all methods, it has some shortcomings, but it has proved practical for obtaining good results in a timely manner in real flight test and research environments. It is the standard method used for routine application at the Dryden Flight Research Facility of NASA Ames Research Center.

One area where output error methods have severe difficulty is in application to highly unstable systems. For aircraft such as the X-29 that are highly unstable in open loop, it has proved difficult and expensive (in labor and computer time) to obtain acceptable results with the output error method. This paper investigates only issues specifically relating to instability. The references discuss numerous other issues that must be addressed in any successful application.

Equation error methods are not subject to the same problems as output error methods with unstable systems. We generally prefer to avoid equation error because of the problems associated with the accurate state reconstruction that it requires.[1,7] Several authors have successfully applied equation error methods to aircraft stability and control estimation[3,7,8] but, in this paper, we investigate other approaches.

It is widely known that filter error methods have stabilizing properties, which should allow them to be easily applied to unstable systems. However, little if anything seems to have been published on this subject.

In this paper, we first describe the difficulties of applying output error to highly unstable aircraft. We then show how the filter error method alleviates these difficulties. We illustrate these discussions with flight data from the X-29 forward-swept-wing demonstator. Finally, we discuss some practical limitations of the filter error method and suggest an alternative ad hoc method of stabilizing output error estimations. The alternative method is similar in computational form to the filter error method, but its ad hoc derivation allows broader applicability.

Output Error Method

The output error method applies to dynamic systems of the form:

$$x(0) = x_0(\xi) \tag{1a}$$

$$\dot{x}(t) = f[x(t), u(t), \xi] \tag{1b}$$

$$z(t_i) = g[x(t_i), u(t_i), \xi] + \eta_i \qquad i = 1,...,N \tag{1c}$$

The input vector u is assumed to be known as a function of time. The response vector z is measured at the discrete time points t_i. The measurement noise η is assumed to be a sequence of independent Gaussian vectors with mean zero and nonsingular covariance GG^*. Throughout this paper, we will assume that GG^* is known; the extension for unknown GG^* is easy but irrelevant to the issues of this paper. The objective is to estimate the values of the unknown parameters in ξ.

The output error method uses the cost function

$$J(\xi) = \sum_{i=1}^{N} [z(t_i) - \tilde{z}_\xi(t_i)]^*(GG^*)^{-1}$$
$$\times [z(t_i) - \tilde{z}_\xi(t_i)] \qquad (2)$$

where $\tilde{z}_\xi$ is obtained by integrating the system equations using the appropriate value of ξ and ignoring the noise. Hence,

$$\tilde{x}_\xi(0) = x_0(\xi) \qquad (3a)$$

$$\dot{\tilde{x}}_\xi(t) = f[\tilde{x}_\xi(t), u(t), \xi] \qquad (3b)$$

$$\tilde{z}_\xi(t_i) = g[\tilde{x}_\xi(t_i), u(t_i), \xi] \qquad i = 1,...,N \qquad (3c)$$

The output error estimate of ξ is the value that minimizes this cost function. If the system is stable, this is the maximum likelihood estimate of ξ. The Gauss-Newton algorithm is suitable for numerically finding the minimizing value of ξ.

We have two computer programs that implement output error parameter estimation: the modified maximum likelihood estimator program, version 3 (MMLE3)[9] which is designed for batch usage on linearized systems, and the more recent parameter estimation program (pEst), which is designed for interactive usage on nonlinear systems. (Formal documentation for pEst is not yet available.) This paper uses results from both programs.

Difficulties with Output Error

For the aircraft stability and control application, we use a model in which the input u to the system consists of control surface positions, such as elevator, aileron, and rudder. The response consists of measurements from rate gyros, accelerometers, and other sensors. This model, therefore, considers only the open-loop response of the airframe to the control surface inputs. Any feedback control system or shaping of the pilot's inputs is external to the analysis model; Fig. 1 illustrates this situation.

This choice of analysis model allows us to concentrate on the aerodynamics, which are what we want to estimate; it avoids the extensive complications of modeling the control system, which is peripheral to our objectives. The control system influences the design of maneuvers to be flown for stability and control derivative estimation, but it does not explicitly enter into the estimation process. Maneuvers obtained with the control system feedback turned off usually give better parameter estimation results, but this is not always practical. This rationale is further elaborated by Maine and Iliff.[2] Output error using such open-loop models has been routinely applied to numerous aircraft with feedback control systems.

Alternatively, closed-loop characteristics can be estimated using equivalent system models that are adequate for some applications. Such equivalent system models do not allow estimation of the pure aerodynamic characteristics that are required for many applications.

When the open-loop airframe is highly unstable, the output error method requires numerical integration of the unstable equations of motion in Eq. (3) with various values of ξ. The closed-loop airplane will presumably be stable or nearly so (otherwise, the project will be mercifully short), but the stabilization achieved through the feedback system is not reflected in the open-loop analysis model. In principle, the output error

method is applicable to unstable models, but there are severe practical difficulties when the model is highly unstable.

