NASA
Reference
Publication
1168

June 1986

Application of Parameter Estimation to Aircraft Stability and Control

The Output-Error Approach

Richard E. Maine
and Kenneth W. Iliff

NASA

NASA
Reference
Publication
1168

1986

Application of Parameter Estimation to Aircraft Stability and Control

The Output-Error Approach

Richard E. Maine
and Kenneth W. Iliff

Ames Research Center
Dryden Flight Research Facility
Edwards, California

NASA
National Aeronautics
and Space Administration

Scientific and Technical
Information Branch

TABLE OF CONTENTS

		Page
	NOMENCLATURE	v
1.0	INTRODUCTION	1
	1.1 SCOPE AND BASIS	1
	1.2 OUTPUT-ERROR METHOD	2
	1.3 AIRCRAFT STABILITY AND CONTROL APPLICATION	2
	1.4 PURPOSE	3
	1.5 FLIGHT TEST PROGRAM	5
2.0	AIRCRAFT EQUATIONS OF MOTION	7
	2.1 SIGNAL DEFINITIONS	7
	2.1.1 Control Surfaces	7
	2.1.2 Body-Axis System	7
	2.1.3 Euler Angles	7
	2.1.4 Angular Rates	8
	2.1.5 Wind-Relative Velocity	8
	2.1.6 Linear Accelerations	9
	2.2 SIX-DEGREE-OF-FREEDOM EQUATIONS OF MOTION	9
	2.2.1 Newtonian Mechanics	9
	2.2.2 External Forces and Moments	10
	2.2.3 Euler Angles	11
	2.2.4 Polar Coordinate Velocity Form	12
	2.2.5 Collected Equations	13
	2.2.6 Spatial Position	14
	2.3 OBSERVATION EQUATIONS	15
	2.4 AERODYNAMIC MODELS	16
3.0	SIMPLIFYING THE EQUATIONS	21
	3.1 NONLINEAR VERSUS LINEARIZED IMPLEMENTATIONS	21
	3.2 COORDINATE SYSTEMS	22
	3.3 SIMPLIFICATION USING MEASURED DATA	22
	3.4 UNCOUPLED STATE EQUATIONS	23
	3.4.1 Longitudinal Equations	23
	3.4.2 Lateral-Directional Equations	24
	3.5 OBSERVATION EQUATIONS	25
	3.5.1 Sensor Position	26
	3.5.2 Coordinate Transformation	26
	3.6 BIASES AND INITIAL CONDITIONS	27
	3.6.1 Perturbation Equations	28
	3.6.2 Physical States	28
	3.6.3 Special Considerations	29
	3.6.4 Aerodynamic Biases	30
	3.7 LINEARIZED EQUATIONS	30
	3.7.1 Linearization Using Measured Data	30
	3.7.2 Axis Transformations	31
	3.8 COLLECTED EQUATIONS	32
	3.8.1 Longitudinal Equations	32
	3.8.2 Lateral-Directional Equations	34
4.0	DETAILED COMPUTATIONAL ESTIMATION EXAMPLES	37
	4.1 DESCRIPTION OF THE MMLE3 PROGRAM	37
	4.2 EQUATIONS FOR A SIMPLE EXAMPLE	37
	4.3 COMPUTATIONAL DETAILS OF MINIMIZATION	40
	4.3.1 Example With No Measurement Noise	40
	4.3.2 Example With Measurement Noise	41
	4.4 COST FUNCTIONS	42
	4.4.1 One-Dimensional Case	42
	4.4.2 Two-Dimensional Case	44
	4.5 ESTIMATION USING FLIGHT DATA	45
	4.5.1 Hand Calculation Example	45
	4.5.2 Cost Function for Full Aircraft Problem	46
	4.5.3 Cramér-Rao Bounds	47
	4.6 SUMMARY	47
5.0	PREFLIGHT DATA	65
	5.1 PREDICTED AERODYNAMIC CHARACTERISTICS	65
	5.1.1 Uses	65
	5.1.2 Theoretical Computations	65
	5.1.3 Wind-Tunnel Tests	65
	5.1.4 Independent Flight Tests	66
	5.1.5 Reference Geometry	66
	5.2 AIRCRAFT MASS CHARACTERISTICS	67
	5.2.1 Nondimensionalization	67
	5.2.2 Moment Reference	69
	5.2.3 Sensor Reference	71
	5.2.4 Kinematics	73
	5.2.5 Sources	73
	5.3 ATMOSPHERIC DATA	75

6.0 FLIGHT TEST MANEUVERS 81
6.1 THE FLIGHT ENVELOPE 81
6.1.1 Envelope Coverage 81
6.1.2 Envelope Expansion 81
6.1.3 Scheduling 82
6.2 SMALL-PERTURBATION MANEUVERS 82
6.2.1 Reasons for Small-Perturbation Maneuvers 82
6.2.2 Signal-to-Noise Ratio 83
6.2.3 Other Sizing Considerations 84
6.2.4 Design Constraints 84
6.2.5 Pilot Involvement 85
6.3 IDENTIFIABILITY 86
6.3.1 Independent Inputs 86
6.3.2 Modal Excitation 87
6.3.3 Frequency Content 87
6.3.4 Feedback Systems 88
6.4 SAMPLE MANEUVERS 89
6.4.1 Longitudinal Maneuvers 89
6.4.2 Lateral-Directional Maneuvers 89

7.0 DATA ACQUISITION SYSTEM 105
7.1 TYPES OF RECORDING SYSTEMS 105
7.2 TIME TAGS 106
7.3 ALIASING AND PREFILTERING 107
7.4 SAMPLE RATE 109
7.5 RESOLUTION 110

8.0 INSTRUMENTATION 123
8.1 GENERAL TRANSDUCER CHARACTERISTICS 123
8.2 CONTROL POSITIONS 124
8.3 ANGULAR RATES 125
8.4 LINEAR ACCELERATIONS 126
8.5 FLOW ANGLES 127
8.5.1 Uses 127
8.5.2 Vanes 128
8.5.3 Pressure Ports 129
8.5.4 Other Sources 129
8.6 EULER ANGLES 130
8.7 AIR DATA 131
8.8 ANGULAR ACCELERATIONS 132
8.9 ENGINE PARAMETERS 133
8.10 CONFIGURATION PARAMETERS 133
8.11 LOADING DATA 134

9.0 EVALUATION OF RESULTS 137
9.1 COMPUTATION OF THE CRAMÉR-RAO BOUND 137
9.1.1 Discrepancy in the Cramér-Rao Bound 137
9.1.2 Previous Attempts at Explanation 138
9.1.3 Explanation of the Discrepancy 139
9.1.4 Suggested Implementations 140
9.2 EXAMPLES OF APPLICATION 141
9.2.1 Example 1 141
9.2.2 Example 2 142
9.2.3 Example 3 142
9.2.4 Example 4 142
9.2.5 Example 5 142
9.2.6 Example 6 143
9.3 MODELING CONSIDERATIONS 143

10.0 CONCLUDING REMARKS 153

BIBLIOGRAPHY 155

NOMENCLATURE

In this document, all equations assume angles in radians, while the text and the figures give angles in degrees.

a_n	normal accelerometer output, g
a_x, a_y, a_z	longitudinal, lateral, and vertical accelerometer output, g
b	reference span, m or ft
C_ℓ, C_m, C_n	coefficients of roll, pitch, and yaw moment
C_L, C_D	coefficients of lift and drag
C_N, C_A	coefficients of normal and axial force
C_X, C_Y, C_Z	coefficients of longitudinal, lateral, and vertical force
c	reference chord, m or ft
F	external applied force, N or lb
F_X, F_Y, F_Z	components of external applied force, N or lb
f(.)	system state function
GG^*	measurement noise covariance
g(.)	system observation function
g	acceleration of gravity, m/sec^2 or ft/sec^2
H	altitude, m or ft
h	angular momentum vector
I_x, I_y, I_z	moments of inertia, kg-m^2 or slug-ft^2
I_{x_e}	engine moment of inertia, kg-m^2 or slug-ft^2
I_{xy}, I_{xz}, I_{yz}	cross products of inertia, kg-m^2 or slug-ft^2
J(.)	cost function
K_α, K_β	upwash and sidewash factors
L_β, L_δ	dimensional rolling moment derivatives
M	external applied moment, N-m or ft-lb
m	mass, kg or slug
m_ξ	prior mean of ξ
$M_\alpha, M_\delta, M_q, M_b$	dimensional pitching moment derivatives
N	engine speed, rad/sec
N_β, N_δ	dimensional yawing moment derivatives
P	prior covariance of ξ
p	roll rate, rad/sec or deg/sec
P_s	static pressure, N/m^2 or lb/ft^2
q	pitch rate, rad/sec or deg/sec
$\bar{q}$	dynamic pressure, N/m^2 or lb/ft^2
r	yaw rate, rad/sec or deg/sec

s	reference area, m^2 or ft^2
T	thrust, N or lb
t	time, sec
u	system input; or body X-axis wind-relative velocity, m/sec or ft/sec
V	total wind-relative velocity, m/sec or ft/sec
$\bar{V}$	velocity vector, components in m/sec or ft/sec
v	body Y-axis wind-relative velocity, m/sec or ft/sec
W_x, W_y	horizontal wind components, m/sec or ft/sec
W_v	vertical wind component, m/sec or ft/sec
w	body Z-axis wind-relative velocity, m/sec or ft/sec
X	spatial x position, m or ft
x	system state
$x_\alpha, x_\beta, x_{a_x}, x_{a_y}, x_{a_n}$	sensor X-axis positions, m or ft
Y	spatial y position, m or ft
$y_\alpha, y_\beta, y_{a_x}, y_{a_y}, y_{a_n}$	sensor Y-axis positions, m or ft
z	system response
$z_\alpha, z_\beta, z_{a_x}, z_{a_y}, z_{a_n}$	sensor Z-axis positions, m or ft
Z_α, Z_β	dimensional vertical force derivatives
α	angle of attack, rad or deg
α_f	flank angle of attack, rad or deg
β	angle of sideslip, rad or deg
γ	flightpath angle, rad or deg
Δ	sample interval, sec
δ	generic control
δ_a	aileron deflection, rad or deg
δ_e	elevator deflection, rad or deg
δ_r	rudder deflection, rad or deg
η_i	measurement noise
θ	pitch attitude, rad or deg
ϕ	roll attitude, rad or deg
ξ	vector of unknown parameters
ρ	air density, kg/m^3 or $slug/ft^3$
τ	lag time constant, sec
ψ	heading angle, rad or deg
ω	angular velocity vector, components in rad/sec

Subscripts

b	bias, or body axes
c	corrected
e	earth axes
flt	flight
m	measured
min	minimum valve
$p,r,\dot{\beta}$	subscripting an aerodynamic coefficient, derivative with respect to the subscript times b/2V
$q,\dot{\alpha}$	subscripting an aerodynamic coefficient, derivative with respect to the subscript times c/2V
ref	reference
wind	wind axes
z	observation
α,β,δ	subscripting an aerodynamic coefficient, derivative with respect to the subscript

Superscripts

~	estimated
-	average

1.0 INTRODUCTION

This document examines the practical application of parameter estimation techniques to the problem of estimating aircraft stability and control derivatives from flight test data.

The field of aircraft stability and control exemplifies a successful application of system identification technology. Workers in the industry accept and use system identification techniques on a routine production basis. There certainly are isolated problems (primarily in extending the application to more difficult situations), but there is little doubt that the basic application is highly successful. Major factors contributing to this success include the following: a well-understood, physically derived model form that is reasonably representative of the true vehicle in most flight regimes; high-quality measurements of several relevant parameters; the ability to apply inputs specifically for system identification; and engineers familiar with both system identification and aircraft stability and control.

Contrast aircraft stability and control with economics, for instance. Economic models are of dubious validity, at best. The measurements are little better than the models (for example, whose numbers for the inflation rate are best) and are often biased by political situations (like lowering the federal budget deficit by redefining methods for computing the deficit). Economic controls are not subject to manipulation for system identification experiments, even assuming that the controls are known. Some economists insist that the Federal Reserve Bank controls major portions of the economy; others insist that it has little influence and merely responds to existing conditions. (In either case, it is not practical to shut down the Federal Reserve Bank for three months to see what happens, changing nothing else for the next four years to avoid confusing the issue.)

1.1 SCOPE AND BASIS

This document does not survey the large number of parameter estimation techniques applicable to aircraft stability and control, but rather thoroughly examines a single approach. This concentration allows us to cover the entire parameter estimation process in some detail, from planning the flight test program to evaluating the results. We discuss the kinds of difficulties encountered in practice, with extensive examples using real flight data. We intend this detailed treatment to serve both as an introduction to the problem and as a guide for practicing engineers. We hope to promote a more complete understanding than that provided by texts that present only brief examples with "correct" answers, or that provided by field education, which often amounts to "Here's a computer program to solve the problem — sometimes it works."

The problems and solutions that we discuss are all real and are based on our own experience at the Dryden Flight Research Facility of NASA's Ames Research Center (Ames-Dryden). Although we sometimes illustrate the principles with simulated data, we have encountered all these issues in real situations. We omit discussions of numerous potentially important issues that have not yet troubled us. For example, although finite computer word length may be an important issue on a smaller machine (we do most of our work on a 60-bit computer) or for a different application, we have yet to encounter finite word length difficulties that did not stem from more basic sources, such as poorly written problem statements. We decline to artificially create difficulties for experimentation. Others have encountered these issues and treated them elsewhere, and we have little to add to the discussion.

By restricting our discussion to the scope of our own experience, we achieve a strength that we feel is lacking in many reports, even those purporting to be practically oriented. We offer a self-critical first-hand view of the estimation process, including numerous mundane problems encountered along the way. Most reports omit mentioning "stupid" errors and other trivial problems that are not relevant to the final reported results, but these problems can account for the majority of the time spent.

In the flight test environment, results are subject to detailed critical review. If our results disagree with predictions, someone will ask where we erred; we need to convincingly defend our results before an often skeptical audience. If we suggest that the simulator be revised based on our results, we must demonstrate why the update is worth the work (and hope the pilot notices that the revised simulator flies more like the airplane). If we suggest that instrumentation errors have occurred, someone will test it and contradict us. If we request more test data, the schedulers will complain about milestones and cost. In some flight regimes the controls and handling qualities group wants assurance that our results are very accurate because they have little margin for error; in other flight regimes they may insist that we must be wrong, because if our results are correct, the control system needs to be redesigned. Throughout this process, few people care if we have an elegant, sophisticated, and innovative method; they simply want good results and they want them immediately.

This flight test environment fosters a self-critical attitude; it is this attitude, more than any specific issue, that we try to promote in this document. It is a corrollary of Murphy's law that although there may be several ways to arrive at acceptable results, there are more ways to get incorrect results. We have seen too many papers (with impressive claims about having analyzed some complicated nonlinear problem) where cursory inspection has shown the results to be wrong (they had, as claimed, analyzed the problem, but the analysis was wrong). We certainly do not claim to be free of errors. In fact, a principal advantage of our experience is that we have made (and fixed) more mistakes than most people. We hope that this document will help achieve the goals of the practicing engineer: to avoid problems where possible, to fix the problems that cannot be avoided, and to recognize the problems that cannot be fixed.

This document emphasizes areas of mature technology. The techniques we discuss have become widely accepted, based on a broad background of practical applications. We avoid detailed discussion of

emerging areas of technology, where there are only a few isolated examples of application (even when some examples are from our own experience). Material on emerging areas rapidly becomes dated and is inappropriate for this document. Such emerging areas of technology include filter-error algorithms, optimal input design, automatic model structure determination, nonlinear and unsteady aerodynamics, and helicopter applications. There are examples of application in these areas, to both simulated and real data, but their appropriate role in routine use is not yet clear. The reference list includes numerous papers addressing these areas.

1.2 OUTPUT-ERROR METHOD

We begin by briefly reviewing the form of the output-error parameter estimation method. (See Maine and Iliff (1984) for a complete treatment of the development and implications of the following equations.)

The aircraft is a continuous-time dynamic system. We will assume that measurements are made at discrete time intervals for analysis on a digital computer. The system equations, in general form, are

$$x(t_0) = x_0(\xi) \tag{1.2-1a}$$

$$\dot{x}(t) = f[x(t),u(t),t,\xi] \tag{1.2-1b}$$

$$z(t_i) = g[x(t_i),u(t_i),t_i,\xi] + n_i \tag{1.2-1c}$$

where x is the state, u the input, and z the response of the system. We assume that the measurement noise n is a sequence of independent Gaussian random vectors with zero mean and covariance GG*. The output-error method does not account for any process noise. We assume that the forms of the f and g functions are known. The parameter estimation task is to estimate the value of the unknown-parameter vector ξ.

The output-error estimate of ξ is the value that minimizes the cost function

$$J(\xi) = \frac{1}{2}\sum_{i=1}^{N} [z(t_i) - \tilde{z}_\xi(t_i)]^*(GG^*)^{-1}[z(t_i) - \tilde{z}_\xi(t_i)] + \frac{1}{2}(\xi - m_\xi)^* P^{-1}(\xi - m_\xi) \tag{1.2-2}$$

where $\tilde{z}_\xi$ is the computed response obtained from

$$\tilde{x}_\xi(t_0) = x_0(\xi) \tag{1.2-3a}$$

$$\dot{\tilde{x}}_\xi(t) = f[\tilde{x}_\xi(t),u(t),t,\xi] \tag{1.2-3b}$$

$$\tilde{z}_\xi(t_i) = g[\tilde{x}_\xi(t_i),u(t_i),t_i,\xi] \tag{1.2-3c}$$

The value m_ξ in the cost function is the mean of the prior distribution of ξ, and P is the covariance. If there is no prior distribution to be considered, then the last term of the cost function is omitted.

The Gauss-Newton algorithm is an effective means of minimizing this cost function (Maine and Iliff, 1984).

1.3 AIRCRAFT STABILITY AND CONTROL APPLICATION

Aircraft stability and control involves controlling the attitude and flightpath of an aircraft. The desired attitude and flightpath are defined by agents external to the stability and control system: pilots, guidance and navigation systems, or autoland systems, for instance. In the simplest systems, pilots decide on a desirable attitude and flightpath and move the control wheels, sticks, pedals, and other controls in a manner that they anticipate will give the desired results. Mechanical cables connect the pilot's controls to movable control surfaces on the airframe or to engine control actuators. Moving the control surfaces changes the aerodynamic forces and moments on the aircraft, thus changing its attitude and flightpath. There is no feedback inherent in these simple control systems; the pilot provides external feedback by monitoring the aircraft response using the cockpit instruments, the view out the window, and the "feel" of the airplane, and by appropriately changing the control stick motions.

In a more complicated system, a guidance and navigation computer generates attitude or flightpath commands needed to maintain the desired trajectory. These commands are in the form of analog or digital signals, which are sent to the stability and control computer (or possibly to a different program running on the same computer). The stability and control computer determines the control surface positions needed to achieve the commanded attitude or flightpath and sends commands to electric or hydraulic surface actuators, which move the aircraft control surfaces. Sensors measure the vehicle motions and feed back this information to the stability and control computer, which modifies its control surface commands accordingly.

The descriptions of these two types of systems are oversimplified, and there are numerous other combinations and variants of the system elements. We make no attempt to cover design or analysis of control systems in this document. That is a major subject in its own right, and there are several references on the subject (Ogata, 1970). The generalizations here are intended only to establish the background and purpose of stability and control parameter estimation.

Stability and control analysis is concerned with the aircraft system at several levels of integration. The first level involves an open-loop study of how the aircraft would respond to given control surface motions. A study at this level ignores all issues of feedback; it involves only vehicle aerodynamics and kinematics. The second level adds considerations of the control system feedbacks, if any. These considerations include issues of sensor characteristics, control system logic, and actuator characteristics. The result of a study at this level is an understanding of how the vehicle and control system, operating together, respond to commands. A third level includes the pilot as part of the system. The aircraft designer must ensure that a pilot can reasonably control the aircraft.

We normally estimate the stability and control derivatives using only the open-loop, bare airframe model. We look at the control surface motions and the aircraft response. The causes of the control surface motions are irrelevant. In particular, the model applies whether or not the control surface motions are a result of feedback at some higher level. There are several related reasons for using the open-loop model. The central reason is that this is the simplest system. The simpler the system, the more likely it is that we can get good estimates. Furthermore, the feedback systems are superfluous to our estimation problem. Maine and Iliff (1984) pointed out that we need to take maximum advantage of previous information in constructing our system model. We should use parameter estimation techniques only to fill the gaps in our knowledge.

The gap that we are filling here is in our knowledge of the aircraft's aerodynamics. Relative to our knowledge of the aerodynamics, we have good models of most components of the feedback systems. We can get accurate data on most actuator and sensor characteristics from both ground and flight test data without using complicated parameter estimation techniques; in exceptional cases, we may need to include the actuator or sensor models as parts of the system identification task. If the control system analysts do not know the control logic, there are very basic problems, which system identification techniques are not likely to solve.

Another big gap is in our knowledge of the pilot model. We can make some general statements about piloting techniques, but accurate quantitative pilot models are as unattainable as good economic models. There have been attempts to use system identification techniques for pilot modeling, but we are not aware of any convincing results. We do not discuss the pilot modeling problem in this document. Whereas we exclude automatic feedbacks from our analysis because simpler techniques are appropriate, we exclude the pilot model because it makes the problem too difficult. Including dubious pilot models in our system would simply corrupt the quality of the stability and control derivative estimates.

This document also ignores, as far as possible, the response of the aircraft structure and propulsion system. The aircraft structure and propulsion system are appropriate areas for the use of system identification techniques, and several applications have been made in these areas. However, we consider them as separate applications that we will not treat extensively in this document. In some cases we cannot separate the structural, propulsion system, and aerodynamic system responses. We must then consider these issues in more detail, and the problem is considerably more difficult.

In summary, we use an open-loop model for estimating aircraft stability and control derivatives. The model ignores structural and propulsion system responses to the extent possible. The model includes kinematics and the aerodynamic forces and moments acting on the aircraft. The kinematics are well known; we estimate only the aerodynamic forces and moments.

The preceding discussions point out the importance of defining the scope of the estimation problem. A complete model of all aspects of an aircraft response would be intractable; we therefore have restricted the scope of the aircraft stability and control problem to that which we can reasonably hope to handle. The other elements of the aircraft are handled separately; designers and control system analysts can assemble the different element models. Many system identification failures stem from poorly defining the scope in the original problem statement, posing a problem that is intractable from the start.

1.4 PURPOSE

As the first step in any practical application, we must establish the purpose of the estimation. The intended use of the results influences every aspect of system identification, including model selection, choice of estimation algorithms, experiment design, and evaluation of the adequacy of the results. In this section we mention some common uses of aircraft stability and control derivative estimates. A particular application can involve several of these and other uses.

One of the earliest uses of aircraft stability and control derivative estimation was in validating wind-tunnel or analytical predictions. Although comprehensive wind-tunnel and analytical testing can give a reasonable estimate of the flight vehicle aerodynamics, there are several potential sources of misprediction, sometimes minor and sometimes major: The wind-tunnel models used for most of the testing are often slightly different from the actual flight vehicle because of last-minute configuration changes. It is difficult or impossible to precisely match actual flight conditions with scaled wind-tunnel models. Reynolds number differences are a standard explanation for discrepancies between flight and wind-tunnel results. Support system (sting) effects are almost always an issue in wind-tunnel tests; major discrepancies have been traced to sting effects in several programs (Ericsson, 1981). Minimization of support effects is still more of an art than a science. Funding and time constraints necessitate shortcuts in the wind-tunnel tests (and in the flight tests). For these and numerous other reasons, it is always wise to at least spot-check wind-tunnel and analytical predictions with flight test data, even for simple configurations.

Typical reporting practices (the tendency to emphasize positive results) play down the role of flight test validation of predictions. After all, the use of flight data in this role is a tool for finding errors. It is natural (and appropriate in many places) to emphasize the consistency of the results after all the errors are corrected rather than to discuss how the errors were found and corrected. The errors can be as simple as typographic or arithmetic mistakes. More complicated errors include failure to account for (or incorrectly accounting for) Reynolds number effects, sting effects, wind-tunnel wall effects, flexibility effects, or propwash (or jetwash) effects. In some cases the data needed to make an accurate prediction have been available before flight; they simply have not been used. For instance, often there are two disagreeing sets of wind-tunnel data, one of which is discarded based on some judgment (funding and time typically prohibit running a third independent wind-tunnel test); in some such cases we have found that the flight results agreed with the discarded set. In most of the situations mentioned in this paragraph, final reports will emphasize that the wind-tunnel data (appropriately corrected) are in good agreement with the flight data.

The tendency to emphasize agreement between data sets is appropriate if the consistency of the results is being used to support the thesis that the values are accurate. To evaluate the utility of estimating the stability and control derivatives from flight data, it is valuable to look at a comparison that is seldom published: the predicted data exactly as used before the first flight compared with the best estimates combining flight data and predictions after the flight test program is completed. Such comparisons would reveal that stability and control derivative estimation from flight data has an important role in correcting simple oversights and otherwise validating predictions. Because of this role, stability and control derivative estimation from flight data is indispensable in any major new aircraft test program, regardless of how thorough the wind-tunnel test was. In spite of the most thorough wind-tunnel and analytical test program in history, the flight data for the space shuttle exhibited significant disagreements with preflight predictions.

Do not take the preceding discussion of prediction errors as reflecting our position on the ubiquitous question of whether wind-tunnel or flight test data are better; in general, we regard the question as irrelevant. It suits our purpose in this section to briefly allude to potential problems with predictions, but we devote several subsequent chapters to the potential problems with flight data analysis. For purposes of our current discussion, the roles of flight and wind-tunnel data are symmetric: The flight data can be useful in finding problems with the wind-tunnel data, and the wind-tunnel data can be useful in finding problems with the flight data. Having the most confidence in the results requires both wind-tunnel and flight data. Thus, flight data are useful in improving confidence, even if they can be summarized by stating that they agree with predictions.

In some test programs, a complete aerodynamic data base must be built using only flight data. This can occur if there are no wind-tunnel tests or if there are such major deficiencies in the wind-tunnel data as to make it easier to discard then and use only the flight data (as, for instance, with the highly maneuverable aircraft technology (HiMAT) program; Matheny and Pangeas, 1981; Iliff and Maine, 1982). In this case much better coverage of the flight envelope is needed. It is no longer sufficient to spot-check a few flight conditions and say that the predictions look reasonable. The end uses of the flight-estimated stability and control derivatives are similar, whether they are used alone or in conjunction with predictions.

Flight estimates of stability and control derivatives play an important role in envelope expansion and safety of flight during flight test programs. As the demonstrated flight envelope of an aircraft expands, engineers examine stability and control derivative estimates for unexpected trends that might make flight unsafe in some regimes. This monitoring can be done during the flight or between flights.

Computer programs for analyzing aircraft stability, handling qualities, and control systems make direct use of stability and control derivative estimates. These programs can analyze the aircraft as it is, or they can analyze the effects of proposed changes in the control system.

High-fidelity simulators are increasingly necessary in modern flight research and test programs and in operational use. For many years, simulators have been used for flight planning and pilot training, both to save expensive flight time and to minimize risky flight operations, such as training for emergency conditions. As aircraft become more complex and as flight envelopes expand to include unconventional regimes, the need becomes greater for using simulators as integral parts of the flight test programs. Control systems and, to the extent possible, handling qualities must be tested on simulators before committing to flight. The design and refinement of complex modern control systems operating over large flight envelopes require higher-fidelity simulators than were needed for pilot training. These high-fidelity simulations require complete stability and control data that give an accurate representation of the actual flight vehicle. For many years, every major flight project at Ames-Dryden involving stability, control, or flying qualities has included a high-fidelity simulator updated with flight-determined stability and control derivatives.

For most of these purposes, we care little about the phenomenology of the aerodynamics; we need only know the end result of how the aircraft responds. Thus, we may be relatively indifferent to some model distinctions. For instance, the parameters C_{m_q} and $C_{m_{\dot{\alpha}}}$ both affect aircraft damping in pitch. In many flight regimes, the aircraft responses are identical (within the accuracy of our measurements) whether the damping arises from C_{m_q} or $C_{m_{\dot{\alpha}}}$. Therefore, if we are restricted to such flight regimes and if our purpose is solely to construct an accurate simulation, we need only estimate the sum $C_{m_q} + C_{m_{\dot{\alpha}}}$. To

attempt to separately estimate the two components would not only be a waste of time, it could well degrade the results.

For other purposes, however, the phenomenology might be of primary interest. We might be using the stability and control derivative estimates as an aid to understanding aerodynamic flow phenomena. In such a case, C_{m_q} and $C_{m_{\dot{\alpha}}}$ are quite distinct because they arise from different flow phenomena. This simple example illustrates how a choice of model form (C_{m_q} and $C_{m_{\dot{\alpha}}}$ summed or distinct) can depend on the purpose of the estimation.

1.5 FLIGHT TEST PROGRAM

In this section we give an overview of the flight test procedure for estimating stability and control derivatives. This overview also serves as an outline of the subjects covered in the rest of this document. Figure 1.5-1 illustrates the major steps of a flight test program. The individual items are elaborated below.

You need to develop a test plan that addresses all the issues mentioned in this section. This plan need not be a formal written document. It may be unreasonable to make detailed plans for the entire project before it starts. For instance, you probably will need to examine some of the raw data before making final decisions about data filtering and preprocessing. At a minimum, you need to make a mental checklist of what needs to be accomplished at each stage of the project. Otherwise it is easy to omit small but important items, such as measuring the instrument positions.

The first element of the test plan is a statement of the objectives. In Section 1.4 we discuss the importance of the objectives and the influence that the objectives can have on the approach. A statement of objectives should be concrete and should include acceptable forms of the results. For instance, it is important to know whether a single model, which describes the aircraft over its entire envelope, is required, or a set of models, each valid in a part of the envelope, is acceptable.

You must establish the requirements for predicted derivatives and a source of data to meet these requirements. You may also need to determine how to obtain the aircraft mass characteristics. Chapter 5 discusses mass characteristics and preflight predictions.

You need to define the necessary flight maneuvers, as described in Chapter 6. Chapters 7 and 8 discuss the requirements for instrumentation and data systems adequate for stability and control analysis. Establishing and implementing instrumentation requirements are among the long-lead-time processes and must be addressed early in the flight test program. Preprocessing of the measured flight data to get it in a form usable for analysis can be a major effort.

You need to define the appropriate equations and analysis methods. Basic forms and some variants of the aircraft equations of motion are discussed in Chapters 2 and 3; analysis methods are discussed throughout this document.

As the final step of the estimation process, you must evaluate and present the results. The form of presentation should be tied to the acceptable forms established in the statement of objectives. Chapter 9 discusses issues related to evaluating and presenting results.

An important facet of the entire process is the interplay between steps. Although we lay out the steps in a linear fashion, the steps can seldom be separated so neatly in practice. Results from each step influence decisions in other steps. In many cases, you will need to revise earlier decisions based on later results. For instance, unexpected trends in the estimates might justify extra instrumentation, additional maneuvers, or alternative analysis methods. Inflexibility and refusal to reevaluate previous decisions invite poor results.

Any realistic flight test plan should include generous allowances for unexpected problems. The positive-sounding term "success-oriented" philosophy is a currently popular management euphemism for ignoring problems and hoping they will go away, usually an attempt to avoid admitting that the program is over budget, behind schedule, and under performance. (A more traditional description of this philosophy is overoptimistic).

In an extreme case, the success-oriented program starts with the lowest quality data system that someone can argue should be acceptable; the test program consists of a quick series of flights with one maneuver at each of a few questioned flight conditions; then the instrumentation is removed, and the airplane is sent off for other uses before any data are analyzed. Such a plan minimizes projected cost and time in the unlikely event that nothing goes wrong. Unfortunately, there is little or no allowance for error. Problems that would have been minor in a better-planned program (perhaps only requiring an extra flight) can completely invalidate the data in a success-oriented program. For example, if the data recorder were to fail (as happened to some recorders on a few shuttle flights), more time and money could have been saved by forgetting about the tests in the first place. More likely, the data will be full of obvious inconsistencies (or worse, unobvious ones). If the data are analyzed and the results are different than predicted, the extreme success-oriented program would disregard the results anyway, hoping that an error was made or that it might not matter. If a test program does not merit better treatment than this, it is not worth doing.

Rather, flight test plans should reflect a problem-oriented philosophy. Assume that there will be problems, and design the tests to maximize the chances of finding and fixing the problems. Much of this document reflects such problem-oriented philosophy. We point out the kinds of problems that commonly occur, how to identify them, and how to fix them or work around them. Our experience is that such a problem-oriented approach is the best way to ensure success, while a success-oriented approach is the best way to ensure lasting problems.

Figure 1.5-2 shows relationships among several of the major areas of aircraft stability and control estimation: input, vehicle, measurements, and analysis. We discuss all of these in this document. The following list of AGARD documents relates to the areas depicted in Figure 1.5-2. The bibliography at the end of this document lists these and other references in the field.

A: INPUT

Optimal input design
AGARD, 1979b, No. 3
AGARD, 1975, No. 12

Flight test program design
AGARD, 1977, Nos. 5, 11, 12, 13

B: VEHICLE

Fixed-wing aircraft
AGARD, 1979b, No. 6

Rotorcraft
AGARD, 1979b, No. 7
AGARD, 1983, No. 16

Wind-tunnel and free-flight models
AGARD, 1979b, No. 10
AGARD, 1983, No. 17A

Extreme flight regimes
AGARD, 1979b, No. 8
AGARD, 1979a, Nos. 1, 14

Closed-loop aspects
AGARD, 1979b, No. 11

Aeroelastic flight testing
AGARD, 1972-1983, No. 9
AGARD, 1983, Nos. 18, 19, 20

Flight/ground testing
AGARD, 1976, Nos. 6, 8, 13
AGARD, 1983, Nos. 3B, 15, 16, 17A

Dynamic stability parameters
AGARD, 1978, Nos. 14, 15, 17, 18
AGARD, 1981, No. 10

C: MEASUREMENTS

Flight test instrumentation
AGARD, 1979b, No. 4
AGARD, 1972-1983, Vol. 1

Signal filtering
AGARD, 1972-1983

D: ANALYSIS

Maximum likelihood methods
AGARD, 1979b, No. 2
Maine and Iliff, 1984

Flightpath reconstruction
linear regression methods
AGARD, 1979b, No. 5

Frequency domain methods
AGARD, 1979b, No. 2

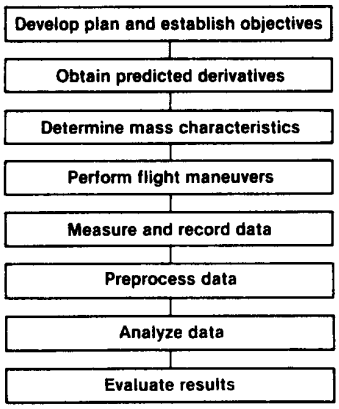

Figure 1.5-1. Steps of a flight test program.

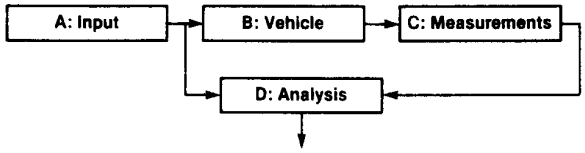

Figure 1.5-2. Aircraft stability and control estimation.

2.0 AIRCRAFT EQUATIONS OF MOTION

This chapter defines the signals pertinent to aircraft stability and control analysis and derives the aircraft equations of motion. These equations form the basis for aircraft stability and control analysis. The derivations of this chapter assume that the aerodynamic characteristics are known. Chapter 3 discusses the application of parameter identification to these equations in order to estimate the stability and control derivatives.

2.1 SIGNAL DEFINITIONS

In this section, we define the input, state, output, and other signals pertinent to the aircraft stability and control problem.

2.1.1 Control Surfaces

The usual inputs to the aircraft stability and conrol equations are the control surface positions. The basic aircraft control surfaces, illustrated in Figure 2.1-1, are the elevator (δ_e), ailerons (δ_a), and rudder (δ_r). The elevator and ailerons can be separate surfaces or can represent symmetric and antisymmetric deflections of a single pair of surfaces usually called elevons (from elevator and aileron). Control surface arrangement is highly configuration dependent: Various combinations of flaps, spoilers, canards, reaction control jets, or other control devices are possible, in addition to or instead of the basic three controls. Events, such as weight dropping, can also constitute inputs.

Our usual conventions for control surfaces are as follows: All surface positions are angular deflections measured normal to the hinge line. Positive deflection is trailing edge down or trailing edge left (port), depending on the surface orientation (as shown in Figure 2.1-1). Aileron deflections and other antisymmetric control combinations are defined as left surface position minus right surface position. These conventions are not universal; in particular, some authors define aileron deflection as one-half of our definition or with the opposite sign. Sign conventions are worth verifying on each aircraft because there is no universal standard.

2.1.2 Body-Axis System

Most of the quantities of interest to us are referenced to aircraft geometric body axes, shown in Figure 2.1-2. The origin of the body-axis system is at the center of gravity (we could work from some other origin, but placing the origin at the center of gravity significantly simplifies the model). The positive X axis points forward, out the nose of the aircraft; its exact orientation is defined for each aircraft model. The positive Y axis points to the right, and the positive Z axis points down to complete a right-handed axis system. The entire axis system moves and rotates with the aircraft.

2.1.3 Euler Angles

The aircraft attitude with respect to the earth is defined by the three Euler angles Ψ (heading angle), Θ (pitch attitude or pitch angle), and Φ (roll attitude or roll angle). These angles define the rotations that transform earth-fixed axes to aircraft body axes at any particular instant of time. The order of rotation is important. Start from earth-fixed axes, where X_1 is north, Y_1 east, and Z_1 down. Rotate first by the angle Ψ about the Z_1 axis to define the rotated axis system (X_2,Y_2,Z_2) (Z_2 is identical to Z_1). If we imagine an aircraft attached to the axis system as it rotates, zero Ψ is nose north, and positive Ψ rotation is clockwise. Then rotate by the angle Θ about the Y_2 axis to define a third axis system (X_3,Y_3,Z_3) (Y_3 is identical to Y_2). Note that the Θ rotation is about Y_2, not Y_1. Zero Θ is nose level, and positive Θ is nose up. Finally rotate by the angle Φ about the X_3 axis to arrive at the aircraft body axes (X,Y,Z) (X is identical to X_3). Zero Φ is wings level, and positive Φ is right wing down. Figure 2.1-3 illustrates the Euler angles.

These definitions allow multiple solutions for (Ψ,Θ,Φ) for any given aircraft attitude; for instance, it is easy to see that a (Ψ,Θ,Φ) of (0°,0°,0°) is equivalent to (180°,180°,180°). We normally define Θ to lie in the range (-90°,90°). The Euler angles are then unique except when Θ is exactly 90° or -90°.

The transformation of a vector from earth axes V_e to body axes V_b is the product of the three rotation matrices.

$$V_b = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\Phi & \sin\Phi \\ 0 & -\sin\Phi & \cos\Phi \end{bmatrix} \begin{bmatrix} \cos\Theta & 0 & -\sin\Theta \\ 0 & 1 & 0 \\ \sin\Theta & 0 & \cos\Theta \end{bmatrix} \begin{bmatrix} \cos\Psi & \sin\Psi & 0 \\ -\sin\Psi & \cos\Psi & 0 \\ 0 & 0 & 1 \end{bmatrix} V_e$$

$$= \begin{bmatrix} \cos\Theta\cos\Psi & \cos\Theta\sin\Psi & -\sin\Theta \\ \sin\Phi\sin\Theta\cos\Psi - \cos\Phi\sin\Psi & \sin\Phi\sin\Theta\sin\Psi + \cos\Phi\cos\Psi & \sin\Phi\cos\Theta \\ \cos\Phi\sin\Theta\cos\Psi + \sin\Phi\sin\Psi & \cos\Phi\sin\Theta\sin\Psi - \sin\Phi\cos\Psi & \cos\Phi\cos\Theta \end{bmatrix} V_e \qquad (2.1\text{-}1)$$

The inverse transformation is

$$V_e = \begin{bmatrix} \cos\Psi & -\sin\Psi & 0 \\ \sin\Psi & \cos\Psi & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos\Theta & 0 & \sin\Theta \\ 0 & 1 & 0 \\ -\sin\Theta & 0 & \cos\Theta \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\Phi & -\sin\Phi \\ 0 & \sin\Phi & \cos\Phi \end{bmatrix} V_b$$

$$= \begin{bmatrix} \cos\Psi\cos\Theta & \cos\Psi\sin\Theta\sin\Phi - \sin\Psi\cos\Phi & \cos\Psi\sin\Theta\cos\Phi + \sin\Psi\sin\Phi \\ \sin\Psi\cos\Theta & \sin\Psi\sin\Theta\sin\Phi + \cos\Psi\cos\Phi & \sin\Psi\sin\theta\cos\Phi - \cos\Psi\sin\Phi \\ -\sin\Theta & \cos\Theta\sin\Phi & \cos\Theta\cos\Phi \end{bmatrix} V_b \qquad (2.1\text{-}2)$$

2.1.4 Angular Rates

We define the body-axis aircraft angular rates as projections of the angular velocity vector (with respect to inertial space) on the body axes. This definition may seem a little awkward in that it first defines the vector in inertial space; this is because the angular velocity of the aircraft with respect to the body axes is zero by definition and thus is not a useful quantity. The roll rate p, pitch rate q, and yaw rate r are the components of the angular velocity in the body X, Y, and Z directions, respectively. The sign conventions follow a right-hand rule: Positive roll rate is right wing moving down; positive yaw rate is nose moving right; positive pitch rate is nose moving up.

2.1.5 Wind-Relative Velocity

For stability and control analysis, we are seldom interested in the velocity of the aircraft relative to earth. The velocity of the aircraft relative to the air, called the wind-relative velocity, is of central importance. It is often convenient to consider this as the velocity with which the relative wind hits the aircraft (the sign convention is then that the vector points in the direction from which the relative wind comes). These two definitions are equivalent. The components of this vector in the X, Y, and Z body axes are u, v, and w, respectively.

For many purposes, it is more convenient to express the wind-relative velocity in a spherical coordinate system instead of by the Cartesian components u, v, and w. In the spherical system, the total wind-relative velocity is

$$V = \sqrt{u^2 + v^2 + w^2} \qquad (2.1\text{-}3)$$

The angle of attack α and the angle of sideslip β are defined by

$$\alpha = \tan^{-1}\frac{w}{u} \qquad (2.1\text{-}4a)$$

$$\beta = \sin^{-1}\frac{v}{V} \qquad (2.1\text{-}4b)$$

Figure 2.1-4 shows the geometry of these definitions.

The angle of sideslip defined by Equation (2.1-4b) is not exactly the quantity measured by a normal sideslip vane (see Section 8.5). A sideslip vane measures the flank angle of attack, defined by

$$\alpha_f = \tan^{-1}\frac{v}{u} \qquad (2.1\text{-}5)$$

The flank angle of attack is related to the angle of sideslip by the equation

$$\tan\beta = \tan\alpha_f \cos\alpha \qquad (2.1\text{-}6)$$

We use the angle of sideslip in the equations derived in this document. It is possible to derive corresponding equations in terms of the flank angle of attack.

We can invert Equations (2.1-3) and (2.1-4) to give

$$u = V\cos\alpha\cos\beta \qquad (2.1\text{-}7a)$$

$$v = V\sin\beta \qquad (2.1\text{-}7b)$$

$$w = V\sin\alpha\cos\beta \qquad (2.1\text{-}7c)$$

Note that Equation (2.1-7) can be written in the form

$$\begin{bmatrix} u \\ v \\ w \end{bmatrix} = \begin{bmatrix} \cos\alpha & 0 & -\sin\alpha \\ 0 & 1 & 0 \\ \sin\alpha & 0 & \cos\alpha \end{bmatrix} \begin{bmatrix} \cos\beta & -\sin\beta & 0 \\ \sin\beta & \cos\beta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} V \\ 0 \\ 0 \end{bmatrix} \qquad (2.1\text{-}8)$$

which is a product of two rotation matrices and the velocity vector in a relative-wind-oriented coordinate system (called wind axes). The order of the two rotations is important.

2.1.6 Linear Accelerations

The definition of body-axis linear accelerations is similar to that of the body-axis angular rates. The body-axis linear accelerations of any point on the aircraft are the projections of the acceleration vector of the point (with respect to inertial space) on the body axes. Note that the accelerations, unlike the angular rates, are different for different points on a rigid aircraft. When the term acceleration is not qualified with a position, it usually refers to acceleration of the center of gravity.

Linear accelerometers are important sensors for stability and control data analysis and for control system feedback. Such accelerometers measure the body-axis accelerations, excluding the component of acceleration due to gravity or other internal forces. Equivalently, we could say that an accelerometer measures the externally applied force on the accelerometer case. Thus, an accelerometer will indicat 1 g of upward acceleration for an aircraft in steady level flight. The actual acceleration in steady level flight is zero, composed of 1 g of upward acceleration from lift plus 1 g of downward acceleration from gravity; an accelerometer will measure only the acceleration from lift in this situation. For many purposes, accelerometer measurements are actually more useful than true acceleration measurements would be; if we had measurements of the true acceleration, we would often need to subtract the gravity contribution.

Aeronautical engineers (including ourselves) often inaccurately use the term "accelerations" for the linear accelerometer outputs, implying that the second derivative of position is the acceleration plus the gravity contribution. The usage is sloppy, but there is not an accepted better term. Terms like "accelerometer output" are awkward; the term "specific force" is used occasionally but is not universal.

The three components of accelerometer output are a_n, a_x, and a_y. The a_n component, called normal acceleration, is positive upward; note that this is in the negative Z direction. Occasionally you will see a_z, positive down, substituted for a_n in order to make the accelerometers consistent in sign with the axis system. However, the use of a_n is far more common, lift always being thought of as positive upward, and we adopt a_n in this document. The a_x component, called longitudinal acceleration, is positive forward, along the X axis. The a_y component, called lateral acceleration, is positive to the right, along the Y axis.

2.2 SIX-DEGREE-OF-FREEDOM EQUATIONS OF MOTION

In this section we briefly outline the derivation of the aircraft equations of motion. Most aircraft dynamics texts (Etkin, 1959) give more detailed versions of these derivations, starting from the equations for a point mass. Gainer and Hoffman (1972) present a more general set of equations than we use here.

2.2.1 Newtonian Mechanics

We take as a starting point the nonrelativistic, rigid-body equations of motion in a nonrotating inertial axis system. This coordinate system could be fixed with respect to either the earth or the air (assuming that the wind velocity is constant in both time and space); for our purposes here, it is best to use a coordinate system fixed relative to the air. The basic equations are the linear momentum and angular momentum equations:

$$F = \frac{d}{dt}(m\bar{V}) \tag{2.2-1}$$

$$M = \frac{d}{dt}(h) \tag{2.2-2}$$

where F is the external applied force, M the external applied moment about the center of gravity, $\bar{V}$ the velocity vector, and h the angular momentum vector about the center of gravity. Nonrotating coordinates are somewhat awkward because our measurements are made primarily in the rotating body-axis system; furthermore, the inertia tensor is a rapidly changing function of time in the nonrotating system.

We need to transform Equations (2.2-1) and (2.2-2) to the rotating aircraft body-axis system. Let ω be the angular velocity vector of the body-axis system with respect to inertial space. The rules for transforming vector derivatives into rotating coordinate systems give

$$F = \frac{\delta}{\delta t}(m\bar{V}) + \omega \times (m\bar{V}) \tag{2.2-3}$$

$$M = \frac{\delta}{\delta t}(h) + \omega \times h \tag{2.2-4}$$

where all quantities are in the rotating body-axis system. The $\delta/\delta t$ operator denotes the vector of derivatives of the vector components; in a rotating axis system this is not the same as the derivative of the vector, because the components in the rotating axis system change due to this rotation even if the vector is constant in inertial space.

The angular momentum is given by

$$h = \begin{bmatrix} I_x & -I_{xy} & -I_{xz} \\ -I_{xy} & I_y & -I_{yz} \\ -I_{xz} & -I_{yz} & I_z \end{bmatrix} \omega \tag{2.2-5}$$

The matrix in Equation (2.2-5) is the inertia tensor (Etkin, 1959) expressed in the body-axis system. The components of ω in the body-axis system are (p,q,r) by definition. The components of V in the body-axis system are (u,v,w) by definition. For aircraft stability and control applications, we can neglect time derivatives of m and the inertia tensor. With these substitutions, we can write Equations (2.2-3) and (2.2-4) in scalar form as

$$F_X = m(\dot{u} + qw - rv) \tag{2.2-6a}$$

$$F_Y = m(\dot{v} + ru - pw) \tag{2.2-6b}$$

$$F_Z = m(\dot{w} + pv - qu) \tag{2.2-6c}$$

$$M_X = \dot{p}I_x - \dot{q}I_{xy} - \dot{r}I_{xz} + qr(I_z - I_y) + (r^2 - q^2)I_{yz} - pqI_{xz} + rpI_{xy} \tag{2.2-7a}$$

$$M_Y = -\dot{p}I_{xy} + \dot{q}I_y - \dot{r}I_{yz} + rp(I_x - I_z) + (p^2 - r^2)I_{xz} - qrI_{xy} + pqI_{yz} \tag{2.2-7b}$$

$$M_Z = -\dot{p}I_{xz} - \dot{q}I_{yz} + \dot{r}I_z + pq(I_y - I_x) + (q^2 - p^2)I_{xy} - rpI_{yz} + qrI_{xz} \tag{2.2-7c}$$

where F_X, F_Y, and F_Z are the components of the external applied forces; M_X, M_Y, and M_Z are the components of the external applied moments. The angular velocities in these equations are in radians per second.

2.2.2 External Forces and Moments

Let us now examine the external forces and moments. We consider two components of the external moment (for more complete equations, see Gainer and Hoffman (1972)). The most significant component is the aerodynamic moment. We write the aerodynamic moments in terms of the nondimensional coefficients (Etkin, 1959):

$$M_{X_{aero}} = \bar{q}sbC_\ell \tag{2.2-8a}$$

$$M_{Y_{aero}} = \bar{q}scC_m \tag{2.2-8b}$$

$$M_{Z_{aero}} = \bar{q}sbC_n \tag{2.2-8c}$$

where $\bar{q}$ is the dynamic pressure, s the reference area, c the reference chord, b the reference span, and C_ℓ, C_m, and C_n are the coefficients of rolling, pitching, and yawing moments. These coefficients are functions of the aircraft state (see Section 2.4).

The other component of the moment that we consider is the gyroscopic moment from the rotating machinery in the engine. We must consider this as an external applied moment because Equations (2.2-3) to (2.2-5) assume that the vehicle is a rigid body with no internal moving parts. The engine gyroscopic coupling is significant for some aircraft. The following equations assume that the engine is oriented along the X axis. Generalization to arbitrary orientation is easy. The total applied moments are

$$M_X = \bar{q}sbC_\ell \tag{2.2-9a}$$

$$M_Y = \bar{q}scC_m + NrI_{xe} \tag{2.2-9b}$$

$$M_Z = \bar{q}sbC_n - NqI_{xe} \tag{2.2-9c}$$

where I_{xe} is the moment of inertia of the rotating mass of the engine and N the engine speed in radians per second.

We consider three components of the applied forces. The first component is the aerodynamic forces; in terms of the nondimensional coefficients they are

$$F_{X_{aero}} = \bar{q}sC_X \tag{2.2-10a}$$

$$F_{Y_{aero}} = \bar{q}sC_Y \tag{2.2-10b}$$

$$F_{Z_{aero}} = \bar{q}sC_Z \tag{2.2-10c}$$

where C_X, C_Y, and C_Z are the coefficients of X force, Y force, and Z force. These coefficients are functions of the aircraft state (see Section 2.4). It is also common to work in terms of the coefficient of normal force C_N and the coefficient of axial force C_A defined as

$$C_N = -C_Z \tag{2.2-11a}$$

$$C_A = -C_X \tag{2.2-11b}$$

The second force component is gravity, which exerts a force mg along the earth Z axis. Using Equation (2.1-1), we can transform the gravity force into body axes:

$$F_{X_{grav}} = -mg \sin \Theta \tag{2.2-12a}$$

$$F_{Y_{grav}} = mg \sin \Phi \cos \Theta \tag{2.2-12b}$$

$$F_{Z_{grav}} = mg \cos \Phi \cos \Theta \tag{2.2-12c}$$

The third force component is the engine thrust T, which we assume is along the body X axis. Generalization to arbitrary orientation is easy. The total applied forces are then

$$F_X = \bar{q}sC_X - mg \sin \Theta + T \tag{2.2-13a}$$

$$F_Y = \bar{q}sC_Y + mg \sin \Phi \cos \Theta \tag{2.2-13b}$$

$$F_Z = \bar{q}sC_Z + mg \cos \Phi \cos \Theta \tag{2.2-13c}$$

2.2.3 Euler Angles

Since body-axis forces are functions of the Euler angles, we will need the equations for evolution of the Euler angles. Assuming a flat, nonrotating earth, the total angular velocity of the aircraft can be expressed as the sum of $\dot{\Psi}$, $\dot{\Theta}$, and $\dot{\Phi}$ components. These components are not mutually orthogonal: The $\dot{\Psi}$ component is in the earth-axis Z direction, the $\dot{\Theta}$ component is in the Y_2 direction, and the $\dot{\Phi}$ component is in the body-axis X direction. We transform all these components into the body-axis system and equate the sum to the body-axis components of the angular velocity as follows (see Equation (2.1-1)):

$$\begin{bmatrix} p \\ q \\ r \end{bmatrix} = \begin{bmatrix} \dot{\Phi} \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \Phi & \sin \Phi \\ 0 & -\sin \Phi & \cos \Phi \end{bmatrix} \begin{bmatrix} 0 \\ \dot{\Theta} \\ 0 \end{bmatrix} + \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \Phi & \sin \Phi \\ 0 & -\sin \Phi & \cos \Phi \end{bmatrix} \begin{bmatrix} \cos \Theta & 0 & -\sin \Theta \\ 0 & 1 & 0 \\ \sin \Theta & 0 & \cos \Theta \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ \dot{\Psi} \end{bmatrix}$$

$$= \begin{bmatrix} 1 & 0 & -\sin \Theta \\ 0 & \cos \Phi & \sin \Phi \cos \Theta \\ 0 & -\sin \Phi & \cos \Phi \cos \Theta \end{bmatrix} \begin{bmatrix} \dot{\Phi} \\ \dot{\Theta} \\ \dot{\Psi} \end{bmatrix} \tag{2.2-14}$$

Inverting this equation gives

$$\dot{\Phi} = p + q \tan \Theta \sin \Phi + r \tan \Theta \cos \Phi \tag{2.2-15a}$$

$$\dot{\Theta} = q \cos \Phi - r \sin \Phi \tag{2.2-15b}$$

$$\dot{\Psi} = r \cos \Phi \sec \Theta + q \sin \Phi \sec \Theta \tag{2.2-15c}$$

2.2.4 Polar Coordinate Velocity Form

For many purposes, it is more convenient to have the equations in terms of α, β, and V than in terms of u, v, and w. The aerodynamic forces and moments are easier to visualize and express in terms of α, β, and V. Furthermore, we can directly measure flow angles closely related to α and β. We do almost all of our analysis in the α, β and V coordinates. The primary disadvantage of the (α,β,V) system is that it is singular at zero velocity, where α and β are not defined; therefore the (α,β,V) system is inappropriate for hover conditions. The (α,β,V) system also has singularities at β of $\pm 90°$, but these are seldom of concern.

To derive the (α,β,V) equations, first differentiate Equations (2.1-3) and (2.1-4) to obtain

$$\dot{V} = \frac{1}{V}(u\dot{u} + v\dot{v} + w\dot{w}) \tag{2.2-16a}$$

$$\dot{\alpha} = \frac{u\dot{w} - w\dot{u}}{u^2 + w^2} \tag{2.2-16b}$$

$$\dot{\beta} = \frac{V\dot{v} - v\dot{V}}{V^2\sqrt{(1 - v^2/V^2)}} = \frac{(u^2 + w^2)\dot{v} - vw\dot{w} - vu\dot{u}}{V^2\sqrt{(u^2 + w^2)}} \tag{2.2-16c}$$

Substituting for $\dot{u}$, $\dot{v}$, and $\dot{w}$ from Equation (2.2-6) and for u, v, and w from Equation (2.1-7) gives

$$\dot{V} = \frac{F_X}{m}\cos\alpha\cos\beta + \frac{F_Y}{m}\sin\beta + \frac{F_Z}{m}\sin\alpha\cos\beta \tag{2.2-17a}$$

$$\dot{\alpha} = \frac{1}{mV\cos\beta}(F_Z\cos\alpha - F_X\sin\alpha) + q - (\tan\beta)(p\cos\alpha + r\sin\alpha) \tag{2.2-17b}$$

$$\dot{\beta} = \frac{\cos\beta}{mV}F_Y + p\sin\alpha - r\cos\alpha - \frac{\sin\beta}{mV}(F_Z\sin\alpha + F_X\cos\alpha) \tag{2.2-17d}$$

Substituting from Equation (2.2-13) for F_X, F_Y, and F_Z gives

$$\dot{V} = -\frac{\bar{q}s}{m}C_{D_{wind}} + g(\cos\Phi\cos\Theta\sin\alpha\cos\beta + \sin\Phi\cos\Theta\sin\beta - \sin\Theta\cos\alpha\cos\beta) + \frac{T}{m}\cos\alpha\cos\beta \tag{2.2-18a}$$

$$\dot{\alpha} = -\frac{\bar{q}s}{mV\cos\beta}C_L + q - (\tan\beta)(p\cos\alpha + r\sin\alpha) + \frac{g}{V\cos\beta}(\cos\Phi\cos\Theta\cos\alpha + \sin\Theta\sin\alpha) - \frac{T\sin\alpha}{mV\cos\beta} \tag{2.2-18b}$$

$$\dot{\beta} = \frac{\bar{q}s}{mV}C_{Y_{wind}} + p\sin\alpha - r\cos\alpha + \frac{g}{V}\cos\beta\sin\Phi\cos\Theta + \frac{\sin\beta}{V}(g\cos\alpha\sin\Theta - g\sin\alpha\cos\Phi\cos\Theta + \frac{T}{m}\cos\alpha) \tag{2.2-18c}$$

where the stability-axis force coefficients are

$$C_L = -C_Z\cos\alpha + C_X\sin\alpha \tag{2.2-19a}$$

$$C_D = -C_X\cos\alpha - C_Z\sin\alpha \tag{2.2-19b}$$

and to simplify the form of the equations, we have defined the wind-axis coefficients

$$C_{D_{wind}} = C_D\cos\beta - C_Y\sin\beta \tag{2.2-20a}$$

$$C_{Y_{wind}} = C_Y\cos\beta + C_D\sin\beta \tag{2.2-20b}$$

2.2.5 Collected Equations

We have now derived the body-axis six-degree-of-freedom equations of motion. We have two equivalent forms of the equations, depending on whether the air-relative velocity is expressed in rectangular or polar coordinates. For the rectangular coordinate form, collect Equations (2.2-6), (2.2-7), and (2.2-15), substituting from Equations (2.2-9) and (2.2-13).

$$\dot{u} = \frac{\bar{q}s}{m} C_X - qw + rv - g \sin \Theta + \frac{T}{m} \tag{2.2-21a}$$

$$\dot{v} = \frac{\bar{q}s}{m} C_Y - ru + pw + g \sin \Phi \cos \Theta \tag{2.2-21b}$$

$$\dot{w} = \frac{\bar{q}s}{m} C_Z - pv + qu + g \cos \Phi \cos \Theta \tag{2.2-21c}$$

$$\dot{p}I_x - \dot{q}I_{xy} - \dot{r}I_{xz} = \bar{q}sbC_\ell + qr(I_y - I_z) + (q^2 - r^2)I_{yz} + pqI_{xz} - rpI_{xy} \tag{2.2-21d}$$

$$\dot{p}I_{xy} + \dot{q}I_y - \dot{r}I_{yz} = \bar{q}scC_m + rp(I_z - I_x) + (r^2 - p^2)I_{xz} + qrI_{xy} - pqI_{yz} + NrI_{xe} \tag{2.2-21e}$$

$$-\dot{p}I_{xz} - \dot{q}I_{yz} + \dot{r}I_z = \bar{q}sbC_n + pq(I_x - I_y) + (p^2 - q^2)I_{xy} + rpI_{yz} - qrI_{xz} - NqI_{xe} \tag{2.2-21f}$$

$$\dot{\Phi} = p + q \tan \Theta \sin \Phi + r \tan \Theta \cos \Phi \tag{2.2-21g}$$

$$\dot{\Theta} = q \cos \Phi - r \sin \Phi \tag{2.2-21h}$$

$$\dot{\Psi} = r \cos \Phi \sec \Theta + q \sin \Phi \sec \Theta \tag{2.2-21i}$$

For the polar coordinate form, use Equation (2.2-18) in place of Equations (2.2-21a) to (2.2-21c), giving

$$\dot{V} = -\frac{\bar{q}s}{m} C_{D_{wind}} + g(\cos \Phi \cos \Theta \sin \alpha \cos \beta + \sin \Phi \cos \Theta \sin \beta - \sin \Theta \cos \alpha \cos \beta) + \frac{T}{m} \cos \alpha \cos \beta \tag{2.2-22a}$$

$$\dot{\alpha} = -\frac{\bar{q}s}{mV \cos \beta} C_L + q - (\tan \beta)(p \cos \alpha + r \sin \alpha) + \frac{g}{V \cos \beta}(\cos \Phi \cos \Theta \cos \alpha + \sin \Theta \sin \alpha) - \frac{T \sin \alpha}{mV \cos \beta} \tag{2.2-22b}$$

$$\dot{\beta} = \frac{\bar{q}s}{mV} C_{Y_{wind}} + p \sin \alpha - r \cos \alpha + \frac{g}{V} \cos \beta \sin \Phi \cos \Theta + \frac{\sin \beta}{V}(g \cos \alpha \sin \Theta - g \sin \alpha \cos \Phi \cos \Theta + \frac{T}{m} \cos \alpha) \tag{2.2-22c}$$

$$\dot{p}I_x - \dot{q}I_{xy} - \dot{r}I_{xz} = \bar{q}sbC_\ell + qr(I_y - I_z) + (q^2 - r^2)I_{yz} + pqI_{xz} - rpI_{xy} \tag{2.2-22d}$$

$$-\dot{p}I_{xy} + \dot{q}I_y - \dot{r}I_{yz} = \bar{q}scC_m + rp(I_z - I_x) + (r^2 - p^2)I_{xz} + qrI_{xy} - pqI_{yz} + NrI_{xe} \tag{2.2-22e}$$

$$-\dot{p}I_{xz} - \dot{q}I_{yz} + \dot{r}I_z = \bar{q}sbC_n + pq(I_x - I_y) + (p^2 - q^2)I_{xy} + rpI_{yz} - qrI_{xz} - NqI_{xe} \tag{2.2-22f}$$

$$\dot{\Phi} = p + q \tan \Theta \sin \Phi + r \tan \Theta \cos \Phi \tag{2.2-22g}$$

$$\dot{\Theta} = q \cos \Phi - r \sin \Phi \tag{2.2-22h}$$

$$\dot{\Psi} = r \cos \Phi \sec \Theta + q \sin \Phi \sec \Theta \tag{2.2-22i}$$

Either of these systems of nine coupled nonlinear differential equations describes the aircraft motion. The inertia tensor is always an invertible matrix (it is usually nearly diagonal), and thus Equations (2.2-21d) to (2.2-21f) are solvable for the derivatives of p, q, and r. The explicit inversion is messy, and it is more convenient to write the equations in the form shown here and use a numerical inversion routine for the solution.

The equations assume nonrelativistic mechanics, a rigid vehicle, and a flat, nonrotating earth. The equations are valid in a constant wind but do not account for wind shears or fluctuations. The time rate of change of mass and inertia is assumed negligible, and fuel sloshing effects are ignored. There are no small-angle approximations, but the equations have singularities at $\Theta = \pm 90°$. The polar coordinate form also has singularities at zero velocity and at $\beta = \pm 90°$. Engine inertia and thrust terms are included, assuming that the engine alignment and thrust vectors are along the X axis. The equations are given in terms of body axes referenced to the vehicle center of gravity.

The force and moment coefficients are functions of the aircraft state. We must know these functions before we can integrate the equations, but we will put off that issue until Section 2.4.

2.2.6 Spatial Position

Each set of equations in the previous section included only 9 of the 12 differential equations describing the rigid-body motion of the aircraft. Three differential equations for the spatial position are needed for a complete set. Although the spatial position equations are seldom relevant to aircraft stability and control analysis, we present them here for completeness. These equations are necessary for simulations that include navigational considerations. In special situations, such as in ground effect or with an autopilot driven by navigational signals, the spatial position equations can also become pertinent to stability and control.