The most obvious difficulty is that the integrated solution is quite likely to exceed the limits of the machine or program, causing the analysis program to terminate. Unless the program is extremely cautious about detecting and recovering from such problems, the termination is likely to be an abrupt abort, destroying the information about any progress previously made toward a solution. Such behavior makes analysis extremely time-consuming.

For aircraft as unstable as the X-29, the integration of the equations of motion is extremely sensitive to even minuscule changes in "unimportant" parameters. Figure 2 shows a longitudinal maneuver of the X-29 and a match obtained (with much effort) using an output error method. For the GG^* we were using, the cost function value corresponding to this match is 0.6161. Figure 3 shows the same maneuver and the computed response that resulted from changing the initial condition of the angle-of-attack state by 0.01 deg. This small change causes the solution to diverge severely. The cost function value corresponding to this match is 754,150.

In some simulated data cases, we had difficulty saving and restoring the match because the derivatives were saved only to ten significant figures; restoring the derivative values changed in the tenth significant digit caused changes of several orders of magnitude in the cost function value. These simulated cases were based on actual X-29 flight conditions.

In the most extreme cases, it is possible that no parameter values exist within the finite resolution of the computer that give a nondivergent solution to the equations of motion. We did not encounter any cases quite this extreme, but the possibility is real.
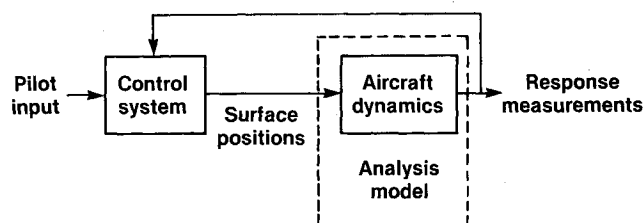

Fig. 1 Block diagram of aircraft with control system.
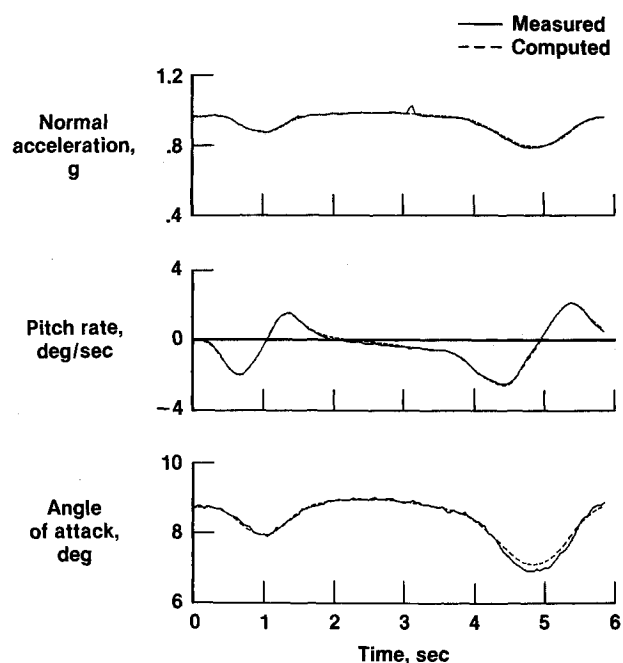

Fig. 2 Match of X-29 maneuver using output error method.
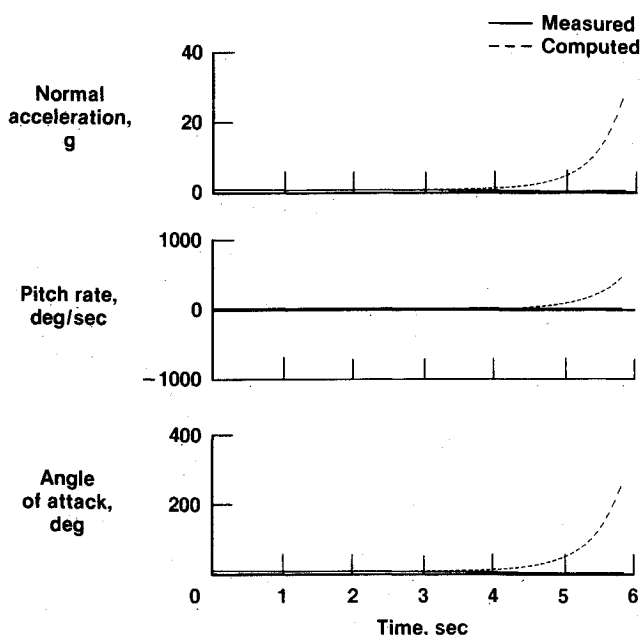
Fig. 3 Match of X-29 maneuver after insignificant change.