To derive the spatial position equations, simply use Equation (2.1-2) to transform the body-axis velocity (u,v,w) into earth axes. We will also allow for a constant wind. Let X and Y be the position in the earth-fixed X and Y coordinates. For the third spatial position variable, we use altitude H, which increases in the negative Z direction. The resulting (X,Y,H) coordinate system is not right-handed and therefore invites sign errors. Unfortunately, the (X,Y,Z) coordinate system also invites sign errors from people falsely assuming the positive Z direction is up, which seems natural. The equations are

$$\dot{X} = u \cos \Psi \cos \Theta + v(\cos \Psi \sin \Theta \sin \Phi - \sin \Psi \cos \Phi) + w(\cos \Psi \sin \Theta \cos \Phi + \sin \Psi \sin \Phi) + W_x \tag{2.2-23a}$$

$$\dot{Y} = u(\sin \Psi \cos \Theta) + v(\sin \Psi \sin \Theta \sin \Phi + \cos \Psi \cos \Phi) + w(\sin \Psi \sin \Theta \cos \Phi - \cos \Psi \sin \Phi) + W_y \tag{2.2-23b}$$

$$\dot{H} = u \sin \theta - v \cos \Theta \sin \Phi - w \cos \Theta \cos \Phi + W_v \tag{2.2-23c}$$

where W_x and W_y are the wind components blowing toward the north and east, respectively, and W_v is the vertical wind component, positive for updrafts. If the polar coordinate form of the wind-relative velocities is used, then substitute from Equation (2.1-7) for the u, v, and w in Equation (2.2-23) to obtain

$$\dot{X} = V \cos \alpha \cos \beta \cos \Psi \cos \Theta + V(\sin \beta)(\cos \Psi \sin \Theta \sin \Phi - \sin \Psi \cos \Phi) + V(\sin \alpha \cos \beta)(\cos \Psi \sin \Theta \cos \Phi + \sin \Psi \sin \Phi) + W_x \tag{2.2-24a}$$

$$\dot{Y} = V \cos \alpha \cos \beta \sin \Psi \cos \Theta + V(\sin \beta)(\sin \Psi \sin \Theta \sin \Phi + \cos \Psi \cos \Phi) + V(\sin \alpha \cos \beta)(\sin \Psi \sin \Theta \sin \Phi - \cos \Psi \sin \Phi) + W_y \tag{2.2-24b}$$

$$\dot{H} = V \cos \alpha \cos \beta \sin \Theta - V \sin \beta \cos \Theta \sin \Phi - V \sin \alpha \cos \beta \cos \Theta \cos \Phi + W_v \tag{2.2-24c}$$

A common quantity related to spatial position is the flight path angle γ, defined by

$$\sin \gamma = \frac{\dot{H}}{\sqrt{\dot{X}^2 + \dot{Y}^2 + \dot{H}^2}} \tag{2.2-25}$$

2.3 OBSERVATION EQUATIONS

The measurable aircraft response variables are α, β, V, p, q, r, Θ, Φ, Ψ, a_n, a_x, a_y, $\dot{p}$, $\dot{q}$, $\dot{r}$, X, Y, and H. The spatial positions (X,Y, and H) are of little relevance to most stability and control analysis, so we usually omit them from our equations. (There are exceptions, most notably for H.) We also usually ignore the V response equation, instead using measured velocity. The simplest set of observation equations assumes that the instruments are at the center of gravity (or the measurements have been corrected to it), are perfectly aligned and calibrated, and have no dynamics. It also ignores the distinction between the flank angle of attack and the angle of sideslip. We then have

$$\alpha_z = \alpha \tag{2.3-1a}$$

$$\beta_z = \beta \tag{2.3-1b}$$

$$p_z = p \tag{2.3-1c}$$

$$q_z = q \tag{2.3-1d}$$

$$r_z = r \tag{2.3-1e}$$

$$\Theta_z = \Theta \tag{2.3-1f}$$

$$\Phi_z = \Phi \tag{2.3-1g}$$

$$\Psi_z = \Psi \tag{2.3-1h}$$

$$a_n = \frac{\bar{q}s}{mg} C_N \tag{2.3-1i}$$

$$a_x = -\frac{\bar{q}s}{mg} C_A + \frac{T}{mg} \tag{2.3-1j}$$

$$a_y = \frac{\bar{q}s}{mg} C_Y \tag{2.3-1k}$$

$$\dot{p}_z = \dot{p} \tag{2.3-1ℓ}$$

$$\dot{q}_z = \dot{q} \tag{2.3-1m}$$

$$\dot{r}_z = \dot{r} \tag{2.3-1n}$$

where the subscript z distinguishes observations (elements of the z vector of Equation (1.2-1c)) from corresponding state variables.

We seldom have sensors exactly at the center of gravity (particularly flow-angle sensors, which must be exposed to the airflow). For a rigid aircraft, the sensed attitudes, rates, and angular accelerations are independent of the sensor position, but the sensed flow angles and linear accelerations vary with sensor position. The effects of sensor position can be included in the observation equations as shown in Equation (2.3-2).

Calibration bias and slope errors, known or unknown, can also be included in the observation equations. For the most part, we advise against using parameter estimation for estimating calibration slope errors. Calibrating the instruments in the lab gives more accurate data and avoids identifiability problems. For the following equations, we allow unknown biases on all measurements.

The observation equations, with arbitrary instrument positions, scale factors on angles of attack and sideslip, and biases on all measurements, are

$$\alpha_z = K_\alpha\left(\alpha - \frac{x_\alpha}{V} q + \frac{y_\alpha}{V} p\right) + \alpha_b \tag{2.3-2a}$$

$$\beta_z = K_\beta\left(\beta - \frac{z_\beta}{V} p + \frac{x_\beta}{V} r\right) + \beta_b \tag{2.3-2b}$$

$$p_z = p + p_b \tag{2.3-2c}$$

$$q_z = q + q_b \tag{2.3-2d}$$

$$r_z = r + r_b \tag{2.3-2e}$$

$$\Theta_z = \Theta + \Theta_b \tag{2.3-2f}$$

$$\Phi_z = \Phi + \Phi_b \tag{2.3-2g}$$

$$\Psi_z = \Psi + \Psi_b \tag{2.3-2h}$$

$$a_n = \frac{\bar{q}s}{mg} C_N + \frac{x_{a_n}}{g} \dot{q} + \frac{z_{a_n}}{g} (q^2 + p^2) - \frac{y_{a_n}}{g} \dot{p} + a_{n_b} \tag{2.3-2i}$$

$$a_x = -\frac{\bar{q}s}{mg} C_A + \frac{z_{a_x}}{g} \dot{q} - \frac{x_{a_x}}{g} (q^2 + r^2) - \frac{y_{a_x}}{g} \dot{r} + \frac{T}{mg} + a_{x_b} \tag{2.3-2j}$$

$$a_y = \frac{\bar{q}s}{mg} C_Y - \frac{z_{a_y}}{g} \dot{p} + \frac{x_{a_y}}{g} \dot{r} - \frac{y_{a_y}}{g} (p^2 + r^2) + a_{y_b} \tag{2.3-2k}$$

$$\dot{p}_z = \dot{p} + \dot{p}_b \tag{2.3-2ℓ}$$

$$\dot{q}_z = \dot{q} + \dot{q}_b \tag{2.3-2m}$$

$$\dot{r}_z = \dot{r} + \dot{r}_b \tag{2.3-2n}$$

The subscript b indicates biases ($\dot{q}_b$ is the bias of $\dot{q}$, not the derivative of q_b). Note that $\dot{p}_z$ is the observation of $\dot{p}$, not the derivative of p_z (there is a difference because of the biases). The quantities x_α, y_α, x_β, z_β, x_{a_n}, y_{a_n}, z_{a_n}, x_{a_x}, y_{a_x}, z_{a_x}, x_{a_y}, y_{a_y}, and z_{a_y} are the instrument positions relative to the center of gravity. The K_α and K_β parameters are upwash factors on the angles of attack and sideslip. It is debatable whether the position corrections for α and β should be multiplied by the upwash factors or not; neither alternative is completely correct, but the effect is second order. The angular rate corrections to α and β in Equation (2.3-2) use small-angle approximations, which are adequate for most situations.

Note that β in these equations is angle of sideslip, not the flank angle of attack measured by typical β vanes. If a β vane measuring flank angle of attack is used, substitute Equation (2.1-6) into Equation (2.3-2b) and use a small-angle approximation to obtain

$$\alpha_f = \frac{K_\beta}{\cos \alpha} \left(\beta - \frac{z_\beta}{V} p + \frac{x_\beta}{V} r \right) + \beta_b \tag{2.3-3}$$

Alternatively, the measurements can be preprocessed, multiplying the vane reading by $\cos \alpha$ to obtain an equivalent measurement of β.

There are many other terms that can enter the observation equations in special situations. For instance, the observation equations can be modified to reflect sensor misalignments or cross-axis sensitivity, known or unknown (preferably known). Known misalignments can alternatively be corrected in the data preprocessing; where to handle such corrections is largely a matter of convenience.

Instrumentation lags are a major issue in some projects. Such lags can be incorporated into the system model by adding extra lag states. For instance, a lag in the sensed angle of attack might be modeled by the equations

$$\dot{\alpha}_\ell = \frac{1}{\tau}\left[K_\alpha\left(\alpha - \frac{x_\alpha}{V} q + \frac{y_\alpha}{V} p\right) + \alpha_b - \alpha_\ell\right] \tag{2.3-4a}$$

$$\alpha_z = \alpha_\ell \tag{2.3-4b}$$

where α_ℓ stands for α lagged and τ is a time constant, known or unknown. In most cases it is preferable to precorrect sensor lags rather than introduce the complication of lag states.

If the (u,v,w) form of the state equations is used instead of the (α,β,V) form, replace the α, β, and V in the observation equations by substitutions from Equations (2.1-3) to (2.1-5). Alternatively, u, v, and w "measurements" can be precomputed from the measured α, β, and V. We discuss this implementation issue in Chapter 3.

2.4 AERODYNAMIC MODELS

The coefficients C_A, C_N, C_Y, C_ℓ, C_m, and C_n in the equations of motion are functions of the aircraft states and controls. Estimating these functions is the primary objective of the stability and control tests. In general, the functions are nonlinear, but in this section we largely restrict ourselves to simple linearized forms of the functions adequate for small perturbations about stabilized conditions in nonseparated flow. (Etkin (1959) discusses the aerodynamics behind the equations.)

For the simplest equations, we assume that the longitudinal aerodynamic coefficients are functions only of angle of attack, pitch rate, and longitudunal control surface positions; likewise, we assume that the lateral-directional aerodynamic coefficients are functions only of angle of sideslip, roll rate, yaw rate, and lateral-directional control surface positions. The local linear expansions of the longitudinal coefficients are then

$$C_N = C_{N_\alpha}\alpha + C_{N_q}\frac{qc}{2V} + C_{N_\delta}\delta + C_{N_b} \tag{2.4-1a}$$

$$C_A = C_{A_\alpha}\alpha + C_{A_q}\frac{qc}{2V} + C_{A_\delta}\delta + C_{A_b} \tag{2.4-1b}$$

$$C_m = C_{m_\alpha}\alpha + C_{m_q}\frac{qc}{2V} + C_{m_\delta}\delta + C_{m_b} \tag{2.4-1c}$$

and those of the lateral-directional coefficients are

$$C_Y = C_{Y_\beta}\beta + C_{Y_p}\frac{pb}{2V} + C_{Y_r}\frac{rb}{2V} + C_{Y_\delta}\delta + C_{Y_b} \tag{2.4-2a}$$

$$C_\ell = C_{\ell_\beta}\beta + C_{\ell_p}\frac{pb}{2V} + C_{\ell_r}\frac{rb}{2V} + C_{\ell_\delta}\delta + C_{\ell_b} \tag{2.4-2b}$$

$$C_n = C_{n_\beta}\beta + C_{n_p}\frac{pb}{2V} + C_{n_r}\frac{rb}{2V} + C_{n_\delta}\delta + C_{n_b} \tag{2.4-2c}$$

where, following common notation for aircraft stability and control, the subscripts indicate a derivative with respect to the subscript quantity (for instance C_{m_α} is the derivative of C_m with respect to α). Etkin (1959) discusses the reasons for the c/2V and b/2V factors in the rate terms; the notation is inconsistent, but it is standard. The δ in these equations is a generic notation for controls; add one term of the form shown for each relevant control. The subscripted quantities in these equations are called the stability and control derivatives and are the primary parameters to be estimated.

The subscript b indicates biases. These biases are closely related to, but not identical to, the often-used 0-subscript coefficients. For example, C_{N_b} is related to C_{N_0} as commonly defined. The C_{N_0} notation is sometimes used for both quantities, but that invites confusion. Figure 2.4-1 illustrates the relationship in one dimension (α dependence only). The coefficient C_{N_0} is the value of C_N at $\alpha = 0$ and $\delta = 0$. (Some authors (Etkin, 1959) alternatively define perturbation equations in which C_{N_0} is the value of C_N at a reference α and δ.) The coefficient C_{N_b} is a linear extrapolation from the average α and δ of the maneuver to the zero point. The quantities C_{N_0} and C_{N_b} are equal if C_N is linear between zero and the conditions of the maneuver. Corresponding comments apply to the other b-subscript coefficients.

Equations (2.4-1) and (2.4-2) assume that the aircraft is symmetric and flying at zero sideslip and that any lateral-directional motions are small enough to have negligible effect on the longitudinal aerodynamics, and conversely. We therefore neglect the dependence of the longitudinal aerodynamic coefficients on lateral-directional parameters and the dependence of the lateral-directional aerodynamic coefficients on longitudinal parameters. If these assumptions are violated, we can add cross derivatives to the expansions. In adding lateral-directional terms to the longitudinal expansions, we must consider whether the modeled effects should be symmetric, antisymmetric, or neither. We would expect longitudinal β effects to be symmetric for symmetric aircraft in most situations (there are exceptions; Orlik-Rückemann, 1977). Therefore, derivatives with respect to $|\beta|$ or β^2 are usually more appropriate than derivatives with respect to β.

Equations (2.4-1) and (2.4-2) also assume that during a maneuver the Mach number, Reynolds number, dynamic pressure, velocity, engine parameters, and other flight condition parameters change little enough that their effects on the nondimensional coefficients are negligible. Note that we are concerned here only with the effects on nondimensional coefficients. The dimensional coefficients (Etkin, 1959) are directly proportional to dynamic pressure, for instance, but the nondimensional coefficients are usually insensitive to dynamic pressure over large ranges (until high dynamic pressure causes structural deflections).

Equations (2.4-1) and (2.4-2) omit $\dot{\alpha}$ and $\dot{\beta}$ terms such as $C_{m_{\dot{\alpha}}}$ for the reasons outlined in Section 1.4. This does not constitute neglect of the $\dot{\alpha}$ and $\dot{\beta}$ terms; their effects are subsumed in the other terms. Explicit $\dot{\alpha}$ and $\dot{\beta}$ derivatives can be added to these equations when appropriate. (Maine and Iliff (1979) discusses this issue in more detail.)

The final assumption in Equations (2.4-1) and (2.4-2) is that of linearity. Over large ranges of the parameters, particularly angle of attack, the aerodynamics will not be linear. The equations therefore restrict the allowable maneuvers to those small enough for a locally linear model to be a reasonable approximation. The range of applicability can be expanded by adding nonlinear terms to Equations (2.4-1) and (2.4-2); subsequent chapters discuss the advantages and problems of modeling nonlinearities in this way. In some situations, notably in separated flow, local linearity is not a good assumption even for very small maneuvers; the coefficients in separated flow can be discontinuous functions and can exhibit hysteresis. Thus, conditions of separated flow can require significantly different forms than these equations.

Equations (2.4-1) and (2.4-2) are starting points from which you can expand as needed, adding terms for any parameter that has a measurable effect on the aircraft dynamics during the maneuver.

Given the body-axis coefficients C_N, C_A, and C_Y, the coefficients C_L, $C_{D_{wind}}$, and $C_{Y_{wind}}$ appearing in Equation (2.2-22) are given by

$$C_L = C_N \cos \alpha - C_A \sin \alpha \tag{2.4-3a}$$

$$C_{D_{wind}} = C_A \cos \alpha \cos \beta + C_N \sin \alpha \cos \beta - C_Y \sin \beta \tag{2.4-3b}$$

$$C_{Y_{wind}} = C_Y \cos \beta + C_A \cos \alpha \sin \beta + C_N \sin \alpha \sin \beta \tag{2.4-3c}$$

We could conversely expand C_L, $C_{D_{wind}}$, and $C_{Y_{wind}}$ as functions of the parameters and write C_N, C_A, and C_Y as

$$C_N = C_L \cos \alpha + C_{D_{wind}} \sin \alpha \cos \beta + C_{Y_{wind}} \sin \alpha \sin \beta \tag{2.4-4a}$$

$$C_A = C_{D_{wind}} \cos \alpha \cos \beta + C_{Y_{wind}} \cos \alpha \sin \beta - C_L \sin \alpha \tag{2.4-4b}$$

$$C_Y = C_{Y_{wind}} \cos \beta - C_{D_{wind}} \sin \beta \tag{2.4-4c}$$

Note that if C_N, C_A, and C_Y are linear in α, then by Equation (2.4-3) C_L, $C_{D_{wind}}$, and $C_{Y_{wind}}$ are nonlinear; conversely, if C_L, $C_{D_{wind}}$, and $C_{Y_{wind}}$ are linear in α, then C_N, C_A, and C_Y are nonlinear. Therefore, the body-axis and wind-axis expansions are not completely equivalent, although they are close. One might be valid over a larger α range than the other.

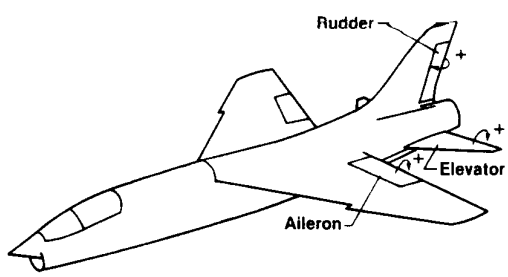

Figure 2.1-1. Control surfaces.

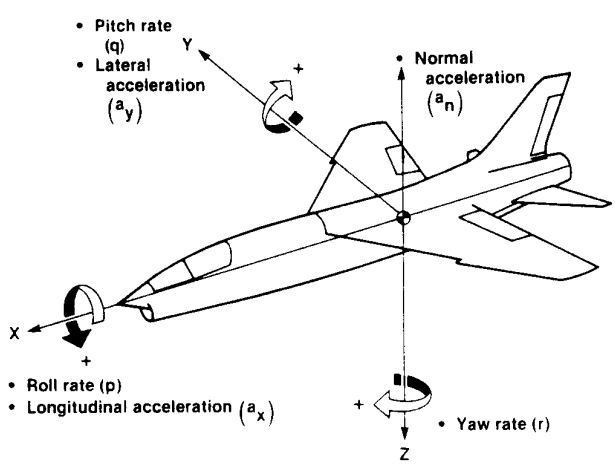

Figure 2.1-2. Body-axis system.

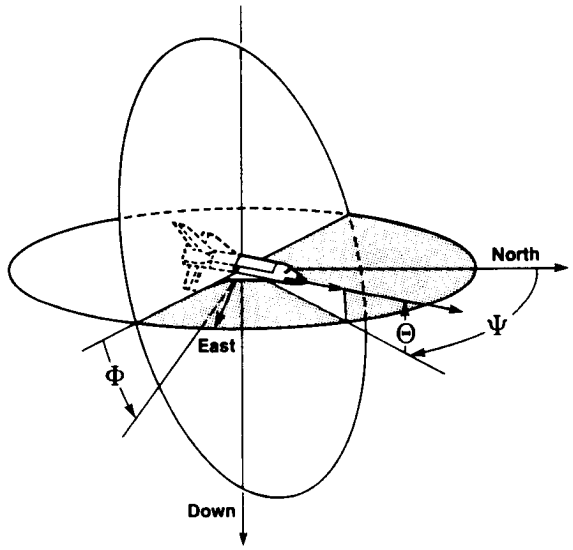

Figure 2.1-3. Euler angles.

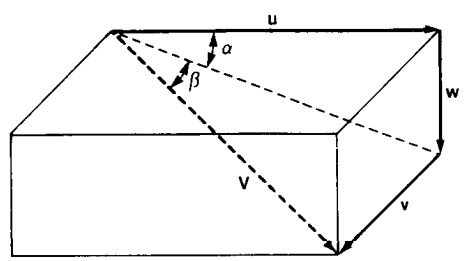

Figure 2.1-4. Flow angles.

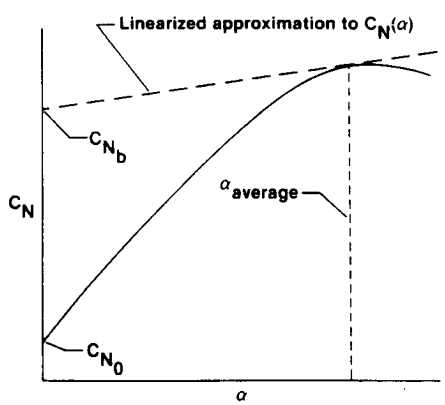

Figure 2.4-1. Relationship of C_{N_b} and C_{N_0}.

3.0 SIMPLIFYING THE EQUATIONS

Chapter 2 introduces the aircraft stability and control problem and develops the basic system equations describing aircraft motion. In Chapter 3 we discuss application of parameter estimation methods to these equations. We simplify the general nonlinear equations of Chapter 2, obtaining forms amenable to practical use.

The six-degree-of-freedom equations derived in Chapter 2 are much too complicated for convenient application of parameter estimation. Although it is certainly possible to work with the full set of equations, we seldom wish to do so in practice. Recall our emphasis on the importance of simple model forms. One of the most important engineering judgments in parameter estimation lies in striking a balance between models that are too complicated to analyze and models that are too simple to represent the system. The full six-degree-of-freedom equations lead toward models too complex for most of our purposes. Therefore, we investigate several simplifications of these equations.

Several of the simplification techniques that we use are common to many fields and require little elaboration here. We separate the equations into (nearly) independent subsets, linearize, neglect small terms, and restrict the maneuvers so as to keep more terms small enough to neglect.

Many of the implementation issues are strongly interconnected and must be considered together from a systems viewpoint. You will not get very far considering each implementation issue divorced from the others. The process of implementation is iterative, rather than linear. You make tentative decisions on early issues and see how the ideas fit together. You will likely later find that some of the tentative decisions merit revision. The structure of this chapter somewhat reflects the interdependent nature of the implementation decisions. For instance, we need to define a coordinate system before we can write much of anything concrete, yet the choice of coordinate system is strongly affected by the ease of linearizing the equations, which is one of the last issues addressed. Because of this interrelationship there is a significant amount of forward referencing in this chapter for details.

Throughout this chapter we assume the use of an output-error estimator. The equations of Chapter 2 form the basis for all analysis methods. Several of the recommendations of this chapter, however, are strongly dependent on the choice of analysis method. For example, the issue of whether to include various terms in the output equation or in the data preprocessing is irrelevant to output-equation error analysis (often just called equation error; Maine and Iliff, 1984); the two approaches are equivalent in output-equation error. The equations presented in this chapter are not necessarily good choices for analysis using an output-equation error estimator.

3.1 NONLINEAR VERSUS LINEARIZED IMPLEMENTATIONS

There are numerous trigonometric and multiplicative nonlinearities in the aircraft stability and control equations derived in Chapter 2. If we use the (α,β,V) equations instead of the (u,v,w) equations, there are also inversions of V. The dynamic pressure $\bar{q}$, which multiplies all the aerodynamic coefficients, hides another nonlinearity because it is defined as

$$\bar{q} = \frac{1}{2}\rho V^2 \tag{3.1-1}$$

We have not yet even considered the nonlinearities in the aerodynamic coefficients.

There are two significantly different approaches to applying the output-error method to nonlinear systems. The first approach is to use an output-error program designed for analysis of nonlinear systems. There are no assumptions of linearity inherent in the output-error method. Several programs capable of handling general nonlinear systems exist (Jategaonkar and Plaetschke, 1983).

The second approach is to linearize the equations for use in an output-error program designed for linear systems. Several widely used output-error programs are restricted to linear (or linearized) systems; for applications where linearization is feasable, these programs are more efficient and easier to use.

Both nonlinear and linearized implementations of output error have roles in practical analysis. Programs restricted to linearized models are appropriate for routine batch analysis of large amounts of data in cruise flight regimes. In such circumstances, computational efficiency is important, and nonlinearities play a minimal role. The general nonlinear programs are appropriate for intensive interactive analysis of data where nonlinear effects are important, such as at high angles of attack. In such situations, computational efficiency is less important than the ease of implementing nonlinearities.

This chapter emphasizes the linearized implementation of output error, which is the more difficult of the two approaches. We specifically treat the use of the Iliff-Maine code, MMLE3 (Maine and Iliff, 1980), although the principles discussed are applicable to most output-error programs for linear systems. Several of the issues discussed are important to both nonlinear and linearized implementations, particularly the uncoupling of the state equations, the treatment of biases and initial conditions, and the handling of sensor corrections. Just because it is easy, for instance, to implement the full six-degree-of-freedom equations in a nonlinear program does not mean that it is a good idea. Section 3.7 is largely

irrelevant to analysis with nonlinear programs. The choice of axis systems, discussed several places in the chapter, is significantly influenced by linearization considerations; therefore, the best choice might be different for nonlinear analysis.

3.2 COORDINATE SYSTEMS

The choice of coordinate systems is a major decision affecting the implementation of the aircraft equations of motion. In this section we outline some of the issues to be considered in making this decision.

There is little debate about the best choice of coordinate system for rotational degrees of freedom. The body-axis angular rates and the Euler angles are the most naturally measured quantities. The choice of these quantities as states allows the use of the trivial observation equations (2.3-2c) to (2.3-2h) and (2.3-2ℓ) to (2.3-2n). Furthermore, the state equations (2.2-21d) to (2.2-21f) are reasonably simple in this coordinate system. Although the state equations may at first appear complicated, they simplify considerably under reasonable conditions (see Section 3.4).

The choice of coordinate systems for the translational degrees of freedom is more difficult. The normally used measurements are expressed in two different coordinate systems. Wind-relative velocity is most commonly measured in terms of α, β, and V; this naturally suggests the use of the (α,β,V) form of the state equations (2.2-22), which involve wind-axis force coefficients. The accelerations, on the other hand, are measured in the body-axis system, therefore Equation (2.3-2) involves body-axis force coefficients. This suggests the use of the (u,v,w) form of the state equations (2.2-21), which likewise involve body-axis force coefficients.

Thus, both the (α,β,V) and the (u,v,w) forms of the state equations have objectionable features when used with the usually available measurements. The (α,β,V) form gives simple observation equations for α, β, and V, but introduces a nonlinear relationship between the wind-axis force coefficients in the state equations and the body-axis coefficients in the acceleration observation equations. The (u,v,w) form results in consistent use of all force coefficients in the body-axis system, but gives nonlinear observation equations for α, β, and V.

In principle, either the (α,β,V) or the (u,v,w) form is usable. The nonlinearities somewhat complicate the implementation, but these complications can be handled in several ways. (Sections 3.5.2 and 3.7.2 discuss this issue further.) Of course, with a program designed for nonlinear systems, the nonlinearities present no inherent problems. Neither form stands out as the obvious best choice for all conceivable situations.

In this document we concentrate on the use of the (α,β,V) form of the state equations, although many of the principles apply equally well to the (u,v,w) form. Our output-error program for linear systems more naturally accommodates the linearizations necessary in the (α,β,V) form than those necessary in the (u,v,w) form. We have found the (α,β,V) form adequate for all applications we have addressed.

In many cases, the (u,v,w) form should be equally suitable, but we have not had the need to try it extensively. Ross and Foster (1976) presents one of the several computer programs that have been successfully applied with the (u,v,w) form. For hovering flight, the (u,v,w) form is required because of the singularities in the (α,β,V) form.

3.3 SIMPLIFICATION USING MEASURED DATA

One simplification technique that merits discussion is the use of measured data. In this technique we eliminate the equations required to compute some variable and substitute the measured value of that variable wherever it appears in other equations. This technique requires that measured data be available; it thus does not generally apply to simplifying simulators, design studies, or many other uses of the equations of motion. However, it can be a powerful tool in simplifying parameter estimation problems.

The use of measured data is closely related to the equation-error method of parameter estimation. We use measurements as though they were exact noise-free values. The method is therefore sensitive to noise in the measurements so used. As we eliminate more differential equations in favor of measured values, the algorithm takes on more characteristics of equation error.

Our general approach is a compromise. We use our knowledge of aircraft dynamics to identify the few equations that dominate the characteristics we are studying, and we center our analysis around these equations. Substituting measured data is one of the tools used to eliminate the equations peripheral to the study.

For example, aircraft velocity is only weakly related to the longitudinal short-period mode. The short-period characteristics definitely depend on the velocity, so we must consider velocity in short-period analysis. The coupling, however, is almost entirely static; there is little dynamic coupling between the velocity and the short-period mode. Therefore, we use measured velocity in our short-period analysis, accounting for the velocity effects without integrating the $\dot{V}$ equation.

The validity of this example, as with several of the simplifications we introduce in this chapter, is configuration dependent. In a configuration with very low wing loading, such as a human-powered or solar-powered aircraft, the velocity might be closely coupled with the short-period mode; indeed there might not be conventional short-period and phugoid modes. We will emphasize the simplifications that can be made for conventional flight vehicles. For unconventional configurations, many of the assumptions must be reevaluated on a case-by-case basis. This continual reevaluation is an unavoidable consequence of making assumptions and simplifications. Unfortunately, any set of equations general enough to handle all conceivable situations will be so complicated as to be unusable. Careful reevaluation of assumptions and examination of the data for potential violations of the assumptions are two of the hallmarks of thorough data analysis as opposed to uncritical number crunching.

3.4 UNCOUPLED STATE EQUATIONS

Most aircraft are symmetric about the X-Z plane and fly at small sideslip angles. We can use this symmetry and small-angle approximations to separate the equations of motion into two largely independent sets describing the longitudinal and lateral-directional motions of the aircraft. The unknown parameters also separate into longitudinal and lateral-directional sets. Thus, the stability and control derivative estimation problem reduces to two smaller problems, each with about half as many differential equations and half as many unknowns as the combined problem. The simplifications from this separation are so great that we go to considerable length to find ways to use the separated forms, even when they superficially appear inapplicable. The separation is important to both nonlinear and linearized implementations of output error.

3.4.1 Longitudinal Equations

The longitudinal motions are rotations about the body Y axis and translations along the X and Z axes. The longitudinal aerodynamic coefficients are thus C_N, C_A (or equivalently, C_L, C_D), and C_m. The longitudinal state equations are (2.2-22a), (2.2-22b), (2.2-22d), and (2.2-22h). We can separate these equations from the others by making the following assumptions: Assume that the vehicle is symmetric about the X-Z plane; thus, I_{xy} and I_{yz} are zero. Further assume that during the maneuver analyzed, β, p, r, and Φ are all constant at zero (or small enough to be unimportant). Then

$$\dot{\alpha} = -\frac{\bar{q}s}{mV} C_L + q + \frac{g}{V}\cos(\Theta-\alpha) - \frac{T\sin\alpha}{mV} \tag{3.4-1a}$$

$$\dot{q} I_y = \bar{q}scC_m \tag{3.4-1b}$$

$$\dot{\Theta} = q \tag{3.4-1c}$$

$$\dot{V} = -\frac{\bar{q}s}{m} C_D + g\sin(\alpha-\Theta) + \frac{T}{m}\cos\alpha \tag{3.4-1d}$$

In many cases, the restrictions on the applicability of Equation (3.4-1) are too severe. These equations do not apply, for instance, in a steady turn. The use of measured data for the lateral-directional signals allows us to eliminate the lateral-directional differential equations without assuming that the lateral-directional motions are negligible. A more widely applicable set of longitudinal equations is thus

$$\begin{aligned}\dot{\alpha} = &-\frac{\bar{q}s}{mV\cos\beta} C_L + q - (\tan\beta)(p\cos\alpha + r\sin\alpha) \\ &+ \frac{g}{V\cos\beta}(\cos\Phi\cos\Theta\cos\alpha + \sin\Theta\sin\alpha) - \frac{T\sin\alpha}{mV\cos\beta}\end{aligned} \tag{3.4-2a}$$

$$-\dot{p}I_{xy} + \dot{q}I_y - \dot{r}I_{yz} = \bar{q}scC_m + rp(I_z - I_x) + (r^2 - p^2)I_{xz} + qrI_{xy} - pqI_{yz} + NrI_{xe} \tag{3.4-2b}$$

$$\dot{\Theta} = q\cos\Phi - r\sin\Phi \tag{3.4-2c}$$

$$\begin{aligned}\dot{V} = &-\frac{\bar{q}s}{m} C_{D_{wind}} \\ &+ g(\cos\Phi\cos\Theta\sin\alpha\cos\beta + \sin\Phi\cos\Theta\sin\beta - \sin\Theta\cos\alpha\cos\beta) \\ &+ \frac{T}{m}\cos\alpha\cos\beta\end{aligned} \tag{3.4-2d}$$

where we use measured data for the lateral-directional signals β, p, r, and Φ. Although Equation (3.4-2) eliminates the lateral-directional differential equations, it does not completely eliminate the lateral-directional aerodynamic coefficients and is thus not yet adequately separated from the lateral-directional problem. The coefficient $C_{D_{wind}}$ in Equation (3.4-2d) is not purely longitudinal, because it involves the lateral force coefficient C_Y (see the definition of $C_{D_{wind}}$ in Equation (2.2-20)). To complete the separation of the longitudinal and lateral-directional problems, we assume that β is small enough that C_D can be substituted as an approximation for $C_{D_{wind}}$ in Equation (3.4-2d). With this approximation, Equation (3.4-2) allows us to concentrate on longitudinal analysis, even in the presence of substantial lateral-directional motions. The only assumption required, in addition to the availability of appropriate lateral-directional measured data, is that β be small enough for the C_D approximation to be reasonable.

We can often further reduce the longitudinal equations by eliminating the $\dot{V}$ equation. In most cases, longitudinal stability and control maneuvers predominantly excite the short-period mode and not the phugoid mode; thus, V is essentially constant during the maneuver. If V does change noticeably, we can use the measured velocity where needed to avoid integrating the $\dot{V}$ equation. In addition to reducing the number of differential equations, this eliminates the issue of approximating $C_{D_{wind}}$ by C_D because that approximation is required only in the V equation.

Further simplification is possible by eliminating the $\dot{\Theta}$ equation and substituting measured Θ. The pitch attitude Θ does change during a typical longitudinal maneuver, but the effect of the motion on the short-period dynamics is small. The Θ variable enters only the gravity terms in the $\dot{\alpha}$ and $\dot{V}$ equations. The $\dot{V}$ equation is eliminated by the use of measured V, and for many maneuvers, the gravity term in the $\dot{\alpha}$ equation is well approximated by setting Θ equal to α. Alternatively, we can use measured Θ in the gravity terms.

The simplest usable set of longitudinal state equations comes from Equation (3.4-1), eliminating the $\dot{V}$ and $\dot{\Theta}$ equations, approximating Θ by α in the gravity term, and neglecting thrust. This gives

$$\dot{\alpha} = -\frac{\bar{q}s}{mV} C_L + q + \frac{g}{V} \tag{3.4-3a}$$

$$\dot{q} I_y = \bar{q}scC_m \tag{3.4-3b}$$

using measured $\bar{q}$ and V.

We most commonly base our analysis on the following equations, using measured data for V, $\bar{q}$, β, p, r, and Φ:

$$\dot{\alpha} = -\frac{\bar{q}s}{mV} C_L + q - (\tan\beta)(p\cos\alpha + r\sin\alpha) + \frac{g}{V}(\cos\Phi\cos\Theta\cos\alpha + \sin\Theta\sin\alpha) - \frac{T\sin\alpha}{mV} \tag{3.4-4a}$$

$$\dot{q} = \bar{q}scC_m + rp(I_z - I_x) + (r^2 - p^2)I_{xz} + NrI_{xe} \tag{3.4-4b}$$

$$\dot{\Theta} = q\cos\Phi - r\sin\Phi \tag{3.4-4c}$$

These equations assume a symmetrical vehicle and drop some cos β factors from the $\dot{\alpha}$ equation; they assume small β (up to 10° is generally acceptable) but otherwise allow arbitrary lateral-directional motion. For a nonsymmetric vehicle, extra terms can be added to Equation (3.4-4b).

The equations presented in this section assume the use of the (α,β,V) form of the state equations. If the (u,v,w) form of the state equations is used, similar simplifications apply, with the $\dot{u}$ equation substituted for the $\dot{V}$ equation and the $\dot{w}$ equation substituted for the $\dot{\alpha}$ equation.

3.4.2 Lateral-Directional Equations

The lateral-directional motions are rotations about the body X and Z axes and translations along the body Y axis. The lateral-directional aerodynamic coefficients are thus C_Y, C_ℓ, and C_n. The lateral-directional state equations are (2.2-22c), (2.2-22d), (2.2-22f), (2.2-22g), and (2.2-22i). It is not realistically possible to have purely lateral-directional motion, even with a symmetric aircraft. There will be some q excitation from the rp and $r^2 - p^2$ terms in Equation (2.2-22e) unless $I_{xz} = 0$ and $I_x = I_z$,

which are unlikely conditions. There will be α and Θ excitation in Equations (2.2-22b) and (2.2-22h) regardless of the inertias. For moderate lateral-directional motions, however, the longitudinal excitations will be small.

To get a workable set of lateral-directional equations, we must approximate $C_{Y_{wind}}$ in Equation (2.2-22c) by C_Y. This approximation is similar in purpose and effect to the approximation of $C_{D_{wind}}$ by C_D in the longitudinal equations. The coefficient $C_{Y_{wind}}$, as defined in Equation (2.2-20), involves the longitudinal coefficient C_D. The C_D involvement is harder to avoid here than was the C_Y involvement in the longitudinal equations because β is important to the lateral-directional dynamics, so eliminating the β equation is seldom reasonable. The maneuvers must have small β for the approximation to be reasonable. This restriction is seldom a problem, because it is difficult for most aircraft to achieve β large enough for the approximation to fail, even if you try.

We usually also neglect the sin β term in Equation (2.2-22c) as not being worth the bother; it is not really difficult to include, it just adds extra algebra. The heading angle Ψ adds very little information and does not feed back into any of the other state equations, so we usually ignore Equation (2.2-22i). For a symmetric vehicle, several terms in Equations (2.2-22d) and (2.2-22f) drop out.

These simplifications give the lateral-directional state equations in the form that we most commonly use:

$$\dot{\beta} = \frac{\bar{q}S}{mV} C_Y + p \sin \alpha - r \cos \alpha + \frac{g}{V} \sin \Phi \cos \Theta \tag{3.4-5a}$$

$$\dot{p} I_x - \dot{r} I_{xz} = \bar{q} S b C_\ell + q r (I_y - I_z) + p q I_{xz} \tag{3.4-5b}$$

$$\dot{r} I_z - \dot{p} I_{xz} = \bar{q} S b C_n + p q (I_x - I_y) - q r I_{xz} - N q I_{xe} \tag{3.4-5c}$$

$$\dot{\Phi} = p + q \tan \Theta \sin \Phi + r \tan \Theta \cos \Phi \tag{3.4-5d}$$

These equations use measured data for the longitudinal variables α, q, Θ, V, and $\bar{q}$. For nonsymmetric vehicles, extra terms can be added to Equations (3.4-5b) and (3.4-5c).

Equation (3.4-5) is based on the use of the (α,β,V) coordinates for wind-relative velocity. The (u,v,w) equations simplify in a similar manner, the $\dot{v}$ equation substituting for the $\dot{\beta}$ equation. The issue of approximating $C_{Y_{wind}}$ by C_Y does not arise in the (u,v,w) form.

3.5 OBSERVATION EQUATIONS

Section 2.3 presents the basic observation equations, both longitudinal and lateral directional, for the aircraft stability and control problem. We discuss some issues of implementation throughout this document. The actual sensor measurements are sometimes complicated functions of the vehicle states. The main implementation issues concern whether to precompute signals corresponding to a simple sensor model or to use a complicated sensor model more representative of the actual measurements. Several of these issues are related to specific sensor characteristics and are treated in more detail in Chapter 8.

In general, whether to use raw sensor data or to precompute more convenient forms depends on several factors, including the complexity of the functions, how well the functions are known, the number of sensors involved, and the general signal qualities and noise levels. This decision also depends on the analysis technique (which may be influenced by similar factors). In particular, the question is moot for output-equation error analysis, where the two choices are equivalent. The discussions in this section assume the use of an output-error method.

The advantages of precomputation are in simplifying the observation equations (or in some cases, the state equations). Precomputation is sometimes the only practical approach for sensors that are extremely complicated functions of the vehicle states and for systems that involve multiple raw sensors to compute one signal. An example of such complication is sensing angle of attack using pressure ports (see Section 8.5.3); the raw sensor data are pressure measurements, not angle of attack. In this case it is preferable to precompute a single angle-of-attack "measurement" from the pressure data rather than to use the raw pressure measurements in the analysis.

The disadvantages of precomputation surface when there are sensor problems. Sensor problems are much easier to diagnose using the raw sensor data rather than some quantity computed as a complicated function of data from several sensors. Complete failure of a sensor is often obvious in either case. The difficulties arise when there are more subtle sensor problems. Precomputation that combines data from several sensors also has the problem that the computed signal tends to combine the worst characteristics of the constituent measurements. Data combined from seven perfect sensors and one noisy sensor often produce a noisy computed signal. The extreme case is that if one of the sensors fails, you cannot do the computation and thus can make no use of the good data from the others. (You can sometimes get around such

problems, but it usually means a lot of work.) If you are using raw sensor data, or at least if the precomputations do not combine data from different sensors, then the failure of one sensor does not alter your ability to use other sensors.

Instrument calibration errors, misalignments, cross-axis sensitivity, and similar factors are usually easiest to manage by precorrecting the data unless the exact corrections are unknown. The corrections are usually small enough that they do not change the basic character of the signals. Sensor problems will seldom be disguised by the small correction terms. The choice of where to resolve these errors is mostly a matter of convenience. We discuss the subjects of sensor position and coordinate transformation in Sections 3.5.1 and 3.5.2. (Instrument biases present special problems discussed in Section 3.6.)

3.5.1 Sensor Position

Several of the observation signals used in aircraft stability and control analysis are functions of the position of the sensors in the aircraft. The simplest equations (such as Equation (2.3-1)) assume that the sensors are exactly at the aircraft center of gravity; in a real airplane we are seldom so fortunate. Equation (2.3-2) shows the major effects of sensor position on the data.

There are two approaches to accommodating sensors offset from the center of gravity. The first approach is to correct the sensor measurements for the offset, giving a signal that simulates a measurement at the center of gravity. This approach works well with angle-of-attack measurements because the correction terms are proportional to pitch rate and roll rate, for which we generally have accurate, low-noise measurements. The largest accelerometer corrections, however, are proportional to the angular accelerations. The measurements of angular accelerations are often noisy, if they are acceptable at all. We can differentiate the angular rates, but such differentiation accentuates noise and causes several other problems. Therefore, careless linear accelerometer corrections can introduce unacceptable noise.

The second approach is to include the effects of sensor position in the observation equations. This approach avoids the use of possibly noisy measured angular accelerations to correct the measured linear accelerations. Instead, we use the computed angular accelerations to correct the computed linear accelerations to the actual sensor positions. The computed angular accelerations are relatively noise-free because they result from integrating the system equations; they do not involve numerical differentiation. There are two other relatively minor benefits of this approach. First, it is more representative of the real instrumentation system. This benefit shows up mainly when you have instrumentation problems. If a rate gyro fails, you lose only one signal; it does not also corrupt the linear accelerations. For similar reasons, instrumentation problems are easier to diagnose. The second benefit of modeling the sensor position in the observation equations is that, in principle, we can treat unknown instrument positions (the instrument position simply being another parameter to estimate). In most situations it is far better to measure the sensor position and treat it as known (an application of the general principle of minimizing the number of parameters to estimate). There are occasions, however, where the ability to estimate sensor positions proves useful.

We generally find it more convenient to model the instrument position in the observation equation than to correct the measurements to the center of gravity, but we have used both approaches successfully.

3.5.2 Coordinate Transformation

Another type of precomputation is the transformation of sensor data to a coordinate system other than that of the raw measurements. We prefer to avoid such transformations because of the previously described difficulties in diagnosis and management of sensor problems. Sensor coordinate transformations can, however, significantly simplify some of the observation equations. This advantage can outweigh the disadvantages in some situations. This section discusses two possible sensor coordinate transformations that give alternatives to the equations we usually use.

The first possibility is to transform the wind-relative velocity from the usually measured (α,β,V) form into the (u,v,w) form. The basic equation for the transformation is Equation (2.1-7). Transforming the wind-relative velocity into (u,v,w) coordinates allows the use of the (u,v,w) form of the state equations without resorting to nonlinear observation equations for α, β, and V. The nonlinearities are transferred from the observation equations, where they cause analysis difficulties, to the data preprocessing, where managing nonlinearities is straightforward in principle (assuming perfect measurements). The disadvantages lie in the mixing of the data from the α, β, and V sensors. The flow-angle sensors typically have different problems than the pressure sensors used for velocity measurements (see Chapter 8). The computed u, v, and w signals combine the problems of both sensors.

The apparent simplicity of Equation (2.1-7) is deceptive. All quantities in this equation relate to the free-stream relative wind at the center of gravity. Application of this transformation is complicated in practice by the fact that the α, β, and V sensors measure local flow conditions relative to the sensor positions. Further, each of the sensors is at a different position (although the position differences are sometimes small enough to neglect). Before applying Equation (2.1-7), you should correct the sensors for local flow effects (upwash and sidewash) and transform the data to the center of gravity. A simple bias error in the angle-of-attack measurement takes on a more complicated character after passing through the nonlinear transformation into (u,v,w) coordinates.

For small-perturbation maneuvers at low angles of attack, the issue of whether to use the measurements in (α,β,V) or (u,v,w) form is largely moot. Implementation is equally easy with or without the transformation. For these conditions, the transformation is nearly linear, as are the observation equations for α, β, and V. This obviates many of the problems of using the transformation and also many of the prob-

lems the transformation is intended to solve. The issues become pressing only at high angles of attack and sideslip.

The second transformation we will consider has a similar objective: to eliminate the nonlinearities in the observation equations. This transformation, however, assumes the use of the (α,β,V) form of the state equations. The principal observation nonlinearities of this form are in the accelerations. The transformation that removes these nonlinearities is the rotation of the body-axis acceleration measurements to the wind-axis coordinate system.

Although this transformation sounds attractive, its practical problems are severe, and we recommend against its use. Before applying the rotation, you must correct the accelerometers to a common reference point, preferably the center of gravity. (Any other body-fixed point has wind-axis coordinates that are a function of the angles of attack and sideslip; this dependence can be handled, but is a nuisance.) The problems of correcting accelerometer measurements for sensor position are discussed in Section 3.5.1.

The most severe problems, however, lie in the fact that the axis transformation equations use angles of attack and sideslip. Accelerometers are reliable, good-quality, high-frequency sensors with excellent linearity and accurate calibrations (see Section 8.4). The transformation corrupts the accelerometer measurements with the biases, scale factor errors, and other problems of the flow-angle measurements (see Section 8.5).

3.6 BIASES AND INITIAL CONDITIONS

Biases and initial conditions are usually nuisance parameters. A nuisance parameter is an unknown parameter that is not really of interest. In many cases you must estimate the value of the nuisance parameter, whether you care about its value or not, in order to get usable estimates of the other parameters. You must carefully consider the treatment of biases and initial conditions, nuisance parameters or not. Poor treatment of biases and initial conditions is a common source of problems, common enough to merit this separate section on the issue. The initial-condition problem is discussed in general terms in section 8.2 of Maine and Iliff (1984). As we see in the following discussion, initial conditions and biases are closely related.

Although you must estimate some bias and initial-condition terms in order to get reasonable stability and control derivative estimates, you cannot estimate all such terms, because of an identifiability problem. To see the problem, consider the time-invariant system

$$\dot{x} = Ax + s_b \tag{3.6-1a}$$

$$z = Cx + z_b \tag{3.6-1b}$$

$$x(0) = x_0 \tag{3.6-1c}$$

with unknown state equation bias s_b, measurement bias z_b, and initial condition x_0. This system is unidentifiable because we can define an equivalent system with state x' by letting

$$x' = x - x_b \tag{3.6-2}$$

for any constant x_b. Then, substituting Equation (3.6-2) into Equation (3.6-1),

$$\dot{x}' = Ax' + s_b' \tag{3.6-3a}$$

$$z = Cx' + z_b' \tag{3.6-3b}$$

$$x_0' = x_0 - x_b \tag{3.6-3c}$$

where

$$s_b' = s_b + Ax_b \tag{3.6-4a}$$

$$z_b' = z_b + Cx_b \tag{3.6-4b}$$

$$x_0' = x_0 - x_b \tag{3.6-4c}$$

Therefore, you can expect to estimate only some subset of the initial conditions and biases. For a system with ℓ states and m observations, there are $2\ell + m$ parameters in s_b, z_b, and x_0. This argument shows that you can independently estimate no more than $\ell + m$ such parameters. In Sections 3.6.1 and 3.6.2 we will present two classes of reasonable subsets for aircraft stability and control analysis; we have used variants of both on various occasions.

3.6.1 Perturbation Equations

The first subset results from the use of perturbation equations. Define $x_b = x_0$, so that x_0' is zero; then estimate s_b' and z_b'. This is the simplest choice in many ways and is applicable to almost any time-invariant linear system. The main disadvantage is that you must work with x' instead of x. Therefore, if the state x consists of physically meaningful parameters, you lose some of the meaning. The method also does not extend well to nonlinear or time-varying systems (sometimes it can be done, but it creates complications).

Specifying x_0 is sufficient to eliminate bias indentifiability problems in most cases. In some circumstances, you can further reduce the number of nuisance bias parameters without causing problems. The main reason for doing this is to save computer time, not to improve identifiability; with x_0 fixed, the remaining bias parameters are usually among the best conditioned parameters in the problem. We discuss three ways of reducing the number of bias parameters in the perturbation equations.

First, if x_0' is zero and the measurement noise at the first time point is small, you can fix the value of z_b' at the first measurement; you can apply this idea to a subset of the elements z_b' if some of the measurements are noisy. You must watch for the occasional wild point if you use this idea; then either remove the wild point, let z_b' be unknown, or (often the easiest solution if it is really an isolated wild point) start your maneuver one point later. You can use a host of other similar methods to arrive at a fixed value for z_b' (for instance, average the first five measurements); however, remember that you are just trying to save time and make the problem simpler. Our attitude is that more complicated methods are not worth the trouble; if the first-point method is inadequate, let z_b' be unknown.

Second, you can subtract the initial measurement from all the measured data and fix z_b' at zero; this is equivalent to fixing z_b' at the initial measured value, except that you must keep track of the difference between your perturbation measurements and the real measurements. The choice of this method or the first is one of preference and convenience; we find it more convenient to use the real measurements.

Third, you can reduce the number of bias nuisance parameters to estimate by fixing s_b' at zero provided that x_0' is zero, the manuever starts at nearly steady state, and the system is stable (neutral stability is not sufficient). We generally advise against this practice because you do not want to restrict yourself to carefully stabilized maneuvers with steady initial conditions. One big advantage of statistical parameter estimation over some of the earlier hand techniques for estimating stability and control derivatives is that the newer methods have less stringent maneuver requirements; you can get more maneuvers in less time and with less pilot workload. Note also that if you use Θ as a state, the system has a neutrally stable eigenvalue (or very close to neutral depending on the exact implementation), and fixing s_b' is asking for trouble (which is usually much more effective than asking for help).

3.6.2 Physical States

For aircraft stability and control analysis, the states are usually physically meaningful measurable quantities. We would like the model states to be close to the true physical states. This makes it easier to handle state nonlinearities because it eliminates the necessity of transforming back and forth between physical and model states (and inevitably using the wrong one somewhere). We can accept small differences (for instance, on the order of the measurement biases in a reasonably calibrated instrumentation system). This suggests the following scheme as an alternative to the perturbation methods discussed in Section 3.6.1.

Choose the subset of the observation vector that represents direct state measurements. We assume that there is such a subset (it could be the entire observation vector). The subset is obvious in Equation (2.3-2): it is the longitudinal measurements α_m, q_m, and Θ_m and the lateral-directional measurements β_m, p_m, r_m, and Φ_m. We could abstractly discuss known invertible partitions of the C matrix, or even introduce pseudoinverses, but if the relevant subset is not obvious, there is probably little benefit to the approach. Define x_b such that the corresponding subset of z_b' is zero (x_b is unknown); we can do this as long as the corresponding partition of C is invertible (which trivially is true for Equation (2.3-2)). This sounds more complicated than it is; all it means is that we ignore the longitudinal quantities α_b, q_b, and Θ_b and the lateral-directional quantities β_b, p_b, r_b, and Φ_b.

Specifying this subset of the z_b' vector is sufficient to resolve most bias identifiability problems. It is completely equivalent to specifying x_0 as zero in the perturbation equation approach; the results from either of these two approaches can be transformed into the other form.

The best way to further reduce the number of bias unknowns is to fix x_0' based on the first measurement. The same cautions apply here as when using the first measurement to fix values in the perturbation equations; the measurement noise on the first point must be small, and you must watch out for wild points.

With this approach, it is not practical to fix s_b at known values, as we could with perturbation equations. This is not a significant limitation, because we recommend against fixing s_b with either approach.

3.6.3 Special Considerations

The preceding discussions of biases and initial conditions are fairly general. You really need to reexamine bias and initial-condition treatment for each new model form, considering the special situation presented by each model. We cannot present every case here; we do discuss two simple situations representative of problems you might encounter.

For the first situation, consider a simple set of longitudinal state equations with states α, q, and Θ; control inputs, lateral-directional coupling, and thrust are irrelevant to this argument, so we assume all such terms are zero. In matrix form, we have

$$\begin{bmatrix} \dot{\alpha} \\ \dot{q} \\ \dot{\Theta} \end{bmatrix} = \begin{bmatrix} Z_\alpha & 1 & 0 \\ M_\alpha & M_q & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \alpha \\ q \\ \Theta \end{bmatrix} + \begin{bmatrix} \frac{g}{V} + Z_b \\ M_b \\ 0 \end{bmatrix} \tag{3.6-5}$$

In the s_b vector, Z_b and M_b are unknown, but there is a known zero in the third element. It is tempting to do the estimation in exactly this form, with the third element of s_b fixed at zero. Although you can sometimes get by with this, we strongly recommend against trying. The Θ state equation is an open-loop integration of q. This is neutrally stable and thus quite sensitive to small errors. Unfortunately, several small errors are inevitable. One such small error is the difference between s_b and s_b' in the physical state approach. This difference is usually small (the measurement noise on the first point) but is not zero. Because the difference is small, it is easy to forget that our model is in terms of s_b' instead of s_b. Although s_b has a third element fixed at zero, s_b' does not; in fact the third element of s_b' should equal the negative of the bias in the pitch rate measurement. Even a very small bias, negligible in most other places, generates significant errors in an open-loop integration. Even if all the measurements are perfect (an unlikely event), the imperfection of numerically integrating the equation can introduce noticeable drift in Θ unless you go to a lot of trouble to avoid it.

In most situations, the Θ state equation adds little information, and you can get along quite well using measured Θ instead. If you do use a Θ state, be sure to allow an unknown bias in the state equation; otherwise the extra state is more likely to degrade the results than to improve them. The unknown bias will not eliminate the problems of neutral stability, particularly if the maneuver is long, but will help reduce the biggest errors.

This discussion has assumed the use of an output-error method. If you use a filter-error method, the measurement feedback makes the filter stable (not just neutrally stable). The filter-error method is therefore considerably less sensitive to small errors, and the preceding discussion does not completely apply. With an equation-error method, you never use a Θ state equation, so the discussion is moot.

Similar principles apply to the Φ state equation in lateral-directional models, except that Φ is important to the lateral dynamics and you are less likely to be satisfied with using measured Φ.

The second situation we consider arises when the K_α of Equation (2.3-2a) is unknown. If K_α is unknown and you are using physical states, the initial condition of α must also be unknown; you cannot just set the initial condition equal to the measured α. (You could reasonably approximate the initial condition by the measured α divided by the unknown K_α, but that is a nuisance to implement. It is easier to let the initial condition be an independent unknown.) This point may seem obvious and unworthy of mention, but we have seen the error made several times (we too have made it once or twice). It is easy to fall into the trap of only adding one more unknown (K_α) without rethinking the initial-condition treatment.

This error is insidious in that the results can appear quite reasonable. If you say that the initial condition is known to equal the initial measured value, this strongly implies that K_α is quite close to unity. The computer program, obedient to your specification of the problem, will (assuming the program works correctly) give an estimate close to unity and a high degree of confidence in the estimate. This conclusion usually is at least plausible; thus, you might be convinced that you had learned something. In fact, one of the symptoms of this problem is estimates more consistent than you should reasonably have expected.

Consider, for example, Figure 3.6-1(a). This figure is a sketch of a typical measured angle-of-attack signal and a corresponding computed signal for a longitudinal maneuver. This computed signal assumes that $K_\alpha = 1$ and the angle-of-attack initial condition is 10°. The shape of the computed signal is similar to that of the measured one, but the amplitude of the computed signal is significantly smaller. Such an angle-of-attack fit, when combined with good fits on the other signals, suggests that the value of K_α used in the computed data might be too small.

If we recompute the computed time history of Figure 3.6-1(a) using $K_\alpha = 2$, while leaving the angle-of-attack initial condition at 10°, the revised response is as shown in Figure 3.6-1(b). This fit is far worse than that with $K_\alpha = 1$ as shown in Figure 3.6-1(a). Thus, if the initial condition is fixed at 10°, the estimator will choose a value of K_α close to unity. If the initial condition is free to vary, the best attainable fit will be with $K_\alpha = 2$ and an initial condition of 5°, shown in Figure 3.6-1(c).

The two situations discussed in this section represent two large classes of similar errors. In both of these situations, the error is omitting a bias or initial-condition term that should be estimated. In the case of the Θ state equation, it is usually quite obvious that something is wrong if you omit a necessary bias. The K_α problem is more insidious in that the results appear reasonable and can lead to false conclusions. You cannot avoid bias problems by estimating all biases and initial conditions as unknowns, because you will then encounter identifiability problems.

Although there are many potential problems due to mishandling biases and initial conditions, most of these problems are easy to avoid; many are as trivial as the K_α example. The difficulties arise mostly through neglect; biases are treated lightly because they are nuisance parameters. For each potential bias or initial condition, you should ask what the consequences of neglecting it would be and what identifiability problems you might introduce by estimating it. If you treat biases with the same care as you should accord all other parameters, you will seldom go wrong because biases, basically, are easy.

3.6.4 Aerodynamic Biases

We use the term "aerodynamic bias" to refer to coefficients like the Z_b or M_b of Equation (3.6-5) or their nondimensional equivalents. These are nuisance parameters in some stability and control analyses; as such, we are not particularly concerned about "true" values. An adequate estimate of a nuisance parameter is one that does not cause problems with the important parameters. For some analyses, particularly in the performance area, the aerodynamic bias coefficients might be of specific interest, rather than being nuisance parameters.

If you desire reasonable estimates of aerodynamic bias parameters, you must be particularly conscious of the differences between s_b and s_b' as defined by Equation (3.6-4). The aerodynamic bias parameters are in s_b and z_b, but we obtain estimates of s_b' and z_b'. Because of the identifiability problem mentioned in the introduction to Section 3.6, we cannot accurately estimate s_b and z_b without making some assumptions. The most common such assumption is that the measurement biases are negligible. Any neglected measurement biases will directly cause biases (errors) in the estimates of the aerodynamic biases. Measurement biases have only second-order effects on the estimates of coefficients in the A matrix and other places (no effect if the system is strictly linear); therefore we are often somewhat cavalier about accurately correcting measurement biases. If, however, you want to accurately estimate aerodynamic biases, you first need to be very careful about removing instrument biases.

3.7 LINEARIZED EQUATIONS

The state and observation equations described in Sections 3.4 to 3.6, combined with the aerodynamic models of Section 2.4, define system models. Even though we have assumed linear aerodynamics, these models are nonlinear because of the trigonometric terms in the state equations and the quadratic terms in the observation equations. Although we could use the nonlinear models as given, linearizing greatly reduces the computational effort and complexity. Most of the kinematic nonlinearities play only a minor role in the dynamics.

This entire section (3.7) is irrelevant to nonlinear implementations of output error and to output-equation error approaches.

3.7.1 Linearization Using Measured Data

We linearize many of the nonlinear terms in the system by using the measured data, similar to the way we eliminate differential equations using measured data. There are four approaches to such linearization, depending on whether we use measured values at each time point or average values and whether we use a zero-order (constant) or first-order (linear) expansion about the measured values.

The first and simplest approach is to use a zero-order expansion about the average measured value; that is, substitute the average measured value as a constant in all nonlinear terms. Before substituting the measured angle of attack, it must be corrected to the center of gravity. The model is time invariant with this approach, allowing very efficient computation. Equation (3.4-2) is an example of this approach with the additional restriction that Equation (3.4-2) assumes the average values are zero. For this approach to work well, the nonlinear terms must be nearly constant during the maneuver; the approach is therefore practically restricted to small angles and steady flight (or nearly so).

A second and slightly more general approach is to use a first-order expansion about the average measured values. For example, with this approach the trigonometric factor in the gravity term of Equation (3.4-3a) is linearly approximated by

$$\cos\Phi\cos\Theta\cos\alpha + \sin\Theta\sin\alpha \cong \cos\bar{\Phi}_m\cos\bar{\Theta}_m\cos\bar{\alpha}_m + \sin\bar{\Theta}_m\sin\bar{\alpha}_m$$
$$+ (\sin\bar{\Theta}_m\cos\bar{\alpha}_m - \cos\bar{\Phi}_m\cos\bar{\Theta}_m\sin\bar{\alpha}_m)(\alpha - \bar{\alpha}_m)$$
$$+ (\cos\bar{\Theta}_m\sin\bar{\alpha}_m - \cos\bar{\Phi}_m\sin\bar{\Theta}_m\cos\bar{\alpha}_m)(\Theta - \bar{\Theta}_m) \qquad (3.7\text{-}1)$$

where a bar over a symbol indicates average measured value. We do not include a $(\Phi - \bar{\Phi}_m)$ term because there is no computed Φ in the longitudinal equations. The α_m used in this equation should ideally be corrected to the center of gravity.

The correction of α_m to the center of gravity is not as crucial as in the first approach, because the measured α is used here only as a point about which to expand the equation; as long as the measured α is close to the computed α, the expansion will be reasonable. This approach is valid for larger maneuvers than the first approach because we approximate the nonlinear terms by linear functions instead of constants. The cost is its somewhat greater complexity. The model with this approach is still time invariant, and although Equation (3.7-1), for instance, may look complicated, it is easy to implement.

The third and fourth approaches are similar to the first two except that we use point-by-point measured values instead of averages. These two approaches are valid for large variations in the values (within the constraints of linearized aerodynamics). For instance, we have analyzed longitudinal maneuvers during 360° aileron rolls (Maine and Iliff, 1979). These two approaches result in time-varying linear models, which require substantially more computer time than time-invariant models (a factor of about 2 or 3 in our program). For the most part, the time-varying models are no more complicated to program than the time-invariant ones; they just run more slowly.

You can combine these four approaches, using different approaches for different terms. Generally, the simpler approaches are adequate for the less important terms. For instance, there is little reason to bother with first-order expansions of the q^2 terms in Equation (2.3-2). These centrifugal force terms are usually too small to measure; we include them for some special cases and because they have negligible cost and are so easy to implement with zero-order expansions.

3.7.2 Axis Transformations

By far the most complicated nonlinearity to implement is the innocuous-looking axis transformation of Equation (2.4-3a). The reason for the complication is that C_N and C_A must be expanded as in Equation (2.4-1), and several of the unknown stability and control derivatives enter into the nonlinearity. In principle, any of the four linearization methods apply to this nonlinearity; application is simply a matter of doing the algebra. The Iliff-Maine code MMLE3 (Maine and Iliff, 1980; Maine, 1981b), cannot use point-by-point values in this linearization. We could extend the program's capability but have not yet encountered requirements sufficient to justify the complication. To date we have not found this restriction a serious limitation, because the linear aerodynamics of Equation (2.4-1) are valid only over limited ranges of α.

We can make a significant simplification in Equation (2.4-3a) for low angles of attack. The simplification is

$$C_L \cong C_N \qquad (3.7\text{-}2)$$

which will actually work over a larger angle-of-attack range than you would initially expect. We generally have success with the simplification at conditions up to 20° angle of attack, and it was adequate at 45° angle of attack on the space shuttle at hypersonic speeds. The approximation holds up so well because the C_L term in Equation (2.2-22b) (the only place C_L appears) usually plays a secondary role. The dominant term in Equation (2.2-22b) is q, particularly at high speeds, as in the shuttle case. Therefore, we can be fairly sloppy about C_L without introducing significant errors. At low speeds and high angles of attack, the C_L term is more important, and the approximation of Equation (3.7-2) is inappropriate.

Another way of stating this argument is that the identifiability of the C_L terms in Equation (2.2-22b) is poor. We do not normally estimate these derivatives independently, but rather we compute them from the C_N and C_A derivatives in the observation equation. Because of this poor identifiability, avoid the temptation to estimate C_A derivatives if an a_x measurement is not available. Theoretically, the distinction between the C_N derivatives in the a_n observation equation and the C_L derivatives in the α state equation gives information about C_A derivatives. We have tried it and, as expected, obtained ridiculous results; you would need much better α measurements than we ever have to make such an idea work. The useful C_N and C_A information comes from the observation equations.

Other possible ways around this nonlinearity include using the (u,v,w) form of the state equations, which introduces different nonlinearities in the α, β, and V observations, or doing the measurement coordinate transformations discussed in Section 3.5.2.

3.8 COLLECTED EQUATIONS

Portions of the aircraft stability and control equations are spread over several of the preceding sections in various forms. This section collects some useful forms of the equations.

3.8.1 Longitudinal Equations

We present two sets of longitudinal equations. The first set (Equations (3.8-1) to (3.8-5)) is adequate for much analysis at low angles of attack. It uses the approximation that $C_L = C_N$, linearizes the trigonometric gravity terms by substituting measured α and Θ, and omits the Θ state equation. Although we write the a_x observation equation in this section, we often omit a_x when using these simplified equations (we then do not estimate C_A derivatives). We omit thrust and all but the most important lateral-directional coupling terms.

The state equations are

$$\dot{\alpha} = -\frac{\bar{q}s}{mV} + (C_L + \dot{\alpha}_b) + q + \frac{g}{V}(\cos\Phi\cos\Theta_m\cos\alpha_c + \sin\Theta_m\sin\alpha_c) \tag{3.8-1a}$$

$$\dot{q}I_y = \bar{q}scC_m \tag{3.8-1b}$$

The observation equations are

$$\alpha_z = K_\alpha\left(\alpha - \frac{x_\alpha}{V}q\right) \tag{3.8-2a}$$

$$q_z = q \tag{3.8-2b}$$

$$a_n = \frac{\bar{q}s}{mg}C_N + \frac{x_{an}}{g}\dot{q} \tag{3.8.2c}$$

$$a_x = -\frac{\bar{q}s}{mg}C_A + \frac{z_{ax}}{g}\dot{q} \tag{3.8-2d}$$

$$\dot{q}_z = \dot{q} + \dot{q}_b \tag{3.8.2e}$$

The aerodynamic coefficient expansions are

$$C_N = C_{N_\alpha}\alpha + C_{N_\delta}\delta + C_{N_b} \tag{3.8-3a}$$

$$C_A = C_{A_\alpha}\alpha + C_{A_\delta}\delta + C_{A_b} \tag{3.8-3b}$$

$$C_m = C_{m_\alpha}\alpha + C_{m_\delta}\delta + C_{m_q}\frac{qc}{2V} + C_{m_b} \tag{3.8-3c}$$

and we approximate

$$C_L \cong C_N \tag{3.8-4}$$

The quantities $\bar{q}$, V, Θ_m, and Φ are all measured data in these equations. They can be either average or point-by-point measurements as required (point-by-point measurement gives better fidelity but requires more computer time). To reduce the number of subscripts, we omit the subscript m from those quantities where no value is defined other than the measurement. The quantity α_c is measured angle of attack, corrected for upwash and sensor position. The unknown parameter vector is

$$\xi = (C_{N_\alpha}, C_{N_\delta}, C_{A_\alpha}, C_{A_\delta}, C_{m_\alpha}, C_{m_\delta}, C_{m_q}, C_{N_b}, C_{A_b}, C_{m_b}, C_{L_b} + \dot{\alpha}_b, \dot{q}_b) \tag{3.8-5}$$

Other commonly used, unknown parameters include K_α and the α and q initial conditions.