The pEst program detects and recovers from extreme state values before they cause overflow aborts. This avoids the time-consuming process of restarting the program and retracing prior steps every time the solution diverges. It does not, however, provide constructive progress toward minimization of the cost function. It simply recovers to a place where the engineer can choose an alternative approach to the minimization. This process is slow and labor-intensive.

We experimented with and successfully implemented schemes for automatically minimizing the cost function, even when the initial attempts to evaluate the cost function failed. Although some of these schemes worked better than others, all used excessive amounts of computer time.

A second difficulty is that the huge calculated response is usually far beyond the bounds within which approximations of linearity in the parameters are reasonable. This is a problem even when the response is small enough to avoid overflow. Efficient minimization algorithms invariably work best in the regions where the response is approximately linear in the parameters (and, thus, the cost function is approximately quadratic). Algorithms robust enough to reliably minimize the cost starting from excessively poor points tend to be excruciatingly slow.

A third difficulty is that the extreme sensitivity in some directions of the parameter plane causes several numerical roundoff problems. It can be difficult to do such simple operations as computing gradients by finite differences. This is because there is no overlap between the range of perturbations small enough to give linearity (a basic assumption of finite-difference computation) and perturbations large enough to avoid underflow and roundoff error. Singularity and ill conditioning may also be encountered in methods that involve matrix inversions.

Figure 4 shows isoclines of the cost function of an X-29 flight maneuver. In order to reduce the data to an easily presentable form, these isoclines are based on the cost function only as a function of the two parameters C_{m_q} and C_{m_α}. The remaining parameters are all constant at their estimated values. The minimum cost function value is 1.584; the outer isocline shown in Fig. 4 is for a value of 50,000. The isoclines are extremely eccentric to the extent that we cannot reasonably show a fully closed isocline on the figure; even the isocline for a cost function value of 2 extends well beyond the edges of Fig. 4.

The broadening of the valley toward the lower left of Fig. 4 reflects the trend of the cost toward a constant asymptote in the direction of higher stability. In most directions, the cost grows rapidly and without bound. However, as the parameters become very stable, the computed response trends toward zero. Thus, the cost stays bounded for a large distance in this direction.

Maximum Likelihood Estimator

Maximum Likelihood estimators for dynamic systems are often divided into three categories. For models with both process and measurement noise, the filter error algorithm is obtained. When there is measurement noise but no process noise, the filter error algorithm simplifies to the form of output error as a special case. When the states can be measured with no measurement noise, the estimator reduces to some form of an equation error algorithm. All three algorithms are discussed in Maine and Iliff.[1]

The categories just mentioned follow the conventional wisdom and do not appear obviously affected by the issue of the system stability. In particular, it is easy to derive the output error algorithm directly, with no reference to filter error, and the issue of stability never arises.

In this paper, we will look at output error as a limiting case of filter error and discover an overlooked subtlety. The mathematical results are not new. The basic issues are studied in some detail for the one-dimensional case in Maine and Iliff[1] and in Balakrishnan;[10] these studies elaborate on the stability and uniqueness conditions derived in Kalman and Bucy.[11] However, the connection between these theoretical results and practical application of parameter estimation methods does not seem to have been explored.

The filter error algorithm is practical only for linear systems. The general form of a linear dynamic system model is

$$x(0) = x_0 \tag{4a}$$

$$\dot{x}(t) = Ax(t) + Bu(t) + n(t) \tag{4b}$$

$$z(t_i) = Cx(t_i) + Du(t_i) + \eta_i \qquad i = 1...N \tag{4c}$$

Cost function value			
A	2	H	500
B	5	I	1,000
C	10	J	2,000
D	20	K	5,000
E	50	L	10,000
F	100	M	20,000
G	200	N	50,000

X denotes minimum

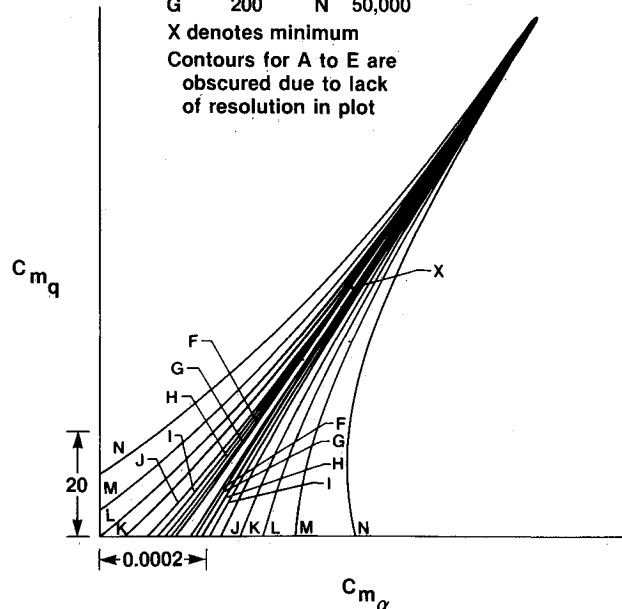Contours for A to E are obscured due to lack of resolution in plot

Fig. 4 Isoclines of output error cost function.