Note that C_{L_b}, C_{N_b}, C_{A_b}, and C_{m_b}, as used in these equations, include measurement bias effects in addition to the aerodynamic biases. See Section 3.6.4 for a discussion of this issue. For instance, there is no separate a_{n_b} unknown, because a_{n_b} is subsumed in C_{N_b}. To keep the notation simple, we

include the aerodynamic C_{N_b} and measurement bias effects in the same symbol because both are nuisance parameters and we cannot independently estimate both anyway. We do explicitly indicate measurement bias effects added to C_L (we use the notation $\dot{\alpha}_b$, indicating a bias in the $\dot{\alpha}$ equation) to make it clear that the measurement bias effects subsumed in C_{L_b} are different from those subsumed in C_{N_b} (for instance, a_{n_b} affects C_{N_b} but not C_{L_b}). The implementation in this case is simpler than the notation; we simply treat C_{N_b} and C_{L_b} as being independent.

Equations (3.8-1) to (3.8-5) are adequate for most longitudinal maneuvers at low angle of attack. The second set of longitudinal equations (Equations (3.8-6) to (3.8-12)) distinguishes C_L from C_N, as needed at high angle of attack. It also includes the Θ state, thrust (aligned with the X axis), and more of the kinematic lateral-directional coupling terms (still assuming symmetry). This set of equations linearizes about measured values in the gravity term and the transformation of C_N and C_A into C_L; this allows you to get by with quite rough measurements of α and Θ, even during large-amplitude maneuvers.

The state equations are based on linearizing Equation (3.4-4)

$$\dot{\alpha} = -\frac{\bar{q}s}{mV}(C_L + \dot{\alpha}_b) + q - (\tan\beta)(p\cos\alpha_c + r\sin\alpha_c)$$

$$-\frac{T\sin\alpha_c}{mV} + \frac{dg_b}{d\alpha}(\alpha - \alpha_c) + \frac{dg_b}{d\Theta}(\Theta - \Theta_m) + g_b \tag{3.8-6a}$$

$$\dot{q}I_y = \bar{q}scC_m + rp(I_z - I_x) + (r^2 - p^2)I_{xz} + NrI_{xe} \tag{3.8-6b}$$

$$\dot{\Theta} = q\cos\Phi - r\sin\Phi + \dot{\Theta}_b \tag{3.8-6c}$$

where the gravity term components are

$$g_b = \frac{g}{V}(\cos\Theta_m\cos\Phi\cos\alpha_c + \sin\Theta_m\sin\alpha_c) \tag{3.8-7a}$$

$$\frac{dg_b}{d\alpha} = \frac{g}{V}(-\cos\Phi\cos\Theta_m\sin\alpha_c + \sin\Theta_m\cos\alpha_c) \tag{3.8-7b}$$

$$\frac{dg_b}{d\Theta} = \frac{g}{V}(-\cos\Phi\sin\Theta_m\cos\alpha_c + \cos\Theta_m\sin\alpha_c) \tag{3.8-7c}$$

Note the important bias term $\dot{\Theta}_b$ in Equation (3.8-6c). Section 3.6.4 discusses the reasons for this term. The observation equations are

$$\alpha_z = K_\alpha(\alpha - \frac{x_\alpha}{V}q + \frac{y_\alpha}{V}p) \tag{3.8-8a}$$

$$q_z = q \tag{3.8-8b}$$

$$\Theta_z = \Theta \tag{3.8-8c}$$

$$a_n = \frac{\bar{q}s}{mg}C_N + \frac{x_{a_n}}{g}\dot{q} + \frac{z_{a_n}}{g}(q_m^2 + p^2) - \frac{y_{a_n}}{g}\dot{p} \tag{3.8-8d}$$

$$a_x = -\frac{\bar{q}s}{mg}C_A + \frac{z_{a_x}}{g}\dot{q} - \frac{x_{a_x}}{g}(q_m^2 + r^2) - \frac{y_{a_x}}{g}\dot{r} + \frac{T}{mg} \tag{3.8-8e}$$

$$\dot{q}_z = \dot{q} + \dot{q}_b \tag{3.8-8f}$$

The aerodynamic coefficient expansions are

$$C_N = C_{N_\alpha}\alpha + C_{N_\delta}\delta + C_{N_b} \tag{3.8-9a}$$

$$C_A = C_{A_\alpha}\alpha + C_{A_\delta}\delta + C_{A_b} \tag{3.8-9b}$$

$$C_m = C_{m_\alpha}\alpha + C_{m_\delta}\delta + C_{m_q}\frac{qc}{2V} + C_{m_b} \tag{3.8-9c}$$

The linearized expression for C_L is

$$C_L = C_{L_\alpha}\alpha + C_{L_\delta}\delta + C_{L_b} \tag{3.8-10}$$

where

$$C_{L_\alpha} = C_{N_\alpha} \cos \alpha_c - C_{A_\alpha} \sin \alpha_c - C_D \tag{3.8-11a}$$

$$C_{L_\delta} = C_{N_\delta} \cos \alpha_c - C_{A_\delta} \sin \alpha_c \tag{3.8-11b}$$

and

$$C_D = C_A \cos \alpha_c + C_N \sin \alpha_c \tag{3.8-12}$$

We could write a complicated expression for C_{L_b} in terms of the C_N and C_A derivatives, but this is unneccessary because we end up estimating $C_{L_b} + \dot{\alpha}_b$ as an independent unknown for the same reasons as in the simpler longitudinal equations. The C_D in Equation (3.8-11a) must be evaluated using only measured values, or the system will not be linear.

As with the simpler longitudinal equations, measured data can be evaluated with either averages or point-by-point measurements. The Iliff-Maine code MMLE3 does not allow the use of point-by-point data in Equation (3.8-11), but this is a limitation of the program rather than of the method.

The vector of unknowns is the same as for the simpler longitudinal equations, with the addition of the bias $\dot{\theta}_b$.

The two sets of equations presented in this section cover a large range of applications. There are numerous other reasonable sets of equations intermediate to these sets. There are essentially an infinite number of terms that you can add to accommodate special situations; nonsymmetric aircraft come to mind as a case requiring extra terms.

3.8.2 Lateral-Directional Equations

We present a single set of lateral-directional equations adequate for most applications. The state equations come from Equation (3.4-5), linearizing the gravity term about the measured Φ, linearizing other terms by substituting measured values, and adding bias unknowns:

$$\dot{\beta} = \frac{\bar{q}s}{mV}(C_Y + \dot{\beta}_b) + p(\alpha_b + \sin \alpha_c) - r \cos \alpha_c$$

$$+ \frac{g}{V} \sin \Phi_m \cos \Theta + \frac{g}{V} \cos \Phi_m \cos \Theta \, (\Phi - \Phi_m) \tag{3.8-13a}$$

$$\dot{p}I_x - \dot{r}I_{xz} = \bar{q}sbC_\ell + qr_m(I_y - I_z) + p_m q I_{xz} \tag{3.8-13b}$$

$$\dot{r}I_z - \dot{p}I_{xz} = \bar{q}sbC_n + p_m q(I_x - I_y) - qr_m I_{xz} - NqI_{xe} \tag{3.8-13c}$$

$$\dot{\Phi} = p + r \tan \Theta \sin \Phi_m + q \tan \Theta \cos \Phi_m \tag{3.8-13d}$$

The observation equations are

$$\beta_z = K_\beta\left(\beta - \frac{z_\beta}{V} p + \frac{x_\beta}{V} r\right) \tag{3.8-14a}$$

$$p_z = p \tag{3.8-14b}$$

$$r_z = r \tag{3.8-14c}$$

$$\Phi_z = \Phi \tag{3.8-14d}$$

$$a_y = \frac{\bar{q}s}{mg} C_Y - \frac{z_{a_y}}{g} \dot{p} + \frac{x_{a_y}}{g} \dot{r} - \frac{y_{a_y}}{g}\left(p_m^2 - r_m^2\right) \tag{3.8-14e}$$

$$\dot{p}_m = \dot{p} + \dot{p}_b \tag{3.8-14f}$$

$$\dot{r}_m = \dot{r} + \dot{r}_b \tag{3.8-14g}$$

The aerodynamic coefficient expansions are

$$C_Y = C_{Y_\beta}\beta + C_{Y_\delta}\delta + C_{Y_b} \tag{3.8-15a}$$

$$C_\ell = C_{\ell_\beta}\beta + C_{\ell_\delta}\delta + C_{\ell_p}\frac{pb}{2V} + C_{\ell_r}\frac{rb}{2V} + C_{\ell_b} \tag{3.8-15b}$$

$$C_n = C_{n_\beta}\beta + C_{n_\delta}\delta + C_{n_p}\frac{pb}{2V} + C_{n_r}\frac{rb}{2V} + C_{n_b} \tag{3.8-15c}$$

We include the term $\dot{\beta}_b$ in Equation (3.8-13a) for the same reasons that $\dot{\alpha}_b$ is included in Equation (3.8-1a); the effect is that the C_{Y_b} in Equation (3.8-13a) is independent of that in Equation (3.8-14e). The parameter α_b in Equation (3.8-13a) also deserves mention: This parameter allows for an unknown bias in the corrected α measurement. We usually do not account for biases in the measured data used to linearize the equations (we account only for biases in the observations we are matching). Biases usually have small effects on the linearizations; accounting for all such biases would entail significant complication and give negligible improvement (indeed, degradation due to identifiability problems would be likely). For instance, we ignore biases in the measured Θ or q used in Equations (3.8-13) and (3.8.14). The (p sin α) term in Equation (3.8-13a) is the one place where we feel that such biases are important. This is because p can be quite large compared with β in aileron maneuvers; therefore, small errors in measured α can have important effects on this term. For instance, the use of average measured α instead of point-by-point values could introduce enough error to be important. We recommend allowing an unknown bias in this term, as shown in Equation (3.8-13a). It is possible to use this term to estimate the angle-of-attack calibration, but the results usually have too much scatter (several degrees) to be very useful for most purposes.

The vector of unknown parameters is

$$\xi = (C_{Y_\beta}, C_{Y_\delta}, C_{\ell_\beta}, C_{\ell_\delta}, C_{\ell_p}, C_{\ell_r}, C_{n_\beta}, C_{n_\delta}, C_{n_p}, C_{n_r}, C_{Y_b}, C_{\ell_b}, C_{n_b}, C_{y_b} + \dot{\beta}_b, \alpha_b, \dot{\Phi}_b, \dot{p}_b, \dot{r}_b) \tag{3.8-16}$$

The equations presented in this section are adequate for most applications. You can add terms as appropriate for special situations.

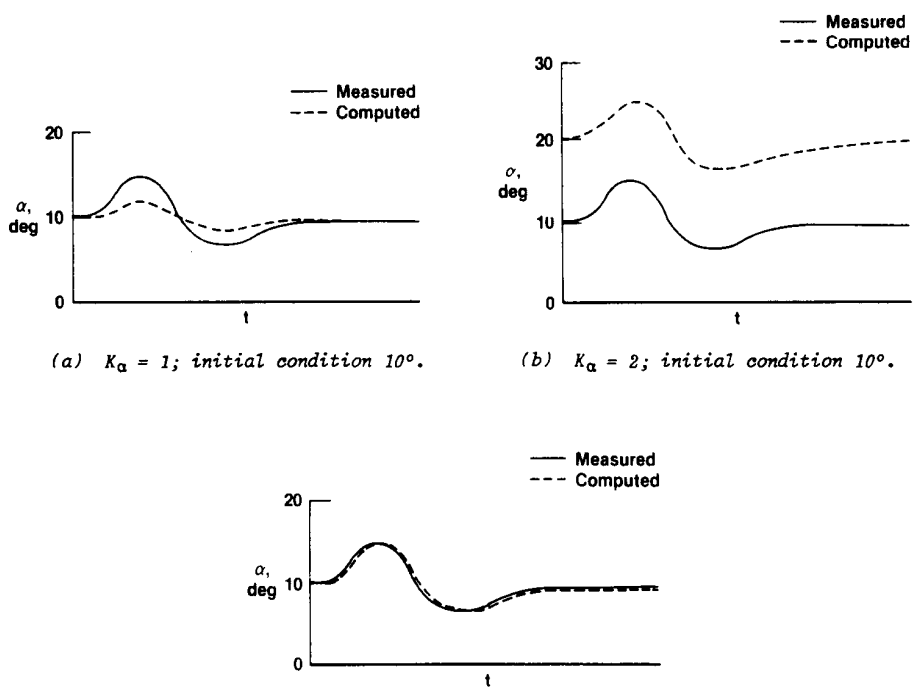

(a) K_α = 1; initial condition 10°.

(b) K_α = 2; initial condition 10°.

(c) K_α = 2; initial condition 5°.

Figure 3.6-1. Angle-of-attack fit.

4.0 DETAILED COMPUTATIONAL ESTIMATION EXAMPLES

In Chapters 2 and 3 we discuss the governing equations that are used for the bulk of the aircraft stability and control estimation that is discussed in the remainder of this document. This chapter deals with the actual estimation of parameters from dynamic data. The purpose of this chapter is to present the parameter estimation steps in enough detail so that someone previously unfamiliar with parameter estimation can fully grasp the method. We begin with a brief description of the computer program used for parameter estimation throughout this document. We then analyze a short computed time history in sufficient detail so that anyone can repeat the steps and obtain the same answers with a simple computer program or even with a pocket calculator. After several variants of this example are discussed, the chapter concludes with a discussion of actual flight data analysis.

4.1 DESCRIPTION OF THE MMLE3 PROGRAM

The Iliff-Maine code (MMLE3 program) is used throughout the remainder of this document to estimate the coefficients of the differential equations of motion. The algorithms used are defined in Maine and Iliff (1984, chapters 7 and 8). The bulk of the analysis in this document is for the estimation of stability and control derivatives from aircraft maneuvers; the MMLE3 program includes the implementation of the aircraft equations described in Chapter 3. The program is fully described in Maine (1981b). Documentation of all the available options and the use of the program are detailed in Maine and Iliff (1980).

Figure 4.1-1 illustrates the maximum likelihood estimation concept for aircraft data as used by MMLE3. The measured response of the aircraft is compared with the estimated response, and the difference between these responses is called the response error. The Gauss-Newton computational algorithm (Maine and Iliff, 1984, section 2.5.2) is used to find the coefficient values that maximize the likelihood functional. Each iteration of this algorithm provides revised estimates of the unknown coefficients based on the response error. These revised estimates of the coefficients are then used to update the mathematical model of the aircraft, providing a revised estimated response and therefore a revised response error. The mathematical model is updated iteratively until a convergence criterion is satisfied. The estimates resulting from this procedure are the maximum likelihood estimates.

The maximum likelihood estimator also provides a measure of reliability of each estimate based on the information obtained from each dynamic maneuver. This measure of reliability, analogous to the standard deviation, is called the Cramér-Rao bound (Maine and Iliff, 1981a; Maine and Iliff, 1984) or the uncertainty level. The Cramér-Rao bound as computed by current programs should generally be used as a measure of relative accuracy rather than absolute accuracy. The bound is obtained from H, the approximation of the information matrix. The information matrix equals the approximation to the second gradient given by Maine and Iliff (1984, equation (8.3-3)). The bound for each unknown is the square root of the corresponding diagonal element of H; that is, for the ith unknown, the Cramér-Rao bound is $\sqrt{H(i,i)}$. The MMLE3 program scales the estimate of the Cramér-Rao bound for the observed noise amplitude by first multiplying the H matrix by a scalar (discussed further in Section 4.3.2).

4.2 EQUATIONS FOR A SIMPLE EXAMPLE

The basic concepts involved in a parameter estimation problem can be illustrated by a simple example representative of a realistic aircraft problem. The example chosen here is representative of an aircraft that exhibits pure rolling motion from an aileron input. This example, although simplified, typifies the motion exhibited by many aircraft in particular flight regimes, such as the F-14 aircraft flying at high dynamic pressure, the F-111 aircraft at moderate speed with the wing in the forward position, and the T-37 aircraft at low speed.

Deriving an equation that describes this motion is straightforward. Figure 4.2-1 depicts an aircraft with the X axis perpendicular to the plane of the figure (positive forward on the aircraft). The rolling moment L', roll rate p, and aileron deflection δ_a are positive as shown. For this example, the only state is p and the only control is δ_a. The result of summing moments is

$$I_x\dot{p} = L'(p,\delta_a) \tag{4.2-1}$$

The first-order Taylor expansion then becomes

$$\dot{p} = L_p p + L_{\delta_a}\delta_a \tag{4.2-2}$$

where

$$L' = I_x L$$

Since the aileron is the only control, it is notationally simpler to use δ for δ_a in discussing this example. Equation (4.2-2) can then be written as

$$\dot{p} = L_p p + L_\delta \delta \tag{4.2-3}$$

An alternative approach that results in the same equation is to combine Equations (3.8-13b) and (3.8-15b) (substituting for C_ℓ) and then eliminate the terms that are zero for our example. This yields

$$\dot{p}I_x = \bar{q}sb\left(C_{\ell p}\frac{pb}{2V} + C_{\ell\delta}\delta\right) \tag{4.2-4}$$

where p is the roll rate and δ the aileron deflection. The equation can be put into the dimensional derivative form of Equation (4.2-3) by rearranging terms.

Equation (4.2-3) is a simple aircraft equation in which the forcing function is provided by the aileron and the damping by the damping-in-roll term L_p. In Sections 4.2 to 4.4 we examine in detail the parameter estimation problem where Equation (4.2-3) describes the system. For this single-degree-of-freedom problem, the maximum likelihood estimator is used to estimate either L_p or L_δ, or both, for a given computed time history.

Now that we have specified the equations describing our simple model, we can examine the characteristics of maximum likelihood estimation in this simple case. Maine and Iliff (1984, chapters 2, 7, and 8) describes in detail maximum likelihood estimation for the general case; for our simple example, we need only some of those results. Where, as in our example, there is no state noise and the equations of motion are linear, the equations are

$$x(t_0) = x_0 \tag{4.2-5}$$

$$\dot{x}(t) = Ax(t) + Bu(t) \tag{4.2-6}$$

$$z(t_i) = Cx(t_i) + Du(t_i) + Gn_i \tag{4.2-7}$$

where x is the state vector, z the observation vector, u the control vector, and the stability and control derivatives defined in Chapter 3 are contained primarily in the A and B matrices.

If there is no state noise and the matrix G is known, then the maximum likelihood estimator minimizes the cost function

$$J(\xi) = \frac{1}{2}\sum_{i=1}^{N}[z(t_i) - \tilde{z}_\xi(t_i)]^*(GG^*)^{-1}[z(t_i) - \tilde{z}_\xi(t_i)] \tag{4.2-8}$$

where GG^* is the measurement noise covariance matrix and $\tilde{z}_\xi(t_i)$ is the computed response estimate of z at t_i for a given value of the unknown parameter vector ξ. The cost function is a function of the difference between the measured and computed time histories.

For the no-state-noise case, the $\tilde{z}_\xi(t_i)$ term of Equation (4.2-8) can be approximated by

$$\tilde{x}_\xi(t_0) = x_0(\xi) \tag{4.2-9}$$

$$\tilde{x}_\xi(t_{i+1}) = \Phi\tilde{x}_\xi(t_i) + \Psi[u(t_i) + u(t_{i+1})]/2 \tag{4.2-10}$$

$$\tilde{z}_\xi(t_i) = C\tilde{x}_\xi(t_i) + Du(t_i) \tag{4.2-11}$$

where

$$\Phi = \exp[A(t_{i+1} - t_i)] \tag{4.2-12}$$

$$\Psi = \int_{t_i}^{t_{i+1}} \exp(A\tau)\, d\tau\, B \tag{4.2-13}$$

To minimize the cost function $J(\xi)$, we can apply the Newton-Raphson algorithm, which chooses successive estimates of the vector of unknown coefficients $\hat{\xi}$. Let L be the iteration number; the L + 1 estimate of $\hat{\xi}$ is then obtained from the L estimate as

$$\hat{\xi}_{L+1} = \hat{\xi}_L - \left[\nabla_\xi^2 J\left(\hat{\xi}_L\right)\right]^{-1}\left[\nabla_\xi^* J\left(\hat{\xi}_L\right)\right] \tag{4.2-14}$$

The first and second gradients are defined as

$$\nabla_\xi J(\xi) = -\sum_{i=1}^{N}\left[z(t_i) - \tilde{z}_\xi(t_i)\right]^*(GG^*)^{-1}\left[\nabla_\xi\tilde{z}_\xi(t_i)\right] \tag{4.2-15}$$

$$\nabla_\xi^2 J(\xi) = \sum_{i=1}^{N}\left[\nabla_\xi \tilde{z}_\xi(t_i)\right]^*(GG^*)^{-1}\left[\nabla_\xi \tilde{z}_\xi(t_i)\right] - \sum_{i=1}^{N}\left[z(t_i) - \tilde{z}_\xi(t_i)\right]^* (GG^*)^{-1}\left[\nabla_\xi^2 \tilde{z}_\xi(t_i)\right] \tag{4.2-16}$$

The Gauss-Newton approximation to the second gradient is

$$\nabla_\xi^2 J(\xi) \cong \sum_{i=1}^{N}\left[\nabla_\xi \tilde{z}_\xi(t_i)\right]^*(GG^*)^{-1}\left[\nabla_\xi \tilde{z}_\xi(t_i)\right] \tag{4.2-17}$$

The Gauss-Newton approximation is computationally much simpler than the Newton-Raphson algorithm because the second gradient of the innovation never needs to be calculated. In addition, the Gauss-Newton approximation can speed the convergence of the algorithm, as is discussed in Section 4.4.1.

Equation (4.2-8) then gives the cost function for maximum likelihood estimation. The weighting GG* is unimportant for this problem, so set GG* = 1. For our example, Equations (4.2-6) and (4.2-7) reduce to $x_i = p_i$ and $z_i = x_i$. Therefore, Equation (4.2-8) becomes

$$J(L_p, L_\delta) = \frac{1}{2}\sum_{i=1}^{N}\left[p_i - \tilde{p}_i(L_p, L_\delta)\right]^2 \tag{4.2-18}$$

where p_i is the value of the measured response p at time t_i and $\tilde{p}_i(L_p, L_\delta)$ is the computed time history of $\tilde{p}$ at time t_i for $L_p = \hat{L}_p$ and $L_\delta = \hat{L}_\delta$. Where computed (not flight) data are used in the remainder of this chapter, the measured time history refers to p_i, and the computed time history refers to $\tilde{p}_i(L_p, L_\delta)$. The computed time history is a function of the current estimates of L_p and L_δ, but the measured time history is not.

The most straightforward method of obtaining $\tilde{p}_i$ is to use Equations (4.2-9) and (4.2-10). In terms of the notation just defined,

$$\tilde{p}_{i+1} = \Phi\tilde{p}_i + \Psi(\delta_i + \delta_{i+1})/2 \tag{4.2-19}$$

where

$$\Phi = \exp(L_p\Delta) \tag{4.2-20}$$

$$\Psi = \int_0^\Delta \exp(L_p\tau)\, d\tau\, L_\delta = \frac{L_\delta[1 - \exp(L_p\Delta)]}{L_p} \tag{4.2-21}$$

and Δ is the length of the sample interval $(t_{i+1} - t_i)$. Simplifying the notation,

$$\delta_{i+1/2} = (\delta_i + \delta_{i+1})/2 \tag{4.2-22}$$

then

$$\tilde{p}_{i+1} = \Phi\tilde{p}_i + \Psi\delta_{i+1/2} \tag{4.2-23}$$

The maximum likelihood estimate is the value that minimizes Equation (4.2-18). The Gauss-Newton method described previously is used for this minimization. Equation (4.2-14) is used to determine successive values of the estimates of the unknowns during the minimization.

The first and second gradients are defined by Equations (4.2-15) and (4.2-16). The only term from Equations (4.2-14), (4.2-15), and (4.2-16) remaining to be defined is $\nabla_\xi z_\xi(t_i)$ or, for our case, $\nabla_\xi \tilde{p}_i$. Equation (24) of Maine and Iliff (1981b) is a corresponding approximation for our problem. In our notation,

$$\nabla_\xi \tilde{p}_{i+1} = \Phi\nabla_\xi \tilde{p}_i + \Psi\left[\left(\nabla_\xi L_p\right)\tilde{p}_{i+1/2} + \left(\nabla_\xi L_\delta\right)\delta_{i+1/2}\right] \tag{4.2-24}$$

For this simple example, the equations can be concisely written in terms of partial derivatives as

$$\left(\frac{\partial \tilde{p}_{i+1}}{\partial L_p}\right) = \Phi\left(\frac{\partial \tilde{p}_i}{\partial L_p}\right) + \Psi p_{i+1/2} \tag{4.2-25}$$

$$\left(\frac{\partial \tilde{p}_{i+1}}{\partial L_\delta}\right) = \Phi\left(\frac{\partial \tilde{p}_i}{\partial L_\delta}\right) + \Psi\delta_{i+1/2} \tag{4.2-26}$$

where

$$\nabla_\xi \tilde{p} = \left[\frac{\partial \tilde{p}}{\partial L_p}, \frac{\partial \tilde{p}}{\partial L_\delta}\right]^* \tag{4.2-27}$$

For this simple example, $\hat{\xi} = [\hat{L}_p, \hat{L}_\delta]^*$, and successive estimates of $\hat{L}_p$ and $\hat{L}_\delta$ are determined by iterating Equation (4.2-14). The first and second gradients of Equation (4.2-14) are defined by Equations (4.2-15) and (4.2-17). The complete set of equations is given in Maine and Iliff (1980).

We can now write the entire procedure for obtaining the maximum likelihood estimates for this simple example. To start the algorithm, initial estimates of L_p and L_δ are needed; these are the value $\hat{\xi}_0$. From Equation (4.2-14), $\hat{\xi}_1$ and subsequently $\hat{\xi}_L$ are defined by using the first and second gradients of $J(L_p,L_\delta)$ from Equation (4.2-18). The gradients for this particular example from Equations (4.2-15) and (4.2-17) are

$$\nabla_\xi J(\hat{\xi}_L) = -\sum_{i=1}^{N}(p_i - \tilde{p}_i)\nabla_\xi \tilde{p}_i \tag{4.2-28}$$

$$\nabla_\xi^2 J(\hat{\xi}_L) \cong \sum_{i=1}^{N}(\nabla_\xi \tilde{p}_i)^*(\nabla_\xi \tilde{p}_i) \tag{4.2-29}$$

Expanded into partial derivative forms, these become

$$\nabla_\xi J(\hat{\xi}_L) = \begin{bmatrix} \sum_{i=1}^{N}(p_i - \tilde{p}_i)\dfrac{\partial \tilde{p}_i}{\partial L_p} \\ \sum_{i=1}^{N}(p_i - \tilde{p}_i)\dfrac{\partial \tilde{p}}{\partial L_\delta} \end{bmatrix} \tag{4.2-30}$$

$$\nabla_\xi^2 J(\hat{\xi}_L) = \begin{bmatrix} \sum_{i=1}^{N}\left(\dfrac{\partial \tilde{p}_i}{\partial L_p}\right)^2 & \sum_{i=1}^{N}\left(\dfrac{\partial \tilde{p}_i}{\partial L_p}\right)\left(\dfrac{\partial \tilde{p}_i}{\partial L_\delta}\right) \\ \sum_{i=1}^{N}\left(\dfrac{\partial \tilde{p}_i}{\partial L_p}\right)\left(\dfrac{\partial \tilde{p}_i}{\partial L_\delta}\right) & \sum_{i=1}^{N}\left(\dfrac{\partial \tilde{p}_i}{\partial L_\delta}\right)^2 \end{bmatrix} \tag{4.2-31}$$

These terms are completely defined for $\tilde{p}_i$ and the partial derivative of $\tilde{p}_i$ by Equations (4.2-23), (4.2-25), and (4.2-26).

4.3 COMPUTATIONAL DETAILS OF MINIMIZATION

In Section 4.2 we specify the equations for a simple example and describe the procedure for obtaining estimates of the unknowns from a dynamic maneuver. In this section we give the computational details for obtaining the estimates. Some basic concepts of parameter estimation are best illustrated by using computed data, where the correct answers are known; therefore, in this section we study two examples involving computed time histories. The first example (Section 4.3.1) is based on data that have no measurement noise; in this case resulting estimates equal the correct values. The second example (Section 4.3.2) is based on data that have significant measurement noise added; consequently, the estimates do not equal the correct values. Throughout this chapter the term "no-noise case" is used for the case with no noise added and "noisy case" for the case where noise has been added.

We desire to keep these examples simple enough so that some or all of the calculations can be completed on a home computer or, with some labor, on a calculator; therefore, we will use only a very small number of data points. For these computed examples, 10 points (time samples) are used. The simulated data (which we call measured data) are based on Equation (4.2-3). We use the same correct values of L_p and L_δ (-0.2500 and 10.0, respectively) for both examples. In addition, the same aileron input δ is used for both examples, the sample interval Δ is 0.2 sec, and the initial conditions are zero. Tables of all significant intermediate values are given for each example. We report these values to 4 significant digits; however, to obtain exactly the same values with a computer or calculator requires the use of 13 significant digits (as we used in computing these tables). If the 4-digit numbers are used in the computation, the answers will differ by a few tenths of a percent, but they will still serve to illustrate the minimization accuracy. In both examples, the initial values of L_p and L_δ (or $\hat{\xi}_0$) are -0.5 and 15.0, respectively.

4.3.1 Example With No Measurement Noise

The measurement time history for the no-noise case is shown in Figure 4.3-1. The aileron input starts at zero, changes to a fixed value, and then returns to zero. The resulting roll-rate time history is also shown. The values of the measured roll rate are given in Table 4.3-1 along with the aileron input.

Table 4.3-2 lists the essential intermediate values for the terms in Equation (4.2-14) for the first four iterations. Except for columns 8 to 12, the columns are self explanatory. Columns 8 and 9 are the intermediate values of the summation for the first gradient terms of Equation (4.2-30). For example, the fifth row is the summation for i = 1 to 5. Likewise, columns 10 to 12 are the intermediate values of the second gradient terms in Equation (4.2-31).

We can obtain $\hat{\xi}_L$ for each value of L (the iteration number) by using Equation (4.2-14) and the values in the last row of the last five columns for each iteration. For the first iteration (from L = 0 to L = 1), we can use Cramér's rule to obtain

$$\left[\nabla^2_{\xi} J(\hat{\xi})\right]^{-1} = \begin{bmatrix} 352.4 & 33.50 \\ 33.50 & 3.701 \end{bmatrix}^{-1} = \frac{\begin{bmatrix} 3.701 & -33.5 \\ -33.5 & 352.4 \end{bmatrix}}{(352.4)(3.701) - (33.50)(33.50)} = \begin{bmatrix} 0.02034 & -0.1841 \\ -0.1841 & 1.936 \end{bmatrix} \tag{4.3-1}$$

$$\nabla J(\hat{\xi}_0) = \begin{bmatrix} -101.0 \\ -12.24 \end{bmatrix} \tag{4.3-2}$$

Equation (4.2-14) becomes

$$\hat{\xi}_1 = \begin{bmatrix} L_p \\ L_\delta \end{bmatrix}_1 = \begin{bmatrix} -0.5 \\ 15.0 \end{bmatrix} - \begin{bmatrix} 0.02034 & -0.1841 \\ -0.1841 & 1.936 \end{bmatrix} \begin{bmatrix} -101.0 \\ -12.24 \end{bmatrix} = \begin{bmatrix} -0.3005 \\ 9.888 \end{bmatrix} \tag{4.3-3}$$

Again, although these equations display 4 significant digits, we have used 13-digit accuracy in the computation.

Table 4.3-3 shows the values for $\hat{L}_p$, $\hat{L}_\delta$, and J for each iteration, along with the values of Φ and Ψ needed for calculating $\tilde{p}_i$. In three iterations the algorithm converges to the correct values (to four significant digits) for both L_p and L_δ. The value of $\hat{L}_\delta$ overshoots slightly on the first iteration and then quickly arrives at the correct answer. The value of $\hat{L}_p$ overshoots slightly on the second iteration.

In Figure 4.3-2 the measured data are compared with the computed data for each of the first three iterations. The match is very good after two iterations and nearly exact after three iterations.

Because we are looking for the maximum likelihood estimate, we expect the value of the gradient to go to zero at the estimate. The gradient is given in Table 4.3-2, in the direction of L_p in column 8 and in the direction of L_δ in column 9. The gradient should go to zero in each of these directions at the maximum likelihood estimate. The gradients for the entire maneuver are shown in Table 4.3-2 as the last row for each iteration in columns 8 and 9; they rapidly approach zero as the algorithm converges. If we continued to iterate, the algorithm would eventually generate very small values.