The system matrices are functions of ξ, although this is not indicated in Eq. (4). The process noise n is a white-noise process with spectral density FF^*. The measurement noise η is a sequence of independent Gaussian vectors with mean zero and nonsingular covariance GG^*.

The maximum likelihood estimator for this system minimizes the cost function:

$$J(\xi) = \sum_{i=1}^{N} [z(t_i) - \bar{z}_\xi(t_i)]^* R^{-1}$$

$$\times [z(t_i) - \bar{z}_\xi(t_i)] \qquad (5)$$

where R is the innovation covariance matrix and $\bar{z}_\xi$ is computed from the Kalman filter. We are interested in the steady-state form of the filter.

$$\hat{x}_\xi(t_0) = x_0 \qquad (6a)$$

$$\bar{x}_\xi(t_i) = \phi\hat{x}(t_i - 1) + \psi Bu\left(t_i - \frac{1}{2}\right) \qquad (6b)$$

$$\bar{z}_\xi(t_i) = C\bar{x}_\xi(t_i) + Du(t_i) \qquad (6c)$$

$$\hat{x}_\xi(t_i) = \bar{x}_\xi(t_i) + K[z(t_i) - \bar{z}_\xi(t_i)] \qquad (6d)$$

The steady-state Kalman gain matrix K is given by

$$K = QC^*(CQC^* + GG^*)^{-1} \qquad (7)$$

where Q is the steady-state covariance of the predicted state, given as the limit of the discrete-time Riccati equation

$$Q_{i+1} = \phi[Q_i - Q_iC^*(CQ_iC^*$$

$$+ GG^*)^{-1}CQ_i]\Phi^* + FF^* \qquad (8)$$

If all unstable modes of the system are (A,F)-controllable and (A,C)-observable, then Eq. (8) is guaranteed to converge to a unique steady-state solution independent of the initial covariance. The resulting filter is also guaranteed to be stable.

We consider the limit as FF^* approaches zero. Substitution verifies that Eq. (8) has a constant solution of Q_i equals zero when FF^* equals zero. This is the solution that reduces to the special form of output error. When all modes of the system are stable, Eq. (8) converges to zero from any starting covariance. Furthermore, the steady-state solutions for nonzero FF^* smoothly approach zero as FF^* approaches zero. Thus, for stable systems with no process noise, the filter error algorithm clearly reduces to output error. The continuity of the solutions ensures that output error is also a good approximation for cases in which the process noise is negligible but nonzero.

For unstable systems with FF^* equal to zero, the zero solution is still valid, but it is not unique. Let us assume that all unstable modes are (A,C)-observable and that there are no modes with precisely neutral stability. In this case, there is also a nonzero steady-state solution, and this solution yields a stable filter. Any starting covariance that is not exactly zero converges to the nonzero steady-state solution. The steady-state solutions for nonsingular FF^* are unique and smoothly approach the same nonzero steady-state solution as FF^* approaches zero.

In short, the zero solution is anomalous. Nothing close to the output error method is obtained unless the process noise covariance is exactly zero and the initial covariance is exactly zero. The output error algorithm is not a close approximation for cases with nonzero process noise, no matter how negligible the covariance. The error in such approximation depends on the degree of instability rather than on the magnitude of the covariance. Inasmuch as few models are absolutely without er-

ror, the output error algorithm is difficult to justify for highly unstable systems.

Implementation of the filter with nonzero covariance is trivial in any program already implementing the filter error algorithm. This is nothing but a particular case of filter error and does not require any special handling. The MMLE3 program previously tested whether F was zero and bypassed the filter computations in that case. For stable systems, the omitted terms were all zero, and this special handling saved computer time. To implement our proposed algorithm, we added an option to do the full computation whether F was zero or not. Nothing was changed in the code that does the computation, just in the code that decides whether to do it.

The MMLE3 program uses a form of Potter's method[12] to solve the steady-state Riccati equation. If the unstable modes are all observable and there are no neutrally stable modes, this method works well with F equal to zero. The observability conditions are reasonable and necessary for a stable filter. If there are neutrally stable modes, the stable filter may not have a steady-state form; this is typically manifested by a defective Hamiltonian matrix, for which the eigen value determination fails.

Figures 5 and 6 show results from a set of X-29 flight data analyzed using both output error and the algorithm with nonzero covariance. Figure 5 shows results from the pEst program using output error. Obtaining these results required extensive engineering effort and computer time from a series of interactive sessions. Figure 6 shows corresponding results from a single batch run of the MMLE3 program using the nonzero covariance. Convergence was rapid and monotone in
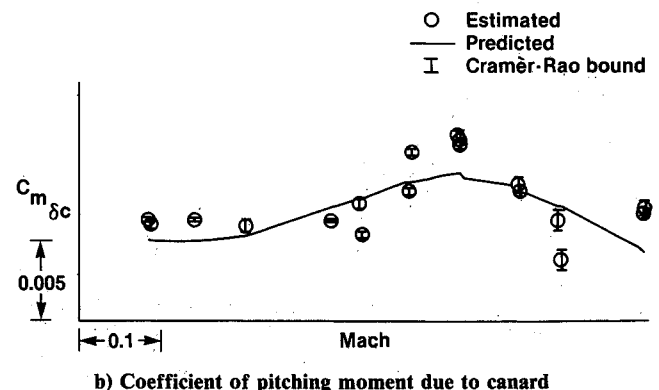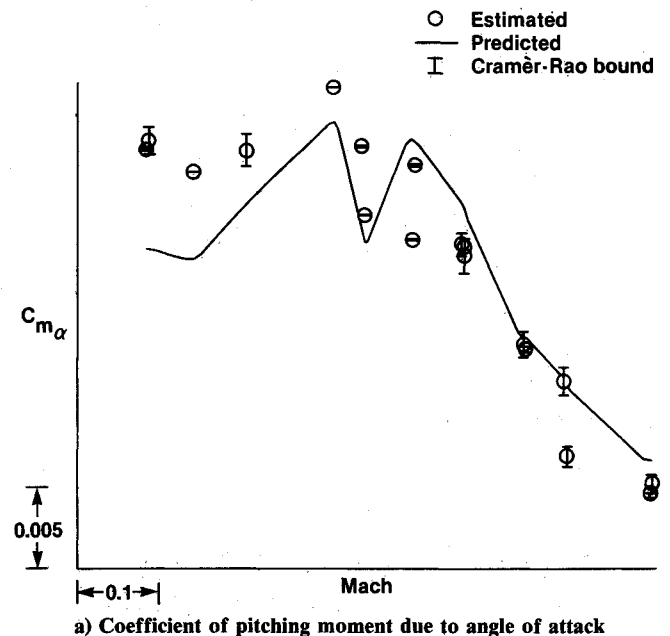
a) Coefficient of pitching moment due to angle of attack

b) Coefficient of pitching moment due to canard

Fig. 5 X-29 derivatives from output error method.

O   Estimated
——   Predicted
I   Cramèr-Rao bound

C_{m_α}

0.005

|←0.1→| Mach

a) Coefficient pitching moment due to angle of attack

O   Estimated
——   Predicted
I   Cramèr-Rao bound

$C_{m_{\delta c}}$

0.005

|←0.1→| Mach
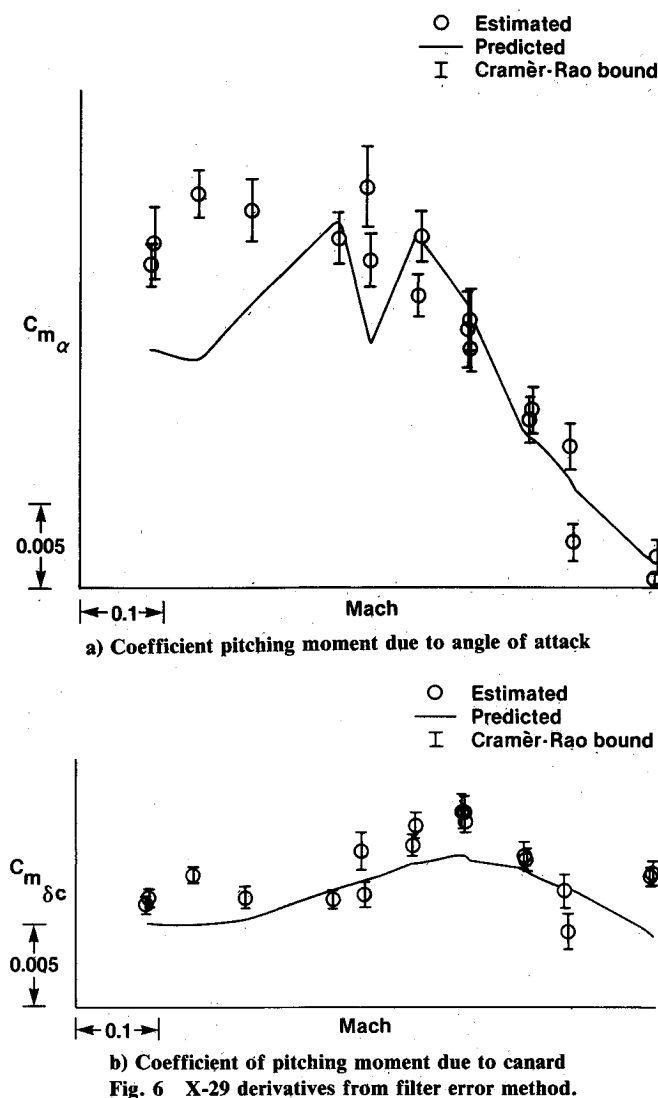
b) Coefficient of pitching moment due to canard
Fig. 6   X-29 derivatives from filter error method.

every case; most cases converged within tight bounds in 5 or 6 iterations, and the worst case required 7 iterations.