Although the algorithm converged with four-digit accuracy in L_p and L_δ, the value of the cost function J continued to decrease rapidly between the third and fourth iterations. This is a consequence of using the maximum likelihood estimator on data having no measurement noise. Theoretically, with infinite-accuracy computation, the value of J at the minimum should be zero. However, with finite accuracy, the value of J becomes small but never reaches zero. The value of J is a function of the number of significant digits being used. For this example, using 13-digit accuracy, the cost eventually decreases to approximately 0.3×10^{-28}.

4.3.2 Example With Measurement Noise

The data used in the noisy case are the same as those used in the no-noise case, except that pseudo-Gaussian noise has been added to the roll rate. The time history is shown in Figure 4.3-3. The signal-to-noise ratio is quite low in this example, as is readily apparent by comparing Figures 4.3-1 and 4.3-3. The values of the time history (to 13-digit accuracy) are shown in Table 4.3-4. Table 4.3-5 shows intermediate values for the terms in Equation (4.2-14) for L = 1 to 4 in the same manner described in Section 4.3-1. The succeeding values of $\hat{L}_p$ and $\hat{L}_\delta$ are obtained as described previously. The values of $\hat{L}_p$, $\hat{L}_\delta$, Φ, Ψ, and J for each iteration are shown in Table 4.3-6. The algorithm converges in four iterations. The behavior of the coefficients as they approach convergence is much like that in the no-noise case. The most notable result of this case is the converged values of $\hat{L}_p$ and $\hat{L}_\delta$, which are somewhat different from the correct values. The match between the measured and computed time histories is shown in Figure 4.3-4 for each iteration. No change is apparent for the last two iterations. The match is very good considering the level of measurement noise.

In Figure 4.3-5, the computed time history for the no-noise estimates of L_p and L_δ is compared with that for the noisy-case estimates. In the noisy case the algorithm converged to values somewhat different than the correct values, so the two computed time histories are similar but not identical.

Table 4.3-5 shows the values of the cost function and the values of the gradients with respect to L_p and L_δ. The values of J for the third and fourth iterations are the same to four significant digits. The gradients in each direction are correspondingly small, indicating that the algorithm has converged. The values of the gradients shown in Table 4.3-5 for L = 4 demonstrate an interesting characteristic that arises from having a very small number of samples for a maneuver with a low signal-to-noise ratio: Between i = 9 and 10 the value of the gradient goes from a fairly large number to a small number. This indicates that with a different noise signal or with more samples the gradient would probably be zero at some values of L_p and L_δ significantly different than those shown for this particular case.

The accuracy of the converged elements can be assessed by examing the Cramér-Rao inequality (Maine and Iliff, 1980, 1981b) discussed previously. The value of the Cramér-Rao bound for each unknown is the square root of the corresponding element in the H matrix. A slightly different approach is used to obtain the Cramér-Rao bound for flight-measured data than is used for computed data, because the measurement covariance matrix GG* is unknown for flight data. In general, the expected value of J_{min} is $\ell(N-1)/2$, where ℓ is the dimension of GG* and N is the number of data points. For measured data, if J_{min} is not approximately $\ell(N-1)/2$, then the original estimate of GG* is probably not good. A first approximation to correcting GG* is to multiply GG* by $2J_{min}/\ell(N-1)$, which has the effect of multiplying H by a factor of $2J_{min}/\ell(N-1)$. For the scalar case we are studying here, $\ell = 1$; so the correction term is $2J_{min}/(N-1)$. Since we are treating our simulated data similarly to how actual flight data would be treated, we will make the correction to the bound. Thus, the Cramér-Rao bound can be obtained from the following approximation to the information matrix:

$$H = 2(J_{min})\left(\nabla_\xi^2 J\right)^{-1}/(N-1) \tag{4.3-4}$$

The Cramér-Rao bounds for L_p and L_δ are the square roots of the diagonal elements of the H matrix, or $\sqrt{H(1,1)}$ and $\sqrt{H(2,2)}$, respectively. The Cramér-Rao bounds are 0.1593 and 1.116 for L_p and L_δ, respectively. The errors in $\hat{L}_p$ and $\hat{L}_\delta$ are less than the bounds.

4.4 COST FUNCTIONS

In Section 4.3 we obtain maximum likelihood estimates for computed time histories by minimizing the values of the cost function. To understand fully what occurs in this minimization, we must study in more detail the form of the cost functions and some of their more important characteristics. The same two time histories studied Section 4.3 are examined in this section: The cost function for the no-noise case is discussed briefly; the cost function for the noisy case is then discussed in more detail. The noisy case is more interesting because it has a meaningful Cramér-Rao bound and is more representative of aircraft flight data.

First, we look at the one-dimensional case, where L_δ is fixed at the correct value, because it is easier to grasp some of the characteristics of the cost function in one dimension. Then we look at the two-dimensional case, where both L_p and L_δ are varying. The cost functions are based on computed time histories defined by Equation (4.2-18). For every time history we might choose (computed or flight data), a complete cost function is defined. For the case of n variables, the cost function defines a hypersurface of n + 1 dimensions. Constructing this surface and looking for the minimum (avoiding the use of the minimization algorithm) is not a reasonable approach, because in general the number of variables is greater than two. Therefore, the cost function can be described mathematically but not pictured graphically.

4.4.1 One-Dimensional Case

To illustrate the many interesting aspects of cost functions, it is easiest to look first at cost functions having one variable. This section studies the one-variable cost function $J(L_p)$, with $L_\delta = 10$; this cost function is more interesting than $J(L_\delta)$. Figure 4.4-1 shows the cost function plotted as a function of L_p for the no-noise case. As expected for this case, the minimum cost is zero and occurs at the correct value of $L_p = -0.2500$. The cost increases much more slowly for a more negative L_p than for a positive L_p. Physically this makes sense: The more negative values of L_p represent cases of high damping, and the positive L_p represents an unstable system. Therefore, the p_i for positive L_p becomes increasingly different from the measured time history for small positive increments in L_p.

Figure 4.4-2 shows the no-noise response for the correct L_p (-0.2500) compared with those obtained for L_p values of -0.25 and 0.75. The response for $L_p = -0.7500$ is closer to the correct response than is the response for $L_p = 0.2500$. The differences affect the cost function even more strongly then they affect the responses because the cost function is proportional to the square of the differences in the responses (Equation (4.2-18)). For very large damping (very negative L_p), the system would show essentially no response. Therefore, large increases in damping change the value of $J(L_p)$ relatively little.

In Figure 4.4-3, the cost function based on the noisy case time history is plotted as a function of L_p. The correct L_p value (-0.2500) and the L_p value (-0.3218) at the cost minimum (3.335) are both indicated on the figure. The general shape of the cost function in Figure 4.4-3 is similar to that shown in Figure 4.4-1. In Figure 4.4-4, the cost functions based on the noisy and no-noise cases are compared. The comments relating to the cost function of the no-noise case also apply to the cost function based on the noisy case. Figure 4.4-4 shows clearly that the two cost functions are shaped similarly but shifted in both the L_p and J directions. Only a small difference in the value of the cost would be expected far from the minimum because the estimated time history is so far from the measured time history that it becomes irrelevant as to whether the measured time history has noise added. This point is clearly seen in Figure 4.4-5, where the response for L_p = 0.2500 is compared with the response for the correct values of L_p for the no-noise and noisy cases. The cost function is the sum of the squares of the difference between the measured and estimated responses (Equation (4.2-18)). In some cases the difference at a given point is smaller for the no-noise response than it is for the noisy case, and sometimes it is larger. However, these differences are small compared with the distance between the correct and estimated responses. Therefore, the value of the cost function J is large for both the no-noise and the noisy cases, and the difference in the two cost functions is small compared with the total cost.

Figure 4.4-6 shows the gradient of $J(L_p)$ as a function of L_p for the noisy case. We discuss finding the zero of this function (or equivalently, the minimum of the cost function) using the Gauss-Newton method. The gradient is zero at L_p = -0.3218, which gives the minimum value of $J(L_p)$.

The difference between the Newton-Raphson method (Equation (4.2-16)) and the Gauss-Newton method (Equation (4.2-17)) of minimization is mentioned in Section 4.2. For this simple one-dimensional case, we can easily compute the second gradient both with the second term of Equation (4.2-16) (Newton-Raphson) and without the second term (Gauss-Newton, Equation (4.2-17)). The Newton-Raphson and the Gauss-Newton approximation second gradients are compared in Figure 4.4-7. The Gauss-Newton second gradient (dashed line) always remains positive because it is the sum of quadratic terms (squared for the one-dimensional example). The Newton-Raphson second gradient can be positive or negative, depending upon the value of the second partial derivative with respect to L_p. Other than the difference in sign for the more negative L_p, the two curves have similar shapes.

As stated previously, the Gauss-Newton method is superior to the Newton-Raphson method in certain cases; we can demonstrate obvious cases of this with our example. Problems with the Newton-Raphson method will occur where the second gradient (slope of the gradient) is near zero or negative; Figure 4.4-6 shows such a region near L_p = -1.0. If we choose a point where the gradient slope is exactly zero, we are forced to divide by zero in Equation (4.2-14) with the Newton-Raphson method; such a point is at L_p = -1.13 in Figure 4.4-7. If the value of the slope of the gradient is negative, then the Newton-Raphson method will generate very negative values of L_p. For very negative values of L_p, the cost becomes asymptotically constant, and the gradient becomes nearly zero. In that region, the Newton-Raphson algorithm diverges to negative infinity. If the slope of the gradient is positive but small, we still have a problem with the Newton-Raphson method. Figure 4.4-8 shows the first iteration, starting from L_p = -0.95, for both Gauss-Newton and Newton-Raphson methods. The Newton-Raphson method selects a point where the tangent of the gradient at L_p = -0.95 intersects the zero line. This results in the selection of $L_p \cong 2.6$ in the first iteration. Many iterations are required to progress from that value to the actual minimum. On the other hand, the Gauss-Newton method selects $L_p \cong -0.09$ and converges to the minimum (to four-digit accuracy) in two more iterations. With more complex examples, a comparison of the convergence properties of the two algorithms becomes more difficult to visualize, but the problems are generalizations of those illustrated here.

The usefulness of the Cramér-Rao bound is discussed in Section 4.3-2. We now digress briefly to discuss some of the ramifications of the Cramér-Rao bound for the one-dimensional case. The Cramér-Rao bound has meaning only for the noisy case. In the noisy example, the estimate of L_p is -0.3218 and the Cramér-Rao bound is 0.0579. The calculation of the Cramér-Rao bound is defined in Section 4.3.2 for both the one-dimensional and two-dimensional examples. The scatter in the estimates of L_p should be about the same magnitude as the Cramér-Rao bound. For the one-dimensional case discussed here, the range (L_p = -0.3218 plus or minus the Cramér-Rao bound 0.0579) nearly includes the correct value of L_p = -0.2500. If noisy cases are generated for many time histories (adding different measurement noise to each time history), then the sample mean and sample standard deviation of the estimates for these cases can be calculated. Table 4.4-1 lists the sample mean, sample standard deviation, and the standard deviation of the sample mean (standard deviation divided by the square root of the number of cases) for 5, 10, and 20 cases. The sample mean, as expected, gets closer to the correct value of -0.2500 as the number of cases increases. This is also reflected by the decreasing values in the third column of Table 4.4-1, which lists estimates of the error in the sample mean. The second column of Table 4.4-1 shows the sample standard deviations, which indicate the approximate accuracy of the individual estimates. This standard deviation, which stays more or less constant, is approximately equal to the Cramér-Rao bound for the noisy case being studied here. In fact, the Cramér-Rao bounds for each of the 20 noisy cases (not shown in the table) differ little from the values found for the noisy case being

studied. Both of these results are in good agreement with the theoretical characteristics (Maine and Iliff, 1981b) of the Cramér-Rao bounds and maximum likelihood estimators in general.

These examples indicate the value of obtaining more sample time histories (maneuvers). Having more samples increases confidence in the estimate of the unknowns. This also holds true when analyzing actual flight data. Thus, it is always advisable to obtain data from several maneuvers at a given flight condition to improve the best estimate of each derivative.

The magnitudes of the Cramér-Rao bounds and of the error between the correct and estimated values of L_p are determined to a large extent by the length of the time history and the amount of noise added to the correct time history. For the case being studied, it is apparent from Figure 4.3-3 that the amount of noise being added to the time history is large. The effect of the measurement noise power (GG*, Equations (4.2-6) and (4.2-7)) on the estimate of L_p for the time history is summarized in Table 4.4-2. The estimate of L_p is much improved by decreasing the measurement noise power. A reduction in the value of G to one-tenth of the value in the noisy case being studied yields an acceptable estimate of L_p. For flight data, the measurement noise is reduced by improving the accuracy of the output sensors.

4.4.2 Two-Dimensional Case

In this section the cost function dependent on both L_p and L_δ is studied. The no-noise case is examined first, followed by the noisy case.

Although the cost function is a function of only two unknowns, it is much more difficult to visualize than is the one-dimensional $J(L_p)$. The cost function over a reasonable range of L_p and L_δ is shown in Figure 4.4-9. The cost increases very rapidly in the region of positive L_p and large values of L_δ. The reason for this sharp increase is just an extension of the argument involving positive L_p given in Section 4.4.1. The shape of the surface can be depicted in greater detail if we examine only the values of the cost function $J(L_p,L_\delta) < 200$ for $L_p < 1.0$. Figure 4.4-10 depicts this restricted surface viewed from the upper end of the surface. The minimum must lie in the curving valley that gets broader toward the far side of the surface. Given this picture of the surface, we can look at the isoclines of constant cost on the L_p-L_δ plane (Figure 4.4-11). The steepness of the cost function in the positive L_p direction is once again apparent. The cost function minimum is indicated by a cross. The more nearly elliptical shape inside the closed isocline indicates that the cost function is nearly quadratic there, so fairly rapid convergence would be expected in this region. The L_p axis becomes an asymptote for the cost function as L_δ approaches zero. The cost is constant for $L_\delta = 0$ because no response would result from any aileron input; the estimated response is zero for all values of L_p, resulting in constant cost.

Figure 4.4-12 shows an expanded view of the cost function near its minimum, which (as seen in the earlier example, Table 4.3-3) occurs at the correct values of $L_p = -0.2500$ and $L_\delta = 10$. This is also evident by looking at the cost function surface shown in Figure 4.4-12. The surface has its minimum at the correct values. As expected, the value of the cost function at the minimum is zero.

As in the one-dimensional case, the primary difference between the cost functions for the no-noise and noisy cases is a shift in the cost function. In the one-dimensional case, the cost function for the noisy case shifted so that the minimum was at a higher cost and a more negative value of L_p. In the two-dimensional case, the cost function exhibits a similar shift in both the L_p and L_δ directions. The shift is small enough that the difference is not visible at the scale shown in Figure 4.4-9 or from the perspective of Figure 4.4-10. Figure 4.4-13 shows the isoclines of constant cost for the noisy case, which look much like the isoclines for the no-noise case shown in Figure 4.4-11. The difference is a shift of about 0.1 in L_p. Intuitively, it can be seen that this would hold true for cases with more than two unknowns; the primary difference between the two cost functions is near the minimum.

We next examine the cost function near the minimum. Figure 4.4-14 shows the same view of the cost function for the noisy case as shown in Figure 4.4-12 for the no-noise case. The shape is roughly the same as that in Figure 4.4-12, but the surface is shifted such that its minimum lies over $L_p = -0.3540$ and $L_\delta = 10.24$, and it is shifted upward to a cost function value of approximately 3.3.

To get a more precise idea of the cost function of the noisy case near the minimum, we once again examine the iscolines. The isoclines in this region (Figure 4.4-15) resemble ellipses much more than those in Figures 4.4-11 and 4.4-13. The estimates from the first few iterations in Table 4.3-6 are also shown in Figure 4.4-15. The first iteration ($L = 1$) brought the values of L_p and L_δ very close to the values at the minimum, and the second essentially arrived at the minimum (viewed at this scale). One of the reasons the convergence is so rapid in this region is that the isoclines are nearly elliptical, demonstrating that the cost function is very nearly quadratic in this region. If we had started the Gauss-Newton algorithm at a point where the isoclines are much less elliptical (as in some of the border regions of Figure 4.4-13), the convergence would have progressed more slowly initially, but it would have proceeded at much the same rate as it entered the nearly quadratic region of the cost function.

Before concluding our examination of the two-dimensional case, we will examine the Cramér-Rao bound. Figure 4.4-16 shows the uncertainty ellipsoid, which is based on the Cramér-Rao bounds defined in Sec-

tion 4.3.2. The relationships between the Cramér-Rao bound and the uncertainty ellipsoid are discussed in Maine and Iliff (1981b). The uncertainty ellipsoid almost encloses the correct values of L_p and L_δ. The Cramér-Rao bound for L_p and L_δ can be determined by projecting the uncertainty ellipsoid onto the L_p and L_δ axes. This should give the same values as analytically calculated in Section 4.3.2, which were $L_p = 0.1593$ and $L_\delta = 1.116$.

4.5 ESTIMATION USING FLIGHT DATA

In Sections 4.2 to 4.4 we examine the basic process of obtaining maximum likelihood estimates from computed examples with one or two unknown parameters. Now that these basics have been established, we can explore the estimation of stability and control derivatives from actual flight data. For the computationally much more difficult situation usually encountered using actual flight data, we obtain the maximum likelihood estimates with the Iliff-Maine code (MMLE3 program) (Maine and Iliff, 1980). The equations of motion that are of interest are presented in Chapter 3; the remainder of the equations are given in Maine and Iliff (1980).

In general, flight data estimation is fairly complex, and codes such as the Iliff-Maine code must usually be used to assist in the analysis. However, one must still be cautious about accepting the results; that is, the estimates must fit the phenomenology, and the match between the measured and computed time histories must be acceptable. This is true in all flight regimes, but particularly in potential problem situations such as (1) separated flow at high Mach numbers or high angle of attack, (2) unusual aircraft configurations, such as the oblique wing (Maine, 1978), or (3) modern high-performance aircraft with high-gain feedback loops. In any of these cases, even small anomalies in the match may indicate ignored terms in the equations of motion, separated flow, nonlinearities, sensor problems, or any of a large number of other problems. Some of these difficulties are discussed in later chapters.

The brief examples described in Sections 4.5.1 to 4.5.3 are intended to show how the caveats given in the preceding paragraph and the computed examples of previous sections can be used to assist in the analysis. In the computed example, we show the desirability of having low-noise sensors, an adequate model, and several maneuvers at a given flight condition.

4.5.1 Hand Calculation Example

Sometimes evaluation of a fairly complex flight maneuver can be augmented with a simple hand calculation. One example can be taken from a flight of the space shuttle. The space shuttle's entry control system consists of 12 vertical reaction-control-system (RCS) jets (6 up-firing and 6 down-firing), 8 horizontal RCS jets (4 left-firing and 4 right-firing), 4 elevon surfaces, a body flap, and a split rudder surface. The locations of these devices are shown in Figure 4.5-1. The vertical jets and the elevons are used for both pitch and roll control. The jets and elevons are used symmetrically for pitch control and asymmetrically for roll control.

The shuttle example we use here is obtained from a maneuver at a Mach number of approximately 21 and an angle of attack of approximately 40°. The controls used for this lateral-directional maneuver are the differential elevons and the side-firing jets (yaw jets). The maneuver is depicted in Figure 4.5-2. Equations (3.8-13) to (3.8-15) describe the equations of motion. Some of the derivatives can be determined by hand using a simplified approach that has been used since the beginning of dynamic analysis of flight maneuvers. In particular, for this maneuver, the slope of the rates can be used to determine the yaw jet control derivatives. This is possible, even for this example, which includes a high-gain feedback system, because the yaw jets are essentially step functions, and the slopes of rates p and r can be determined before the vehicle and differential elevon (aileron) responses become significant. The rolling moment due to yaw jet $L_{\gamma J}$ is particularly important for the shuttle and is generally more difficult to obtain than the more dominant yawing moment due to yaw jet. Therefore, as an example, we determine $L_{\gamma J}$ by hand. Figure 4.5-3 shows yaw jets and smoothed roll rate plotted using expanded scales.

The equation for $L_{\gamma J}$ is given by

$$L_{\gamma J} = \frac{\dot{p} I_x}{(\text{number of yaw jets})} \tag{4.5.1}$$

$$\dot{p} \cong \frac{\Delta p}{\Delta t} = \frac{-0.07}{57.3} \div 0.1 \tag{4.5-2}$$

Therefore, given that $I_x \cong 900{,}000$ slug-ft^2 and the number of yaw jets is 4, then $L_{\gamma J} \cong -2750$ ft-lb per jet.

The same maneuver was analyzed with MMLE3, and the resulting time-history comparison is shown in Figure 4.5-4. The match is very good except for a small mismatch in p at about 6 sec. This small mismatch was studied separately with MMLE3 and found to be caused by a nonlinearity in the aileron derivative. The value from MMLE3 for $L_{\gamma J}$ is -2690 ft-lb per jet, which (for the accuracy used here) is

essentially the same value as that obtained by the simplified method. It would be difficult to determine the aileron derivatives as accurately as the yaw jet derivatives. Although accurate estimates can seldom be obtained with the slope method discussed here, rough estimates can usually be obtained to gain some insight into values obtained with MMLE3 (or any other maximum likelihood program). These rough estimates can then be used to help explain unexpected values of estimates from an estimation program.

Sometimes a flight example becomes too complex to hand calculate anything other than qualitative estimates. Determining the rudder derivative for the F-8 aircraft with the yaw augmentation system on illustrates this difficulty. Figure 4.5-5 shows an example of this difficulty for F-8 aircraft data. This example, taken from Shafer (1982a), includes an aileron pulse and a rudder pulse. Although an independent pilot rudder pulse is input during the maneuver, the rudder is primarily responding to the lateral acceleration feedback. When the rudder is moving, several other variables are also changing, which makes it difficult to use the simplified approach. However, $C_{n_{\delta_r}}$ can be roughly estmated when the rudder moves, approximately 1.7 sec from the start of the maneuver. Most of the yaw acceleration is caused by the rudder input, but a poor estimate would be obtained by the hand calculation method.

4.5.2 Cost Function for Full Aircraft Problem

The analysis of a lateral-directional maneuver obtained in flight typically involves 15 to 25 unknown parameters (as shown in Equations (3.8-13) to (3.8-15)), in contrast with one or two unknowns in the simple aircraft example. This makes detailed examples unwieldy and any graphic presentation of the cost function impossible. Therefore, in this section we examine primarily the estimation procedure and the process of minimization.

For a typical flight example, we examine a lateral-directional maneuver, with both aileron and rudder inputs, that has 17 unknown parameters. The data are from the oblique-wing aircraft (Maine, 1978) with the wing at 0° skew during the maneuver. The time history of the data and the subsequent output of MMLE3 are presented in Maine (1981b). The results of the analysis are shown in Table 4.5-1. The match between the measured time history and the estimated (calculated) time history is shown as a function of iteration in Figure 4.5-6. Figures 4.5-6(a) to 4.5-6(e) correspond to iterations 0 to 4. Table 4.5-1 shows that the cost remains unchanged after four iterations. A similar result was obtained for the two-dimensional simple aircraft example in Figure 4.3-4 and Table 4.3-3.

Two of the many things the analyst must consider in obtaining estimates are how good the match is and how good the convergence is. A satisfactory match and monotonic convergence are desirable, but not sufficient, conditions for a successful analysis. Figure 4.5-6(e) is a very good match, although not perfect. The convergence can best be evaluated by examining the normalized cost in the last row of Table 4.5-1; the cost rapidly and monotonically converges in four iterations, and it remains at the converged value. These factors are convincing evidence that convergence is complete. Therefore, the match and convergence criteria are satisfied in our example. In some cases we might encounter cost that does not converge rapidly (in four to six iterations) or monotonically, or stay "exactly" at the minimum value. These conditions usually indicate that there is at least a small problem in the analysis. These problems can usually be traced to a data problem, an inadequate mathematical model, or a maneuver that contains a marginal amount of information.

Table 4.5-1 also shows that the starting values of all coefficients for the control and bias variables are zero. Wind-tunnel estimates could have been used for starting values, but the convergence of the algorithm is not very dependent on the starting values of these coefficients. As part of the startup algorithm, the MMLE3 program normally holds the derivatives of the state variables constant until after the first iteration, as is evident from Table 4.5-1.

Figure 4.5-6(a) shows a comparison of the measured and computed data for the starting values. The match is very poor because the starting values for the control derivatives are all zero, so the only motion is in response to the initial conditions. The control derivatives and biases are determined on the first iteration, resulting in the much improved match shown in Figure 4.5-6(b). The match after two iterations, shown in Figure 4.5-6(c), is improved as the program further modifies the control derivatives and, for the first time, adjusts the derivatives affecting the natural frequency (C_{n_β} and C_{ℓ_β}). By the third iteration (Figure 4.5-6(d)), where minor adjustments to the frequency are made and the damping derivatives are changed, the match is almost perfect. Figure 4.5-6(e) shows the match when all but the most minor derivatives have ceased to change.

Several general observations can be made based on this well-behaved example. The most important coefficients essentially converged in three iterations. The same effect was seen in the simple example; that is, L_δ converged faster than L_p (Table 4.3-6). Some of the less important or second-order coefficients converged only to two significant digits after three iterations and were still changing in the fourth significant digit at the end of six iterations. Also, even though the sign is wrong after the first iteration for some coefficients (C_{ℓ_r}, $C_{n_{\delta_a}}$, and $C_{\ell_{\delta_r}}$), the algorithm quickly selects reasonable values once the important derivatives have stabilized.

In general, if the analysis of a maneuver goes well, we do not need to spend much time inspecting a table analogous to Table 4.5-1. However, if there are problems in convergence or in the quality of the

fit, a detailed inspection of such a table may be necessary. The table may show that an important coefficient becomes unstable in early iterations, which could cause problems later. If the starting values are grossly in error, the algorithm may generate very unreasonable values and then, for many reasons, not behave well. Occasionally the algorithm selects alternately from two diverse sets of values of two or more coefficients on successive iterations, behaving as if the shape of the cost function were a narrow multidimensional valley analogous to, but more extreme than, the two-dimensional valley shown in Figures 4.4-13 and 4.4-15.

4.5.3 Cramér-Rao Bounds

Sections 4.3 and 4.4 regarding the computed example show that Cramér-Rao bounds are good indicators of the accuracy of estimated parameters. Cramér-Rao bounds can be used in a similar, but somewhat more qualitative, fashion on flight data. The Cramér-Rao bounds included in MMLE3 (as well as in many other maximum likelihood estimation programs) have been useful in determining whether estimates are good or bad. The aircraft example we discuss here has been reported previously (for example, in Iliff (1978b) and Maine and Iliff (1981b)); however, this example of the use of the Cramér-Rao bound in the assessment of flight-derived estimates is pertinent to this document. Figure 4.5-7 shows estimates of C_{n_p} as a function of angle of attack for the PA-30 twin-engine general aviation aircraft (Fink and Freeman, 1969) at three flap settings. There is significant scatter, which reduces confidence in the C_{n_p} information. The data shown are the estimates from the MMLE3 program, which also provides the Cramér-Rao bounds for each estimate. Experience has shown (Iliff, 1978b) that if the Cramér-Rao bound is multiplied by a scale factor (the result sometimes being called the uncertainty level (Maine and Iliff, 1981b) and plotted as a vertical bar with the associated estimate, then the interpretation of flight-determined results is facilitated. Figure 4.5-8 shows the same data as Figure 4.5-7, with the uncertainty levels included as vertical bars. The best estimates are those with the smallest uncertainty levels (Cramér-Rao bounds). The fairing shown in Figure 4.5-8 is drawn through the estimates with small Cramér-Rao bounds and ignores the estimates with large bounds. One can have great confidence in the fairing because it is well defined and consistent when the Cramér-Rao bound information is included. In this particular instance, the estimates with small bounds were from maneuvers where the aileron forced the motion, and those with large bounds were from maneuvers where the rudder forced the motion. Therefore, in addition to aiding in the fairing of the estimates, the Cramér-Rao bounds help show that aileron-forced maneuvers are superior for estimating C_{n_p} for the PA-30 aircraft.

This example illustrates that Cramér-Rao bounds are useful tools for assessing flight-determined estimates, just as they were found useful for the simple aircraft example with computed data.

4.6 SUMMARY

The computed simple aircraft example shows the basics of minimization and the general concepts of cost functions themselves. In addition, this simple example demonstrates the advantages of low measurement noise and multiple estimates at a given condition, the use of Cramér-Rao bounds, and the quality of the match between measured and computed data. The flight data show that many of these concepts apply even though the multidimensional cost function is impossible to plot or visualize.