Although there are detectable differences in the individual estimates, the results from the two methods show the same trends and scatter magnitudes. Unfortunately, the currently obtainable X-29 data are of marginal quality for parameter estimation for reasons peripheral to this paper (primarily control system limitations on the maneuver design). The results are adequate only to show major trends. The two methods appear equally able to show these trends within the limitations imposed by the available data.

Stabilization by Arbitrary Feedbacks

There are several practical complications in applying the filter error method to real data. One of the most obvious problems is that the filter error method is practical only for linear systems. In most cases, it is reasonable to linearize the system to obtain an approximation to the desired filter, but this can be a substantial complication. Application of the output error method to nonlinear systems is straightforward.

Efficiency considerations further constrain the practical application of filter error methods. Efficient implementation requires the use of the steady-state form of the filter. This implies that the system must be time-invariant and that the data must be at a constant sample rate. The sample rate requirement is a minor nuisance accepted in several programs; it primarily affects the handling of data dropouts. However, the time-invariance requirement is quite restrictive in the aircraft stability and control application; it requires dynamic pressure and other flight condition parameters to be held constant

throughout a maneuver. Analysis programs without such restrictions have proved far more useful in application.

We would like to obtain the benefits of the filter error method without such narrow restrictions on the form of the system equations. Toward this end, we consider the filter error method from a purely mechanistic viewpoint rather than a statistical one. We are interested primarily in the benefits of stability and improved convergence. Our presumption is that the parameter values from the output error method are reasonable, but they are too difficult to obtain.

The filter error method differs from the output error method only in the addition of a feedback term. This feedback stabilizes the filter, which results in improved convergence characteristics of the algorithm. The feedback is proportional to the fit error; therefore, an idealized case with zero fit error would give identical results with output error and filter error methods. For a case with small, random fit errors, the differences between output error and filter error methods should be small.

This feedback in the filter should not be confused with any feedback that many be in the aircraft control system. The filter feedback helps stabilize the analysis algorithm and need not bear any relation to the control system feedbacks that help stabilize the aircraft.

We can easily implement a stabilizing feedback of the fit errors without reference to statistical methods. Such a stabilization need not be restricted to linear systems (except insofar as it might be difficult to design a stabilizing feedback for a complicated nonlinear system). As long as the feedback error term is small, we would expect little change from the output error estimates. If the feedback term is large, we anticipate substantial changes in the estimates. The validity of such results would be seriously doubtful.

Although such arbitrary stabilizations do not have the strong theoretical basis of the filter error method, they are more flexible in easily accommodating nonlinearities and other complications. They also allow more direct consideration of practical engineering judgment to improve the results. Naturally, poor exercise of that judgment can degrade the results. In some cases, we can rationalize the arbitrary feedbacks by the observation that they approximate the feedbacks that the filter error method would have with some covariance matrix.

We implemented user-specified stabilizations in the pEst program and applied them to the X-29 data. Although we experimented with several feedbacks, our final choice was to use only a feedback of the pitch rate signal to the pitch rate state (effectively augmenting the damping). The main reason for this choice was that we had substantially higher confidence in the pitch rate measurement than in the angle-of-attack measurement. We were dubious about the wisdom of corrupting the integration by feeding back the questionable angle-of-attack measurements. Feedback of other measurements, such as normal acceleration, is a little more complicated because of variables in their observation equations.

Our initial implementation of the feedback added it to the system differential equations. Although this continuous-time feedback worked successfully, we found it more convenient to use an alternative implementation as a discrete correction at each measurement time point. This alternative implementation closely parallels the form of the discrete-time Kalman filter. The primary benefit of the discrete-time implementation is that high feedback gains can be used without introducing stability problems in the numerical solution of the differential equations. In fact, if all states have discrete-time feedback with a gain of one (which corresponds to infinite gain in the continuous-time implementation), the method transforms into a form of equation error.