TABLE 4.3-1. — VALUES OF ROLL-RATE TIME HISTORY WITH NO MEASUREMENT NOISE

i	δ, deg	p, deg/sec
1	0	0
2	1	0.9754115099857
3	1	2.878663149266
4	1	4.689092110779
5	1	6.411225409939
6	1	8.049369277012
7	1	9.607619924937
8	0	10.11446228200
9	0	9.621174135646
10	0	9.151943936071

TABLE 4.3-2. — INTERMEDIATE VALUES FOR CALCULATING TERMS IN EQUATION (4.2-14) WITH NO MEASUREMENT NOISE

	1	2	3	4	5	6	7	8	9	10	11	12
	i	$\delta_{i+1/2}$	p_i	$\tilde{p}_i$	J_i	$\frac{\partial \tilde{p}_i}{\partial \hat{L}_p}$	$\frac{\partial \tilde{p}_i}{\partial \hat{L}_\delta}$	$\nabla J[0, t]$		$\nabla^2 J[0, t]$		
								$\partial \hat{L}_p$	$\partial \hat{L}_\delta$	$\partial \hat{L}_p^2$	$\partial \hat{L}_\delta^2$	$\partial \hat{L}_p \partial \hat{L}_\delta$
L = 1	1	0.5	0	0	0	0	0	0	0	0	0	0
	2	1	0.9754	1.427	0.01022	0.1358	0.09516	-0.06140	-0.04302	0.01845	0.009956	0.01293
	3	1	2.879	4.146	0.9058	0.6533	0.2764	-0.8897	-0.3935	0.4453	0.08547	0.1935
	4	1	4.689	6.607	2.745	1.614	0.4405	-3.986	-1.238	3.052	0.2795	0.9046
	5	1	6.411	8.833	5.677	2.930	0.5889	-11.08	-2.664	11.64	0.6262	2.630
	6	1	8.049	10.85	9.592	4.524	0.7231	-23.74	-4.687	32.10	1.149	5.902
	7	0.5	9.608	12.67	14.28	6.332	0.8447	-43.13	-7.274	72.19	1.863	11.25
	8	0	10.11	12.89	18.14	8.161	0.8594	-65.79	-9.661	138.8	2.601	18.26
	9	0	9.621	11.66	20.22	9.722	0.7777	-85.66	-11.25	233.3	3.206	25.82
	10	0	9.152	10.55	21.21	10.91	0.7037	-101.0	-12.24	352.4	3.701	33.50
L = 2	1	0.5	0	0	0	0	0	0	0	0	0	0
	2	1	0.9754	0.9596	1.124×10^{-4}	0.09314	0.09705	0.001468	0.001530	0.008675	0.009420	0.009040
	3	1	2.879	2.823	0.001675	0.4548	0.2855	0.02680	0.01743	0.2155	0.09093	0.1389
	4	1	4.689	4.578	0.007890	1.147	0.4630	0.1546	0.06905	1.530	0.3053	0.6697
	5	1	6.411	6.230	0.02433	2.129	0.6301	0.5406	0.1833	6.061	0.7022	2.011
	6	1	8.049	7.786	0.05906	3.365	0.7874	1.427	0.3908	17.38	1.322	4.660
	7	0.5	9.608	9.251	0.1226	4.822	0.9355	3.147	0.7245	40.63	2.198	9.172
	8	0	10.11	9.671	0.2209	6.377	0.9781	5.975	1.158	81.30	3.154	15.41
	9	0	9.621	9.107	0.3531	7.828	0.9210	10.00	1.632	142.6	4.003	22.62
	10	0	9.152	8.576	0.5191	9.087	0.8673	15.24	2.131	225.2	4.755	30.55
L = 3	1	0.5	0	0	0	0	0	0	0	0	0	0
	2	1	0.9754	0.9753	8.095×10^{-9}	0.09515	0.09756	1.210×10^{-5}	1.242×10^{-5}	0.009054	0.009519	0.009284
	3	1	2.879	2.879	1.167×10^{-8}	0.4666	0.2880	-2.737×10^{-5}	-1.195×10^{-5}	0.2267	0.09245	0.1437
	4	1	4.689	4.690	7.128×10^{-7}	1.183	0.4692	-0.001428	-5.676×10^{-4}	1.625	0.3126	0.6985
	5	1	6.411	6.417	5.471×10^{-6}	2.209	0.6417	-0.008241	-0.002547	6.504	0.7243	2.116
	6	1	8.049	8.055	2.175×10^{-5}	3.514	0.8058	-0.02829	-0.007145	18.85	1.374	4.947
	7	0.5	9.608	9.617	6.201×10^{-5}	5.068	0.9620	-0.07377	-0.01578	44.54	2.299	9.823
	8	0	10.11	10.13	1.458×10^{-4}	6.750	1.013	-0.1612	-0.02889	90.10	3.326	16.66
	9	0	9.621	9.638	2.919×10^{-4}	8.352	0.9642	-0.3039	-0.04537	159.9	4.255	24.71
	10	0	9.152	9.173	5.083×10^{-4}	9.784	0.9176	-0.5075	-0.06447	255.6	5.097	33.69
L = 4	1	0.5	0	0	0	0	0	0	0	0	0	0
	2	1	0.9754	0.9754	1.084×10^{-11}	0.09514	0.09754	4.429×10^{-7}	4.541×10^{-7}	0.009052	0.009514	0.009280
	3	1	2.879	2.879	9.377×10^{-11}	0.4664	0.2879	6.450×10^{-6}	4.162×10^{-6}	0.2266	0.09238	0.1436
	4	1	4.689	4.689	2.748×10^{-10}	1.182	0.4689	2.894×10^{-5}	1.308×10^{-5}	1.623	0.3123	0.6977
	5	1	6.411	6.411	5.457×10^{-10}	2.207	0.6411	8.031×10^{-5}	2.801×10^{-5}	6.494	0.7233	2.113
	6	1	8.049	8.049	8.787×10^{-10}	3.510	0.8049	1.709×10^{-4}	4.878×10^{-5}	18.81	1.371	4.938
	7	0.5	9.608	9.608	1.237×10^{-9}	5.061	0.9608	3.064×10^{-4}	7.450×10^{-5}	44.43	2.294	9.800
	8	0	10.11	10.11	1.471×10^{-9}	6.738	1.011	4.523×10^{-4}	9.640×10^{-5}	89.82	3.317	16.82
	9	0	9.621	9.621	1.539×10^{-9}	8.334	0.9621	5.497×10^{-4}	1.076×10^{-4}	159.3	4.243	24.63
	10	0	9.152	9.152	1.543×10^{-9}	9.759	0.9152	5.754×10^{-4}	1.101×10^{-4}	254.5	5.081	33.57

TABLE 4.3-3. — PERTINENT VALUES AS A FUNCTION OF ITERATION

L	$\hat{L}_p(L)$	$\hat{L}_\delta(L)$	$\Phi(L)$	$\Psi(L)$	J_L
0	-0.5000	15.00	0.9048	2.855	21.21
1	-0.3005	9.888	0.9417	1.919	0.5191
2	-0.2475	9.996	0.9517	1.951	5.083×10^{-4}
3	-0.2500	10.00	0.9512	1.951	1.543×10^{-9}
4	-0.2500	10.00	0.9512	1.951	1.062×10^{-14}

TABLE 4.3-4. — VALUES OF COMPUTED TIME HISTORY WITH ADDED MEASUREMENT NOISE

i	δ, deg	p, deg/sec
1	0	0
2	1	0.4875521781881
3	1	3.238763570696
4	1	3.429117357944
5	1	6.286297353361
6	1	6.953798550097
7	1	10.80572930119
8	0	9.739367269447
9	0	9.788844525490
10	0	7.382568353168

TABLE 4.3-5. — INTERMEDIATE VALUES FOR CALCULATING TERMS IN EQUATION (4.2-14) WITH MEASUREMENT NOISE

	1	2	3	4	5	6	7	8	9	10	11	12
	i	$\delta_{i+1/2}$	p_i	$\tilde{p}_i$	J_i	$\frac{\partial \tilde{p}_i}{\partial \hat{L}_p}$	$\frac{\partial \tilde{p}_i}{\partial \hat{L}_\delta}$	$\nabla J[0, t]$		$\nabla^2 J[0, t]$		
								$\partial \hat{L}_p$	$\partial \hat{L}_\delta^2$	$\partial \hat{L}_p^2$	$\partial \hat{L}_\delta^2$	$\partial \hat{L}_p \partial \hat{L}_\delta$
L = 1	1	0.5	0	0	0	0	0	0	0	0	0	0
	2	1	0.4876	1.427	0.4417	0.1358	0.09516	-0.1277	-0.08944	0.01845	0.009055	0.01293
	3	1	3.239	4.146	0.8537	0.6533	0.2764	-0.7207	-0.3404	0.4453	0.08547	0.1935
	4	1	3.429	6.607	5.902	1.614	0.4405	-5.851	-1.740	3.052	0.2795	0.9046
	5	1	6.286	8.833	9.145	2.930	0.5889	-13.31	-3.240	11.64	0.6262	2.630
	6	1	6.954	10.85	16.72	4.524	0.7231	-30.93	-6.055	32.10	1.149	5.902
	7	0.5	10.81	12.67	18.46	6.332	0.8447	-42.73	-7.63	72.19	1.863	11.25
	8	0	9.739	12.89	23.43	8.161	0.8594	-68.46	-10.34	138.8	2.601	18.26
	9	0	9.789	11.66	25.19	9.722	0.7777	-86.69	-11.80	233.3	3.206	25.82
	10	0	7.383	10.55	30.22	10.91	0.7037	-121.3	-14.03	352.4	3.701	33.50
L = 2	1	0.5	0	0	0	0	0	0	0	0	0	0
	2	1	0.4876	0.9781	0.1203	0.09414	0.09626	-0.04618	-0.04721	0.008863	0.009265	0.009062
	3	1	3.239	2.862	0.1913	0.4568	0.2816	0.1260	0.05895	0.2175	0.08859	0.1377
	4	1	3.429	4.606	0.8842	1.142	0.4533	-1.218	-0.4747	1.521	0.2941	0.6553
	5	1	6.286	6.222	0.8863	2.100	0.6123	-1.082	-0.4351	5.930	0.6690	1.941
	6	1	6.954	7.718	1.178	3.286	0.7595	-3.592	-1.015	16.73	1.246	4.437
	7	0.5	10.81	9.103	2.628	4.662	0.8959	4.346	0.5102	38.46	2.048	8.613
	8	0	9.739	9.408	2.683	6.099	0.9259	6.369	0.8172	75.66	2.906	14.26
	9	0	9.789	8.712	3.263	7.392	0.8574	14.33	1.741	130.3	3.641	20.60
	10	0	7.383	8.068	3.497	8.460	0.7940	8.534	1.197	201.9	4.271	27.31
L = 3	1	0	0	0	0	0	0	0	0	0	0	0
	2	1	0.4876	0.9882	0.1253	0.09544	0.09656	-0.04778	-0.04835	0.009106	0.009324	0.009324
	3	1	3.239	2.898	0.1835	0.4642	0.2831	0.1106	0.04825	0.2246	0.08949	0.1406
	4	1	3.429	4.677	0.9624	1.164	0.4570	-1.342	-0.5221	1.580	0.2984	0.6726
	5	1	6.286	6.336	0.9636	2.148	0.6191	-1.449	-0.5528	6.195	0.6816	2.003
	6	1	6.954	7.882	1.394	3.375	0.7702	-4.582	-1.268	17.59	1.275	4.602
	7	0.5	10.81	9.323	2.494	4.807	0.9110	2.547	0.08320	40.70	2.105	8.982
	8	0	9.739	9.678	2.499	6.316	0.9456	2.937	0.1415	80.59	2.999	14.95
	9	0	9.789	9.020	2.791	7.692	0.8814	8.850	0.8190	139.8	3.776	21.73
	10	0	7.383	8.407	3.316	8.852	0.8215	-0.2214	-0.002276	218.1	4.450	29.01
L = 4	1	0.5	0	0	0	0	0	0	0	0	0	0
	2	1	0.4876	0.9891	0.1258	0.09549	0.09654	-0.04789	-0.04842	0.009117	0.009320	0.009218
	3	1	3.239	2.900	0.1833	0.4644	0.2830	0.1096	0.04757	0.2247	0.08942	0.1406
	4	1	3.429	4.679	0.9649	1.164	0.4567	-1.346	-0.5235	1.580	0.2980	0.6724
	5	1	6.286	6.337	0.9662	2.148	0.6186	-1.456	-0.5551	6.195	0.6806	2.001
	6	1	6.954	7.882	1.397	3.374	0.7693	-4.588	-1.269	17.58	1.272	4.597
	7	0.5	10.81	9.321	2.499	4.804	0.9098	2.545	0.08155	40.66	2.100	8.967
	8	0	9.739	9.672	2.501	6.309	0.9441	2.967	0.1448	80.46	2.991	14.92
	9	0	9.789	9.011	2.804	7.681	0.8795	8.944	0.8291	139.5	3.765	21.68
	10	0	7.383	8.394	3.316	8.836	0.8193	0.004819	2.057×10^{-4}	217.5	4.436	28.92

TABLE 4.3-6. — PERTINENT VALUES AS A FUNCTION OF ITERATION

L	$\hat{L}_p(L)$	$\hat{L}_\delta(L)$	$\Phi(L)$	$\Psi(L)$	J_L
0	-0.5000	15.00	0.9048	2.855	30.22
1	-0.3842	10.16	0.9260	1.956	3.497
2	-0.3518	10.23	0.9321	1.976	3.316
3	-0.3543	10.25	0.9316	1.978	3.316
4	-0.3542	10.24	0.9316	1.978	3.316
5	-0.3542	10.24	0.9316	1.978	3.316

TABLE 4.4-1. — MEAN AND STANDARD DEVIATIONS FOR ESTIMATES OF L_p

Number of cases, N	Sample mean, $\mu(\hat{L}_p)$	Sample standard deviation, $\sigma(\hat{L}_p)$	Sample standard deviation of the mean, $\sigma(\hat{L}p)/\sqrt{N}$
5	-0.2668	0.0739	0.0336
10	-0.2511	0.0620	0.0196
20	-0.2452	0.0578	0.0129

TABLE 4.4-2. — ESTIMATE OF L_p AND CRAMÉR-RAO BOUND AS A FUNCTION OF THE SQUARE ROOT OF NOISE POWER

Square root of noise power, G	Estimate of L_p	Cramér-Rao bound
0	-0.2500	---
0.01	-0.2507	0.00054
0.05	-0.2535	0.00271
0.1	-0.2570	0.00543
0.2	-0.2641	0.0109
0.4	-0.2783	0.0220
0.8	-0.3071	0.0457
1.0	-0.3218	0.0579
2.0	-0.3975	0.1248
5.0	-0.6519	0.3980
10.0	-1.195	1.279

TABLE 4.5-1. – STABILITY AND CONTROL DERIVATIVES AS FUNCTIONS OF ITERATION FOR FLIGHT MANEUVER

	ITERATION, L						
	0	1	2	3	4	5	6
C_{Y_β}	-0.008500	-0.008500	-0.007959	-0.008347	-0.008375	-0.008364	-0.008364
C_{ℓ_β}	-0.0002500	-0.0002500	-0.0003141	-0.0003582	-0.0003572	-0.0003571	-0.0003571
C_{n_β}	0.001000	0.001000	0.001159	0.001243	0.001230	0.001230	0.001230
C_{ℓ_p}	-0.2500	-0.2500	-0.3393	-0.3584	-0.3581	-0.3581	-0.3581
C_{n_p}	-0.02500	-0.02500	-0.04356	-0.04537	-0.04612	-0.04599	-0.04600
C_{ℓ_r}	-0.05000	-0.05000	0.06790	0.07044	0.06972	0.06973	0.06974
C_{n_r}	-0.08000	-0.08000	-0.1327	-0.1033	-0.1065	-0.1062	-0.1062
$C_{\ell_{\delta_a}}$	0	0.0008009	0.001000	0.001067	0.001069	0.001069	0.001069
$C_{n_{\delta_a}}$	0	-0.00004604	6.786×10^{-7}	0.00001129	0.00001096	0.00001068	0.00001069
$C_{Y_{\delta_r}}$	0	0.005935	0.002064	0.001456	0.001546	0.001548	0.001548
$C_{\ell_{\delta_r}}$	0	-0.00005068	0.00005764	0.0001043	0.0001059	0.0001055	0.0001055
$C_{n_{\delta_r}}$	0	-0.0007329	-0.0009333	-0.0008875	-0.0008972	-0.0008961	-0.0008961
C_{Y_b}	0	-0.05109	-0.02691	-0.02362	-0.02420	-0.02419	-0.02420
$C_{Y_b} + \dot{\beta}_b$	0	-0.03370	-0.01370	-0.01117	-0.01157	-0.01156	-0.01156
C_{ℓ_b}	0	-0.0007096	-0.001629	-0.002021	-0.002031	-0.002028	-0.002028
C_{n_b}	0	0.005864	.007300	0.007140	0.007175	0.007169	0.007169
$\dot{\phi}_b$	0	0.2121	0.1626	0.1482	0.1506	0.1506	0.1506
J/(N - 1)	731.5	65.00	11.23	4.826	4.701	4.701	4.701

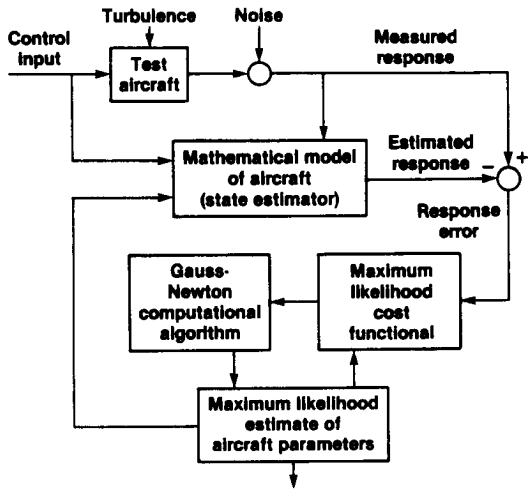

Figure 4.1-1. Maximum likelihood estimation concept.

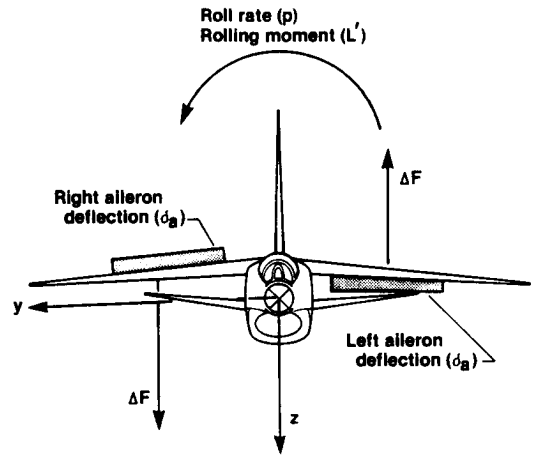

Figure 4.2-1. Simplified aircraft nomenclature.

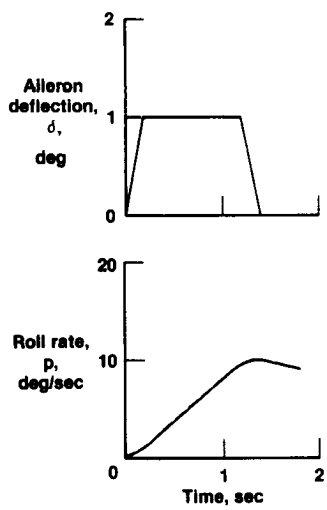

Figure 4.3-1. Time history with no measurement noise.

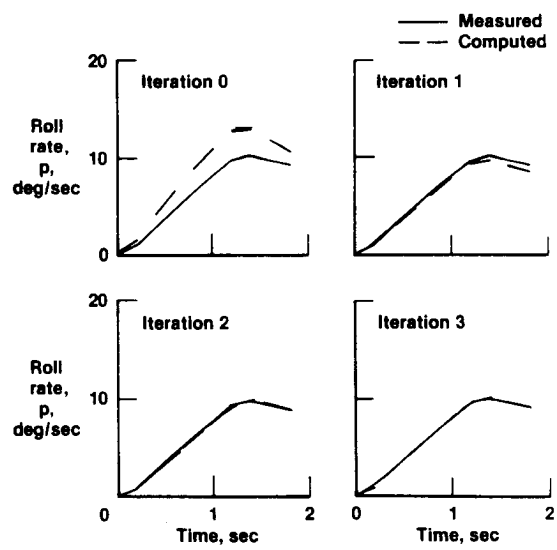

Figure 4.3-2. Comparison of measured and computed data for each of the first three iterations.

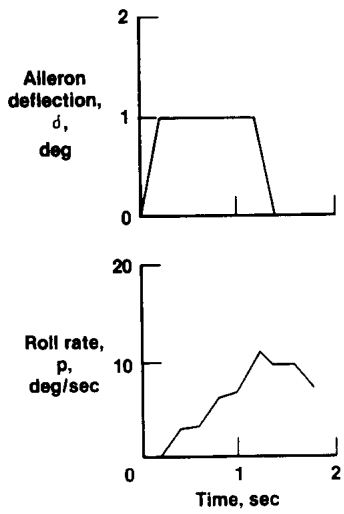

Figure 4.3-3. Time history with measurement noise.

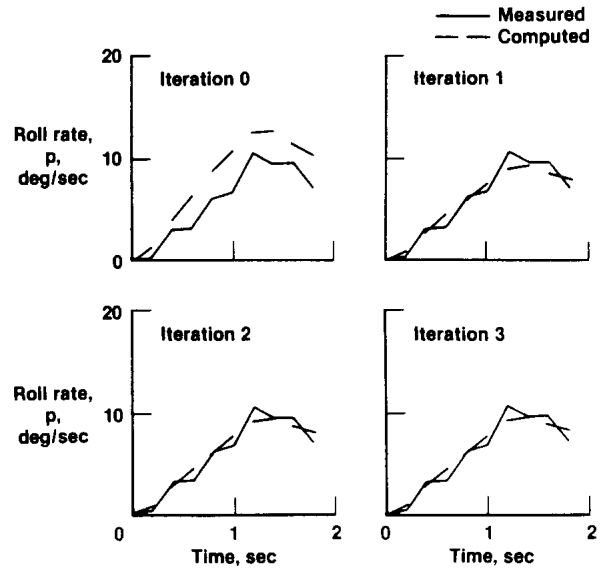

Figure 4.3-4. Comparison of measured and computed data for each iteration.

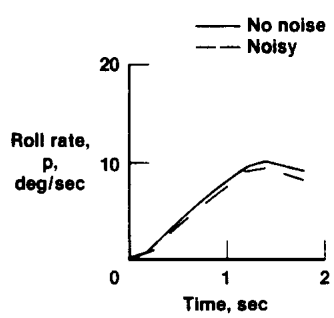

Figure 4.3-5. Comparison of estimated roll rate from no-noise and noisy cases.

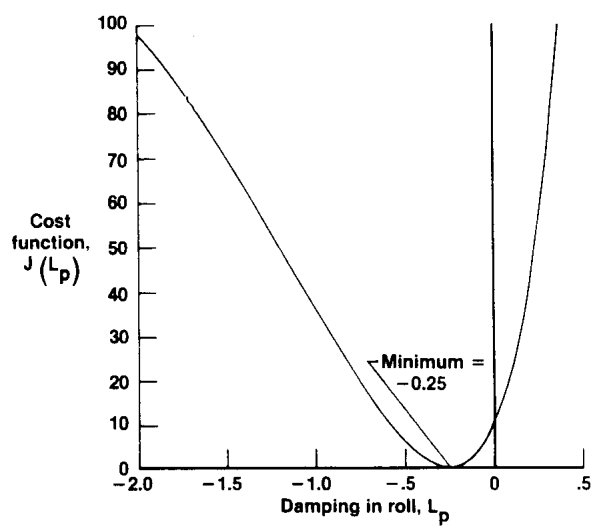

Figure 4.4-1. Cost function $J(L_p)$ as a function of L_p for no-noise case.

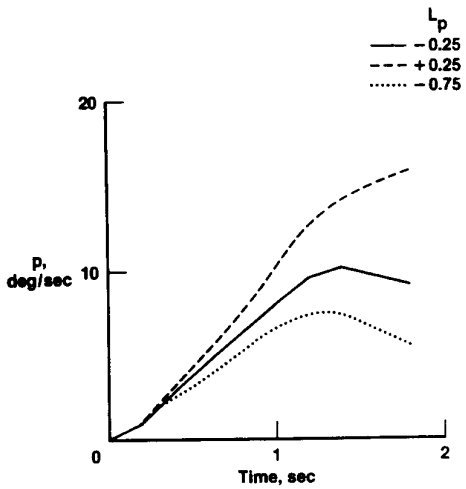

Figure 4.4-2. No-noise response for correct L_p (-0.2500) compared with the response for L_p = -0.75 and L_p = 0.25.

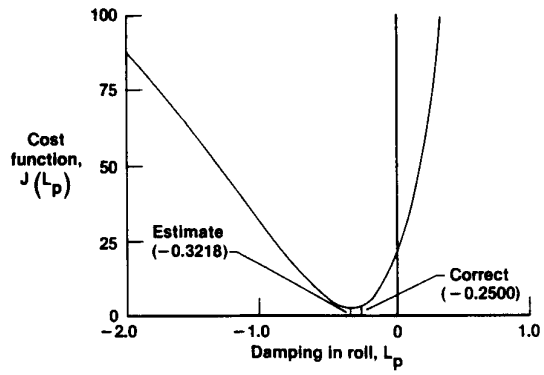

Figure 4.4-3. Cost function as a function of L_p for noisy case.

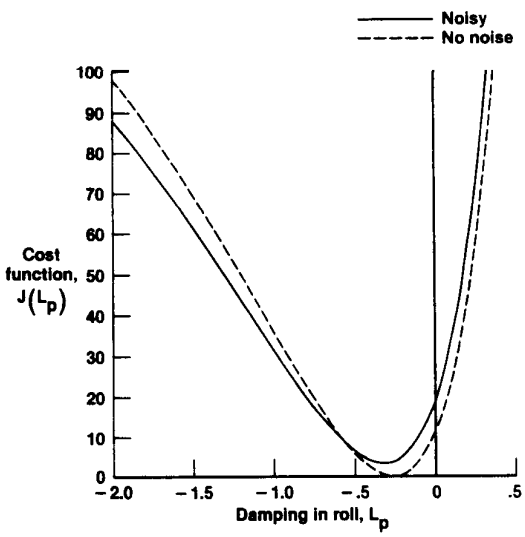

Figure 4.4-4. Comparison of the cost functions for the no-noise and noisy cases.

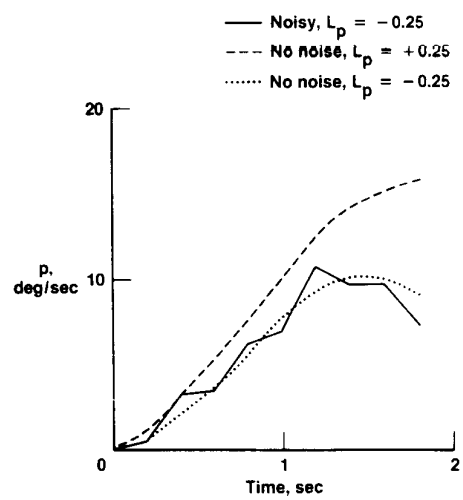

Figure 4.4-5. Response for an L_p of +0.25 compared with the no-noise and noisy response for L_p at the correct value of -0.25.

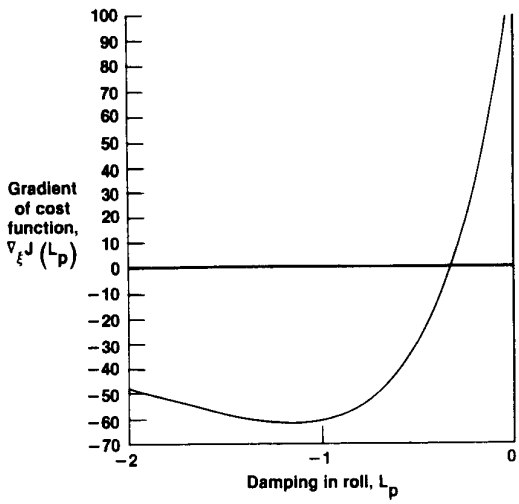

Figure 4.4-6. Gradient of $J(L_p)$ as a function of L_p for noisy case.

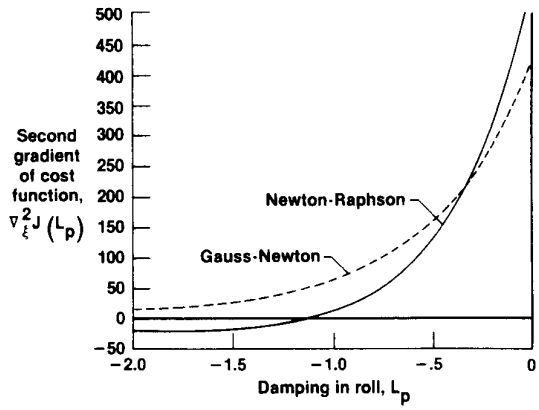

Figure 4.4-7. Comparison of Newton-Raphson and Gauss-Newton values of the second gradient for the noisy case.

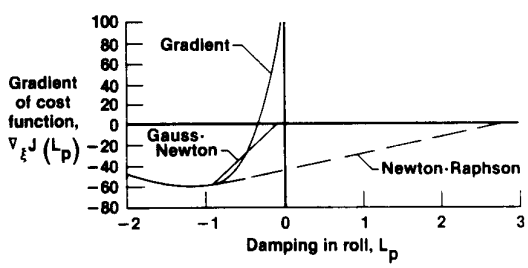

Figure 4.4-8. Comparison of first iteration step size for the Newton-Raphson and Gauss-Newton algorithms for the noisy case.

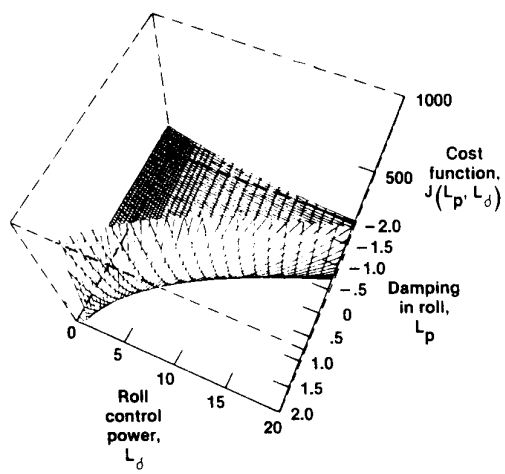

Figure 4.4-9. Large-scale view of cost function surface.

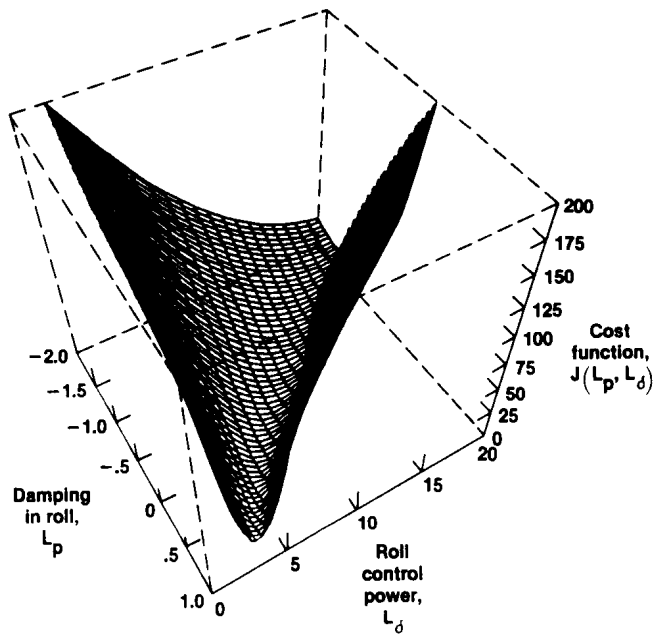

Figure 4.4-10. Restricted view of cost function surface.

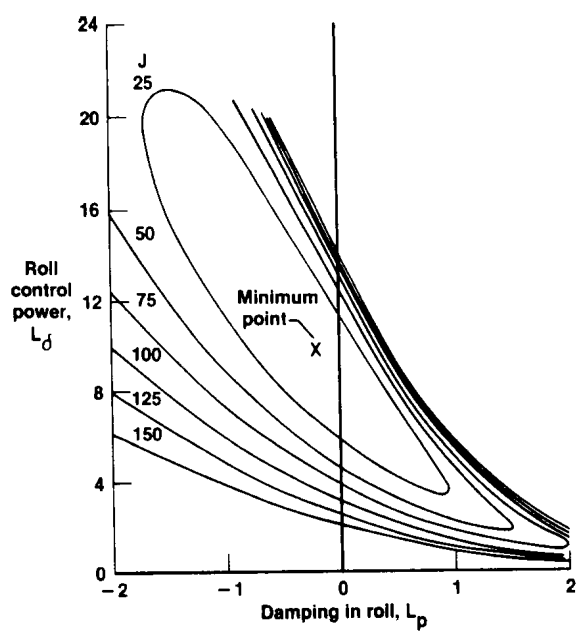

Figure 4.4-11. Isoclines of constant cost of L_p and L_δ for the no-noise case.

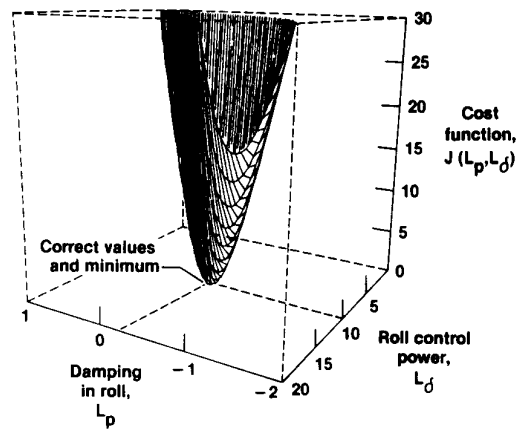

Figure 4.4-12. Detailed view of cost function surface for no-noise case.

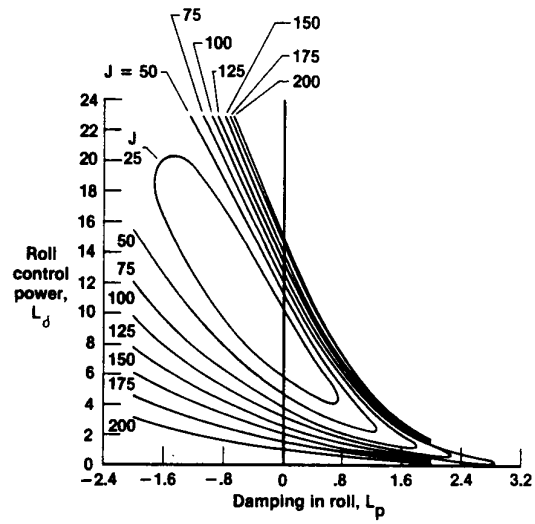

Figure 4.4-13. Isoclines of constant cost in L_p and L_δ for the noisy case.

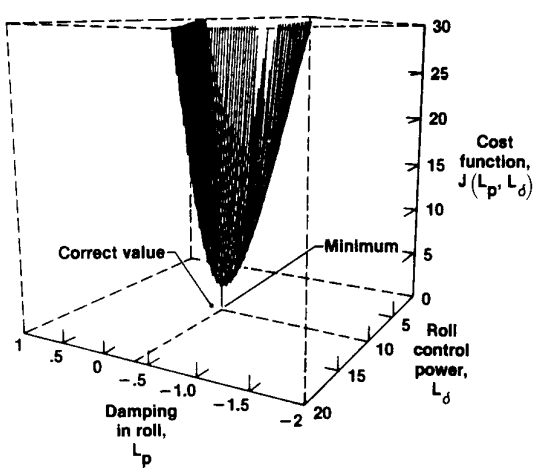

Figure 4.4-14. Detailed view of cost function surface for noisy case.

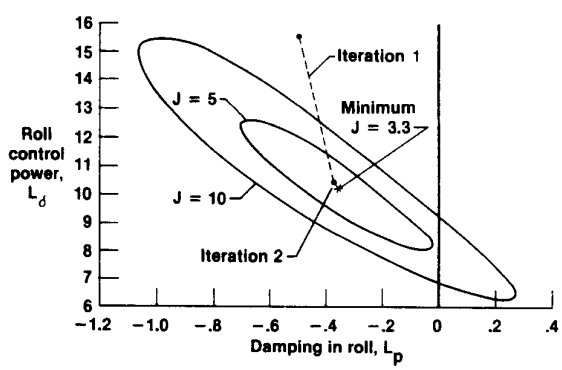

Figure 4.4-15. Isoclines of constant cost for region near minimum for noisy case.

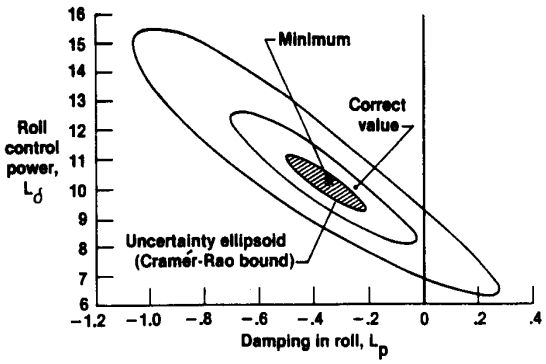

Figure 4.4-16. Isoclines and uncertainty ellipsoid of the cost function for the noisy case.