Figure 7 shows cost function isoclines for several values of the discrete-time feedback gain. These isoclines are based on the same data as used for Fig. 4, which corresponds to a feedback again of zero. The feedback removes the highly eccentric
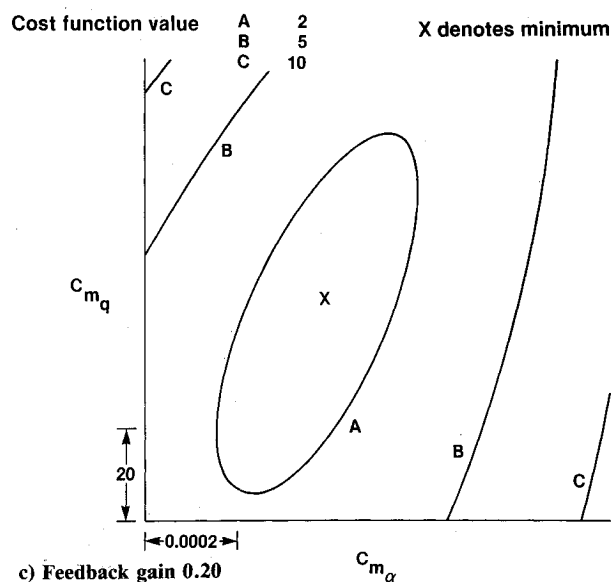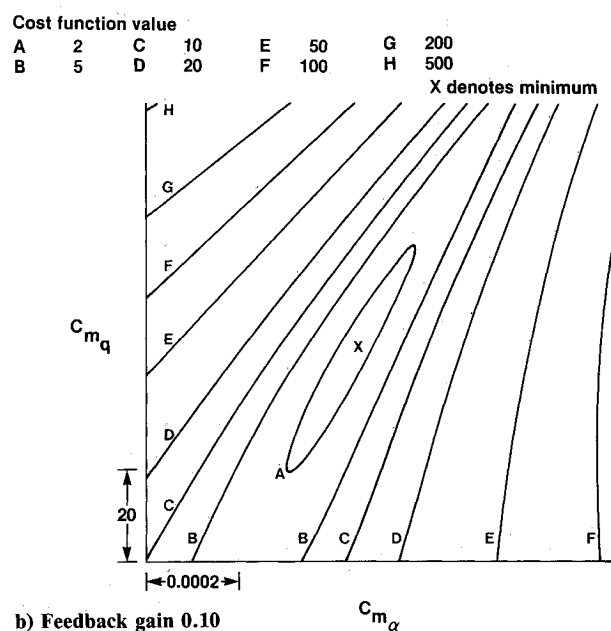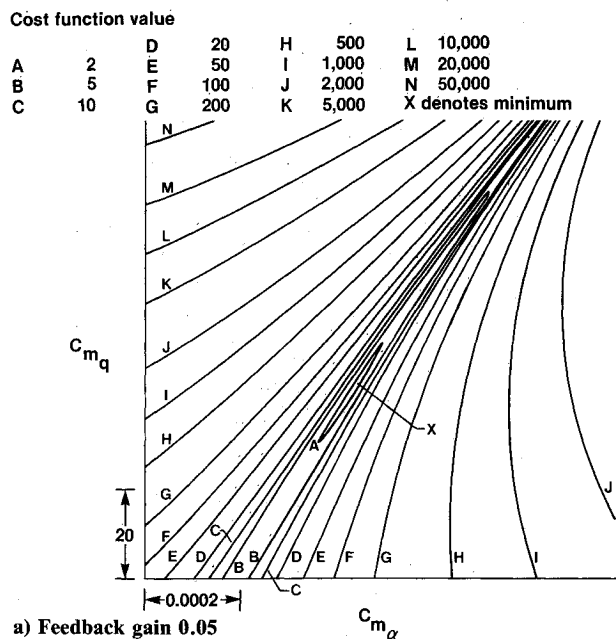
Cost function value

		D	20	H	500	L	10,000
A	2	E	50	I	1,000	M	20,000
B	5	F	100	J	2,000	N	50,000
C	10	G	200	K	5,000	X denotes minimum

a) Feedback gain 0.05

Cost function value

A	2	C	10	E	50	G	200
B	5	D	20	F	100	H	500

X denotes minimum

b) Feedback gain 0.10

Cost function value

A	2
B	5
C	10

X denotes minimum

c) Feedback gain 0.20

Fig. 7   Isoclines of cost functions with pitch rate feedback.

isoclines and steep gradients of the cost function. The isoclines in Fig. 7 show the cost function to be approximately quadratic over a much larger portion of the parameter space when feedback is used. This implies better convergence characteristics. The minimum point of the cost function moves very slightly as a function of the feedback gain. At higher gains than indicated in Fig. 7, the change in the minimum is more significant.

Figure 8 shows results from applying a discrete-time pitch rate gain of 1.0 to the same data as used in Figs. 5 and 6. These results were obtained from a single batch submittal. Convergence in all cases was very rapid with this high feedback gain; most cases converged in two iterations. The results are similar to those shown in Figs. 5 and 6. Figures 5 and 8 are the two extreme cases for the feedback gain; intermediate gain values should give results between those shown in these two figures.

There is little basis for choosing among the data in Figs. 5, 6, and 8. The differences are generally within the accuracy of the data, particularly considering the marginal quality of the available maneuvers. The output error method shows significantly smaller Cramér-Rao bounds than the other methods, but comparison of the Cramér-Rao bounds among different methods can be misleading. (The bounds are most useful for comparing difference maneuvers or different parameters from a single analysis method.) In particular, the Cramér-Rao bounds from the output error method are of dubious validity for unstable systems; the instability invalidates the arguments used to justify the approximations in the derivation and computation of the bound. These problems do not affect the other two methods.
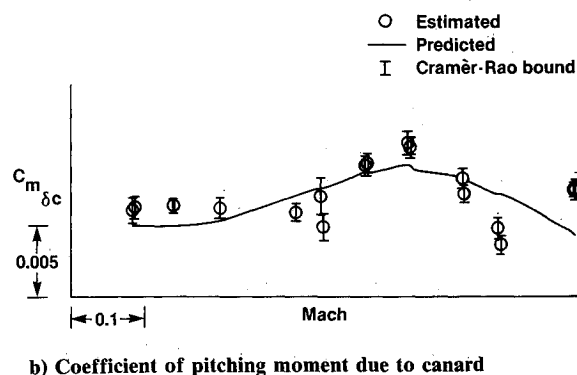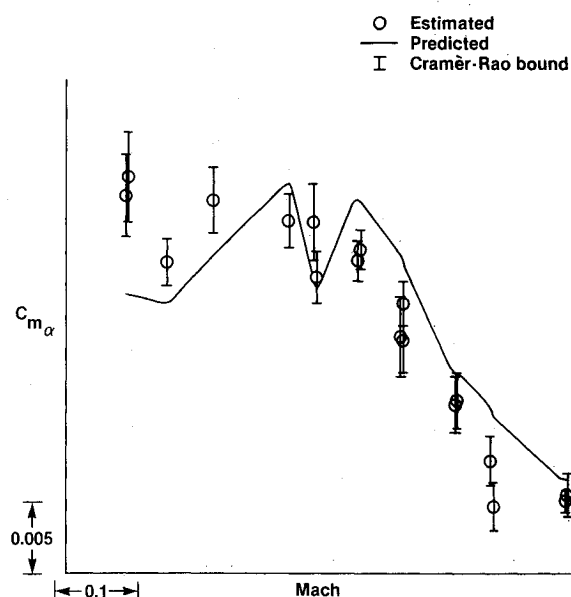
O Estimated
—— Predicted
I Cramér-Rao bound

a) Coefficient of pitching moment due to angle of attack

O Estimated
—— Predicted
I Cramér-Rao bound

b) Coefficient of pitching moment due to canard

Fig. 8   X-29 derivatives with 1.0 pitch rate gain.

Concluding Remarks

Direct application of the output error method to highly unstable systems such as the X-29 is extremely difficult. We have shown that the appropriate maximum likelihood estimator for unstable systems is not the output error method, but the filter error method, which has nonzero gain even when the process noise covariance is zero. We have demonstrated that the filter error method does remove the problems associated with the instability.

The strictly defined filter error method is limited in its applicability. However, it is easy to implement a simple ad hoc method of stabilizing the output error method. The ad hoc method has a form similar to that of filter error and gives essentially the same results but is applicable to many more situations. The main difference is how the feedback gains are obtained. We currently favor the ad hoc method as more readily applicable to our needs, but either method is usable. Both the ad hoc method and the strict filter error method work well with the X-29 test data.

References

[1]Maine, R. E. and Iliff, K. W., "Identification of Dynamic Systems: Theory and Formulation," NASA RP-1138, 1985.

[2]Maine, R. E. and Iliff, K. W., "Application of Parameter Estimation to Aircraft Stability and Control — The Output-Error Approach," NASA RP-1168, 1986.

[3]AGARD, "Parameter Identification," AGARD LS-104, 1979.

[4]Foster, G. W., "A Description of the Weighted Least Squares Output Error Method of Parameter Identification," RAE TM FS 215, 1978.

[5]Jategaonkar, R. and Plaetschke, E., "Maximum Likelihood Parameter Estimation from Flight Test Data for General Non-Linear Systems," DVLR-FB 83-14, 1983.

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

[7]Bach, R. E. and Wingrove, R. C., "Application of State Estimation in Aircraft Flight Data Analysis," AIAA Paper 83-2087, 1983.

[8]Breeman, J. H., Erkelens, L. J. J., and Nieuwpoort, A. M. H., "Determination of Performance and Stability Characteristics from Dynamic Manoeuvres With a Transport Aircraft Using Parameter Identification Techniques," National Aerospace Laboratory (NLR), Amsterdam, the Netherlands, NLR MP 84024 U, 1984.

[9]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.

[10]Balakrishnan, A. V., *Kalman Filtering Theory,* Optimization Software, Inc., New York, 1984.

[11]Kalman, R. E. and Bucy, R. S., "New Results in Linear Filtering and Prediction Theory," *Journal of Basic Engineering,* Vol. 63, March 1961, pp. 95–107.

[12]Potter, J. E., "Matrix Quadratic Solutions," *SIAM Journal on Applied Mathematics,* Vol. 14, May 1966, pp. 496–501.