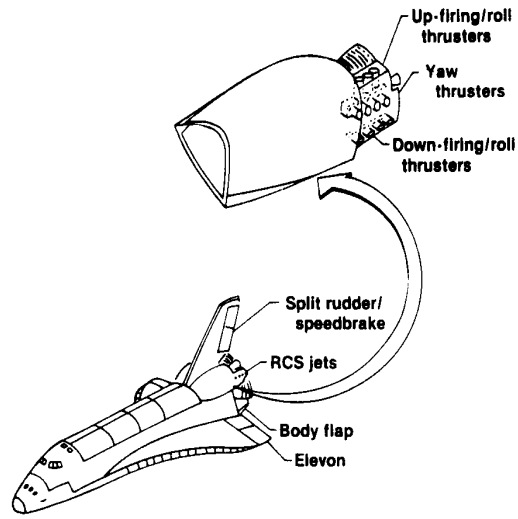

Figure 4.5-1. Space shuttle configuration.

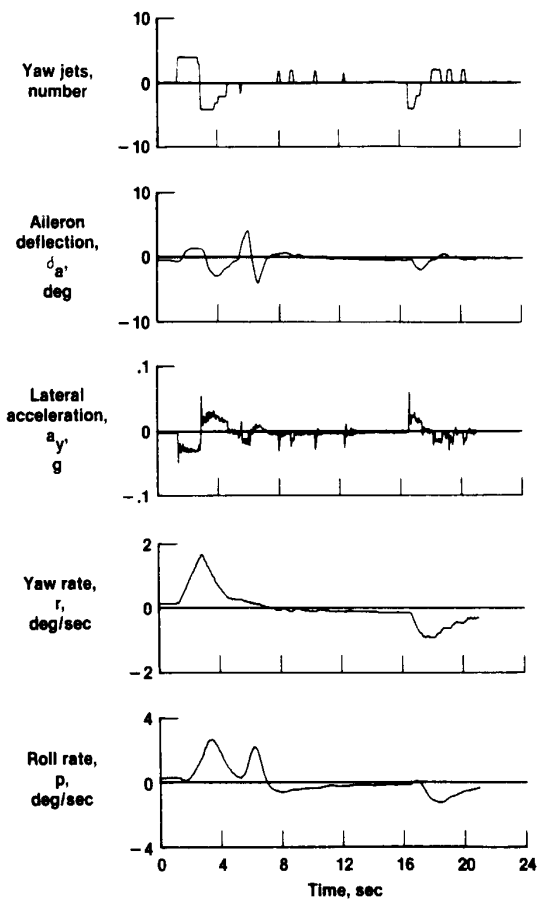

Figure 4.5-2. Lateral-directional space shuttle maneuver at a Mach number of 21.

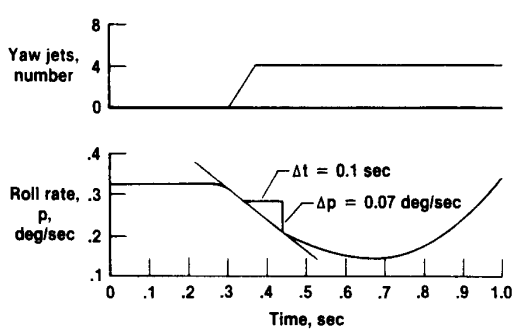

Figure 4.5-3. Examples of obtaining L_{Y_J} by simple calculations for the shuttle data from Figure 4.5-2.

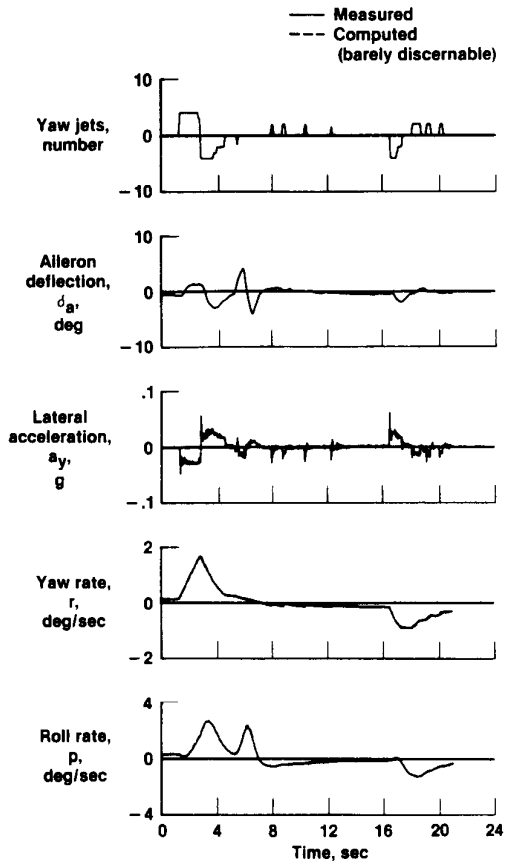

Figure 4.5-4. MMLE3 match of maneuver shown in Figure 4.5-2.

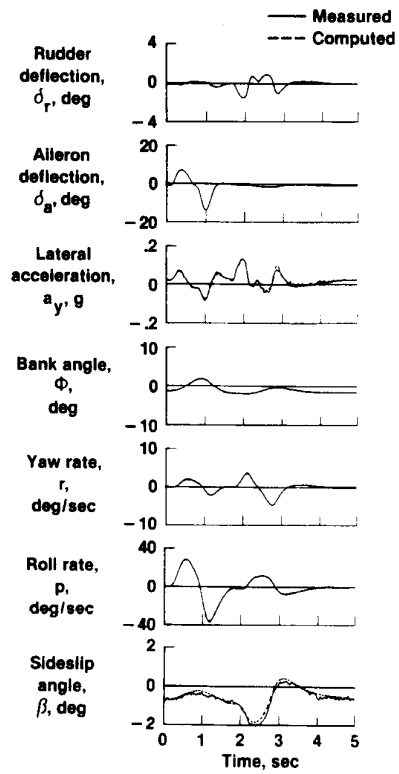

Figure 4.5-5. Lateral-directional maneuver from F-8 aircraft with yaw augmentation on.

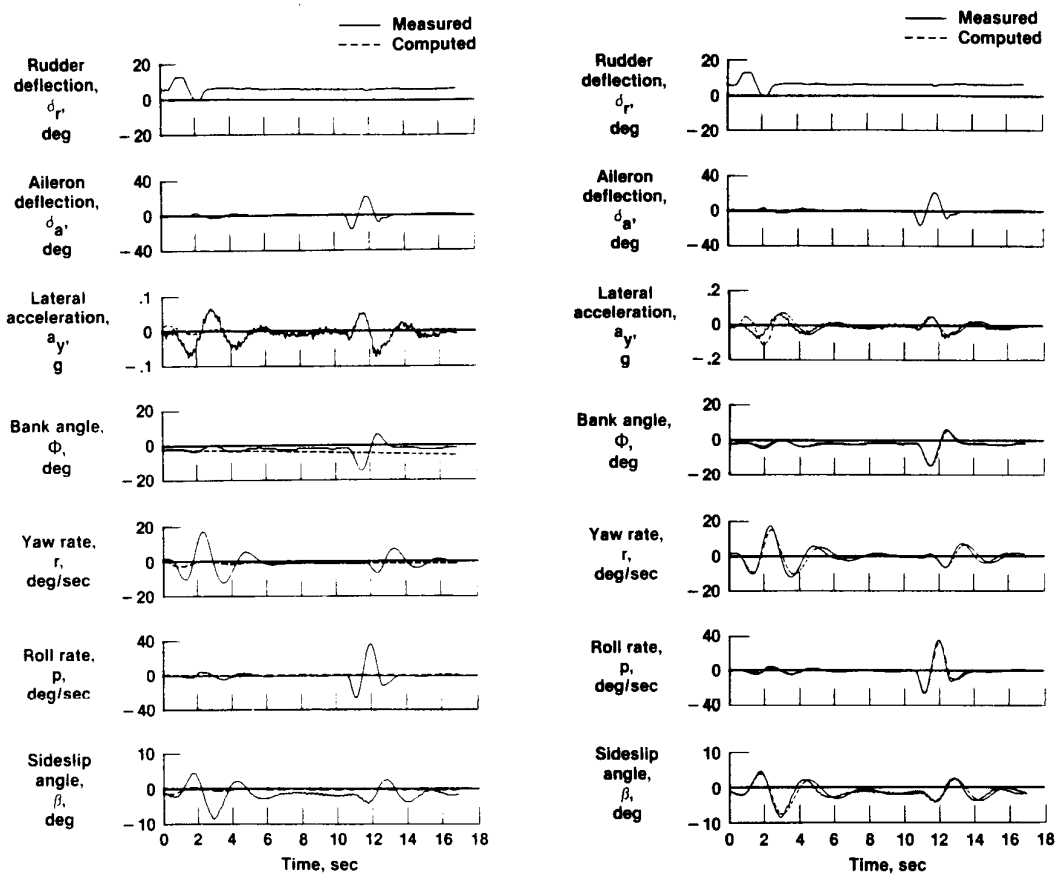

(a) Zero iteration.

(b) One iteration.

Figure 4.5-6. Match between measured and computed time histories as functions of iteration.

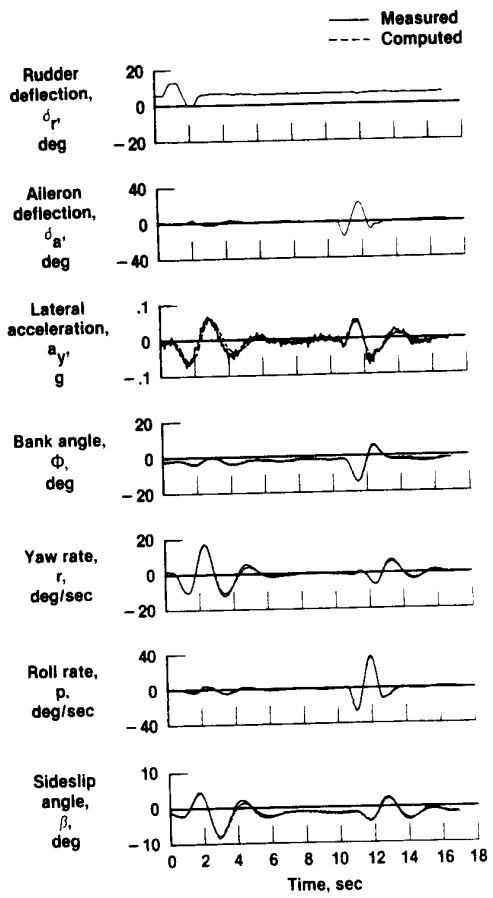

(c) Two iterations.

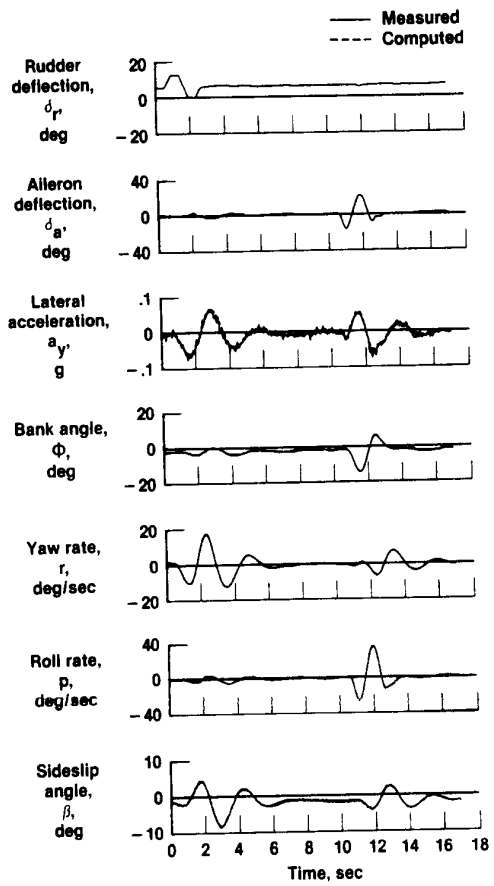

(d) Three iterations.

Figure 4.5-6. Continued.

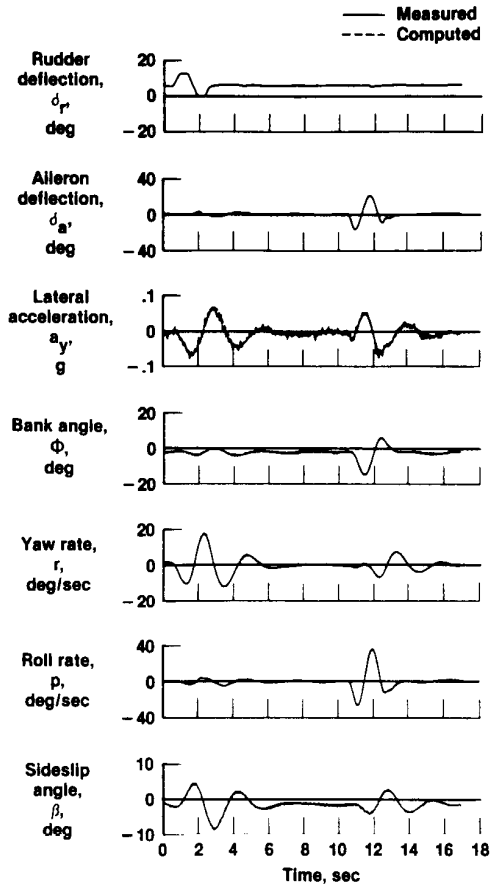

(e) Four iterations.

Figure 4.5-6. Concluded.

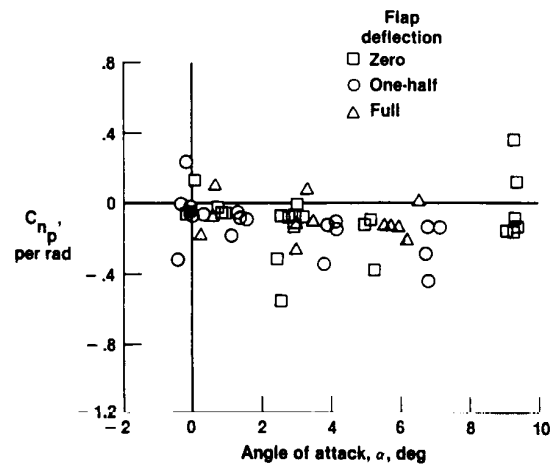

Figure 4.5-7. Variations of C_{n_p} with angle of attack.

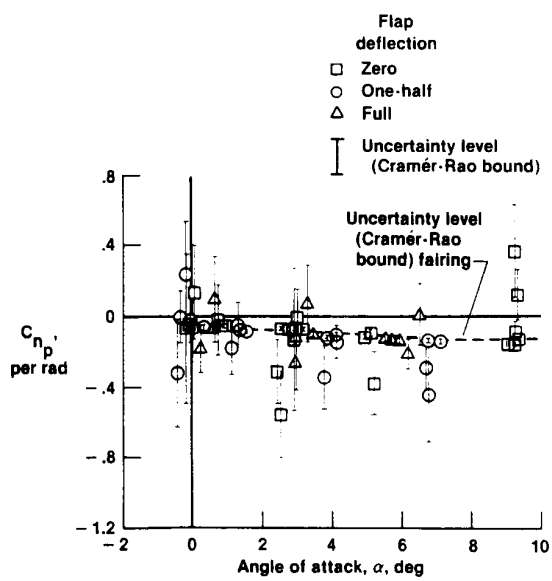

Figure 4.5-8. Variations of C_{n_p} with angle of attack, including uncertainty levels.

5.0 PREFLIGHT DATA

In this chapter we discuss the data gathered from sources other than flight tests. You normally gather the majority of these data before the flight tests, although scheduling problems, configuration changes during the flight tests, rechecking questionable data, or other reasons can necessitate some data gathering between flights or after the flight test program.

We discuss three types of preflight data: predicted aerodynamic characteristics, aircraft mass characteristics, and atmospheric conditions. There are, of course, many other things you need to know about the aircraft prior to flight, but most of these subjects are inextricably tied to discussions in other chapters. For instance, you need a thorough understanding of the instrumentation system, but gathering data about the instrumentation system is an integral part of designing an instrumentation system or evaluating the adequacy of an existing system, subjects covered in Chapters 7 and 8. Similarly, you need to know a lot about the aircraft control system to design flight maneuvers, as discussed in Chapter 6.

We give only general descriptions of the uses and sources of preflight data. Detailed discussions of computational aerodynamics or wind-tunnel techniques, for instance, are outside the scope of this document.

5.1 PREDICTED AERODYNAMIC CHARACTERISTICS

You will need a source of predicted aerodynamic characteristics for almost every flight test program. We use the phrase "predicted aerodynamic characteristics" as a catchall for any aerodynamic data other than that from the current flight tests; the source could be wind-tunnel tests, computations, independent flight tests, or combinations of these.

5.1.1 Uses

The uses of predicted aerodynamic characteristics are varied. In some cases the need for aerodynamic predictions is obvious from the statement of objectives; a common objective of flight testing is the validation of the predictions.

Predictions are also important in validating flight test results; you always want at least a rough idea of what would be reasonable ranges for the flight test results. If you are investigating new flight test methods, validation of the flight test results, partially by comparison with predictions, might be a major objective.

Many estimation methods require predicted aerodynamic characteristics as starting values for iterative algorithms. There are ways to get around these requirements, but you can generally expect better performance and fewer problems if you have a reasonable set of starting values.

There are numerous uses of predicted aerodynamics in planning flight tests. Before the first flight, of course, predicted aerodynamics are the only data available for such areas as control system design, simulation, and test maneuver design. You will often use predictions to select potential problem areas meriting extensive flight testing.

5.1.2 Theoretical Computations

Theoretical computation is one source of predicted aerodynamic characteristics. Theoretical computation comes in several forms. At one extreme are the large computational aerodynamics computer programs for solving the Navier-Stokes equations (the basic equations of fluid flow). Computational aerodynamics is an area of extensive current research activity; some computational aerodynamics problems strain or exceed the capacity of the largest computers in existence.

Some theoretical computations, in contrast, are almost trivially simple. For many purposes, you can get acceptable values with a hand calculator, a sketch of the wing, tail, and control surface planforms, and a few rule-of-thumb equations from elementary aerodynamics texts (Etkin, 1959). Maine and Iliff (1979), for instance, uses such hand calculations.

Between these extremes are many computer programs using various simplifying assumptions, approximations, and rules of thumb. Such programs are common sources of predictions of damping derivatives. Well-known programs of this type include DATCOM (Williams and Vukelich, 1979), VORTEX-LATTICE (Margason and Lamar, 1971), and PAN AIR (Derbyshire and Sidwell, 1982).

5.1.3 Wind-Tunnel Tests

Wind-tunnel tests are the major source of predicted aerodynamic characteristics for most large flight test programs. The flight data analyst often has little influence on the wind-tunnel test program (particularly on large programs) but rather must use what is available.

Static data from wind-tunnel tests are often combined with dynamic derivatives from theoretical computations to provide a complete set of predicted aerodynamic characteristics. The use of dynamic wind-tunnel testing (AGARD, 1978) is also becoming increasingly important. One technique of dynamic wind-

tunnel testing involves free-flying, remotely piloted wind-tunnel models (Hafer, 1978; Krag, 1978a,b; Subke, 1978; Hamel and Krag, 1979; Wilhelm and Verbrugge, 1982); this technique uses the same parameter estimation algorithms as full-scale flight testing.

If you use data from a large wind-tunnel test program, anticipate a significant effort in transforming the data into a form usable for your needs (you might be fortunate, but anticipate a significant effort unless you have specific reason to believe otherwise). If the wind-tunnel tests are extensive, the first large part of the effort might be finding the data pertinent to your analysis amidst volumes of other data.

You might also need to transform the data into a form suitable for your analysis programs. For instance, Figure 5.1-1 shows a typical plot of C_m, C_L, and C_D as functions of angle of attack for several elevator positions. It takes some work to obtain $C_{m_{\delta_e}}$ and C_{m_α}, quantities your analysis programs might need, from these data: The first step is finding the trim elevator position for a given angle of attack. Then C_{m_α} is the slope of the C_m curve as a function of angle of attack for that elevator position. To obtain $C_{m_{\delta_e}}$, you must difference the pitching moment values for elevator positions around the trim point. So far we have assumed that you want small perturbation values about trim conditions. If the data show nonlinearity, you might prefer to get derivatives for typical maneuver amplitudes instead of small amplitudes. If the nonlinearities are large enough to justify explicitly modeling them or if you want to analyze off-trim conditions, then it is much more difficult to decide on an appropriate model form and fit the data to it. This is essentially a system identification problem applied to the wind-tunnel data; you have a large amount of data in tablular form and want a relatively simple model that fits the data.

This example illustrates only a few of the kinds of problems you might encounter in a relatively simple case. For some vehicles with multiple surfaces and complex control systems, the first step of finding the trim condition can be difficult. Even simple data handling can be time consuming if the volume of data is large and the data are not in computer-readable form.

There are numerous manipulations of wind-tunnel data that may be required, and the details vary from program to program. It therefore avails us little to go into much detail here. The important message is that to estimate the effort required to use wind-tunnel data, you must determine exactly what form the available wind-tunnel data take, what form your analysis requires, and what transformations are necessary to obtain the needed form. It is not sufficient merely to determine that wind-tunnel data exist.

5.1.4 Independent Flight Tests

Independent flight tests are often a good source of predicted aerodynamic characteristics. If you are testing modifications to an airplane (for instance, stores testing), flight tests of the unmodified airplane usually provide a good starting point.

For rough approximations, you can use estimates from a different airplane with generally similar geometry. Even aircraft with quite dissimilar geometry often give values adequate for starting estimates in iterative algorithms (provided you are not using an MAP estimator, which gives estimates dependent on the starting values). We have, for instance, done such apparently silly things as using high-performance fighter results as starting values for general aviation aircraft, simply for reasons of convenience. If we have data files for one airplane set up exactly in the required format, the easiest way to start analysis of another airplane is to use the existing files for starting values. We pay a penalty in that convergence is often worse than it would be with better starting values; starting values from a different airplane, however, are usually much better than no starting values at all. We often use such existing files to get started in a hurry, before the wind-tunnel data for the new airplane are available in the required form.

You can also get independent flight test results by using different estimation techniques during the same flight test program. Such results cannot be completely independent, because they use the same instrumentation system, the same mass data, and possibly (depending on the methods involved) the same maneuvers. Nonetheless, using different estimation techniques during the same flight is a valuable check. The more different the techniques, the better the check. Classical static methods (Perkins and Hage, 1949; Greenberg, 1951; Klein, 1977) are quite different from output-error methods and thus could be considered to give fairly independent results. On the other hand, using an equation-error method and an output-error method on the same maneuvers does not give a particularly good check; both methods are subject to similar modeling errors and might give similar but wrong results. However, this comparison does provide a check for some errors and is better than nothing.

5.1.5 Reference Geometry

The predicted aerodynamic characteristics are invariably in nondimensional form. They are thus associated with a particular set of reference geometry values: the reference area, span, and chord. The predictions are also defined relative to a reference point. To make use of the predictions, you need to know the reference geometry values and the reference point, as well as the derivative values.

The main point of this section is that the reference area, span, and chord are just reference values. Although it is conventional, for instance, that the reference chord be the mean aerodynamic chord of the wing, there is no strict requirement that the reference values be tied directly to specific aircraft geometry. In principle, the reference chord could be defined to equal the body length (this is done on the space shuttle orbiter) or anything else that scales with the vehicle. The most important thing is to

be consistent in the use of the reference values; always use the same reference values that were used for the predicted data base (unless you are prepared to adjust the predictions to the revised references and live with the resulting confusions). For instance, if the predictions use the root chord for the reference value (a practice we have seen), resist the temptation to "correct" this error by doing your analysis with the mean aerodynamic chord for a reference; the inconsistency thereby introduced creates problems and solves nothing.

For variable-geometry aircraft, some projects establish reference geometry values that change as a function of configuration. Other projects establish a single set of reference values used with all configurations. There are arguments for both approaches, but you should simply adhere to the conventions established by the project.

5.2 AIRCRAFT MASS CHARACTERISTICS

Accurate mass distribution data are important to flight test analysis. The aircraft mass characteristics relevant to rigid-body motions are weight, center of gravity, and moments of inertia. Weight and longitudinal center-of-gravity position are easy to measure and are routinely monitored during both test and operational flights. Moments of inertia and vertical and lateral center-of-gravity positions require special tests, which are an important part of the flight test program.

In some restrictive circumstances (discussed in detail in Section 5.2.1) you can get by without much of the mass data, but such circumstances are the exception rather than the rule. Assume that you will need mass data unless it is demonstrated otherwise.

The flight test plan should provide for accurate monitoring of the mass characteristics throughout the flight test program. This monitoring can require significant time, money, and effort, particularly for large aircraft; therefore, if you neglect to plan sufficient resources for the task, you may be forced to accept either poor flight results or cost and schedule overruns. In preparing for each test flight you should verify that you have all the required mass data for the flight. You cannot reliably reconstruct mass data after the flight test program is completed, the configuration is changed, equipment in the aircraft is moved, and the records of vehicle changes and loadings are lost (if they ever existed). If you have inaccurate mass data, you will have inaccurate test results, so the mass data deserve the same care and attention to detail as the flight instrumentation system.

There are four principal uses of aircraft mass characteristics for stability and control derivative estimation, having different levels of importance and different accuracy requirements. Sections 5.2.1 to 5.2.4 discuss these uses, and Section 5.2.5 discusses sources of mass data.

5.2.1 Nondimensionalization

The most important use of aircraft mass characteristics is for nondimensionalization of the derivatives. Let us use a simple form of the q state equation to illustrate this use:

$$\dot{q} = \frac{\bar{q}sc}{I_y} C_m = \frac{\bar{q}sc}{I_y} \left(C_{m_\alpha}\alpha + C_{m_\delta}\delta + C_{m_q}\frac{qc}{2V} + C_{m_b}\right) \tag{5.2-1}$$

Defining the dimensional derivatives

$$M_\alpha = \frac{\bar{q}sc}{I_y} C_{m_\alpha} \tag{5.2-2a}$$

$$M_\delta = \frac{\bar{q}sc}{I_y} C_{m_\delta} \tag{5.2-2b}$$

$$M_q = \frac{\bar{q}sc^2}{2VI_y} C_{m_q} \tag{5.2-2c}$$

$$M_b = \frac{\bar{q}sc}{I_y} C_{m_b} \tag{5.2-2d}$$

we can write Equation (5.2-1) as

$$\dot{q} = M_\alpha\alpha + M_\delta\delta + M_q q + M_b \tag{5.2-3}$$

There are no mass data appearing explicitly in Equation (5.2-3). In similar manner, we can write all of the equations of motion in terms of dimensional derivatives (primed dimensional derivatives for lateral-directional equations, as discussed later in this section), eliminating the explicit appearance of mass data in the equations (provided we neglect gyroscopic effects). Therefore, assuming that the dimensional derivatives are constant during a maneuver, we can estimate the dimensional derivatives from the flight data alone, without considering mass data.

Given the dimensional derivative estimates, the inverse of Equation (5.2-2) gives the nondimensional estimates; for instance,

$$C_{m_\alpha} = \frac{I_y}{\bar{q}sc} M_\alpha \tag{5.2-4}$$

The nondimensional derivative estimates are directly proportional to the moments of inertia; therefore, any given percentage error in the inertia data will directly cause the same percentage errors in the nondimensional derivative estimates. By this mechanism, errors in I_x, I_y, I_z, and weight cause proportional errors in nondimensional derivative estimates.

Through a similar mechanism, errors in I_{xz} cause errors in the nondimensional derivative estimates. (Errors in other cross products of inertia would have similar effects, but I_{xz} is the only important cross product for most aircraft.) To see this effect, start with a simple form of the p and r state equations.

$$\dot{p}I_x - \dot{r}I_{xz} = \bar{q}sbC_\ell = \bar{q}sb(C_{\ell_\beta}\beta + C_{\ell_\delta}\delta + \ldots) \tag{5.2-5a}$$

$$-\dot{p}I_{xz} + \dot{r}I_z = \bar{q}sbC_n = \bar{q}sb(C_{n_\beta}\beta + C_{n_\delta}\delta + \ldots) \tag{5.2-5b}$$

The corresponding dimensional equations are

$$\dot{p} - \dot{r}\frac{I_{xz}}{I_x} = L_\beta\beta + L_\delta\delta + \ldots \tag{5.2-6a}$$

$$\dot{r} - \dot{p}\frac{I_{xz}}{I_z} = N_\beta\beta + N_\delta\delta + \ldots \tag{5.2-6b}$$

Equation (5.2-6) still has explicit mass data on the left side. The form independent of mass data is the primed derivative form,

$$\dot{p} = L'_\beta\beta + L'_\delta\delta + \ldots \tag{5.2-7a}$$

$$\dot{r} = N'_\beta\beta + N'_\delta\delta + \ldots \tag{5.2-7b}$$

where the primed derivatives are defined by the equivalence between Equations (5.2-6) and (5.2-7). To get equations for the primed derivatives, solve Equation (5.2-6) for $\dot{p}$ and $\dot{r}$ and equate matching terms. For example,

$$N'_\delta \equiv \frac{N_\delta + L_\delta I_{xz}/I_z}{1 - I_{xz}^2/(I_x I_z)} \tag{5.2-8}$$

The estimates of the primed derivatives are independent of the mass data.

Given the estimates of the primed derivatives, we can find the estimates of the nondimensional derivatives by substituting Equation (5.2-7) for $\dot{p}$ and $\dot{r}$ in Equation (5.2-5) and matching terms. In particular, matching the δ terms gives

$$C_{\ell_\delta} = \frac{I_x}{\bar{q}sb}\left(L'_\delta - \frac{I_{xz}}{I_x}N'_\delta\right) \tag{5.2-9a}$$

$$C_{n_\delta} = \frac{I_z}{\bar{q}sb}\left(N'_\delta - \frac{I_{xz}}{I_z}L'_\delta\right) \tag{5.2-9b}$$

For ailerons, the two terms subtracted in Equation (5.2-9b) may nearly cancel; typically I_{xz} is much smaller than I_z, but L'_{δ_a} is much larger than N'_{δ_a}. The result is that small errors in I_{xz} can cause large percentage errors, or even changes of sign, in the $C_{n_{\delta_a}}$ estimates. By similar reasoning, $C_{\ell_{\delta_r}}$ estimates are sensitive to errors in I_{xz}. The estimates of $C_{\ell_{\delta_a}}$ and $C_{n_{\delta_r}}$ are not measurably influenced by I_{xz}, because in these cases the I_{xz} term is negligible compared to the other term.

We have explained the mechanisms by which errors in the mass data cause errors in the estimates of the nondimensional derivatives. Any such errors in the nondimensional derivative estimates will be consistent for almost all flight conditions, maneuvers, and analysis techniques. It is seldom practical to

find mass data errors based on internal evidence in the flight data (consistent disagreements with predictions are the best clues that such errors exist). Because of this, you must be particularly careful about checking the computations of the mass data enough so that you are confident that the data are accurate. You cannot try out the data to see how they work.

Ross (1980) documents a case in which the initial analysis used erroneous inertia values. Consistent differences between the predicted moment derivatives and those estimated in flight led to suspicion of possible errors in the inertia values used. Subsequent ground-based tests verified significant errors in the moments and cross products of inertia; the roll moment of inertia was in error by about 30 percent, and I_{XZ} had the wrong sign. Ross then used equations similar to those in this section to correct the initial derivative estimates. The revised nondimensional estimates agreed significantly better with the predictions than had the initial estimates.

For some purposes, dimensional derivative estimates might be adequate, in which case you might get by with far less data on mass characteristics. In particular, if dimensional derivative estimates are sufficient, you probably do not need moment-of-inertia data (unless gyroscopic effects are important, and even then accuracy requirements are less stringent). If your sole objective is to model the aircraft motions at exactly the same configuration, flight condition, and loading as tested, then dimensional derivative estimates can satisfy this objective. We have seen test programs, analyses, and simulators based solely on dimensional derivatives. The dimensional approach is most likely to be appropriate for a program with limited objectives.

For most programs, estimation of dimensional derivatives is unacceptable for three reasons. First, all comparisons with values from any other source require nondimensional derivatives. The dimensional derivatives are useless for comparing with wind-tunnel or computed values (unless you dimensionalize the wind-tunnel or computed values, which requires knowledge of the mass data). Dimensional values early and late in the same flight are different, even at the same flight condition, because of fuel usage. It is difficult to present meaningful plots of dimensional data because there are fewer comparable points on any one plot, and of course you need many more plots, or more complexity on each plot, to describe the loading applicable to each point.

The second, closely related, reason is that dimensional derivatives are difficult to extrapolate to other conditions and loadings (unless you know the mass data, which gets back to the nondimensionalization problem). Thus, you need to test more conditions, and you have less preflight assurance that a test at a new condition will be safe than if you used nondimensional derivatives.

The third reason for using nondimensional derivatives is that test maneuver requirements for estimating dimensional derivatives are much more stringent. Dimensional derivatives are strong functions of dynamic pressure and velocity. Therefore, if dynamic pressure or velocity change much during a maneuver, then the dimensional derivatives are not constant and cannot be accurately estimated by most parameter estimation techniques. Changes of as little as 1 or 2 percent in dynamic pressure can sometimes cause noticeable problems in longitudinal cases; a 5 percent change is usually tolerable for lateral-directional cases. (You can work with somewhat larger variations than these, but you must spend a lot more time on the analysis and will get more scatter in the results.) These requirements demand carefully designed and accurately controlled maneuvers initiated from stabilized flight. The nondimensional derivatives are insensitive to dynamic pressure and velocity over large ranges; we have successfully analyzed maneuvers where dynamic pressure varied by a factor of two or more. In principle, these same comments apply to variations in weight and inertia during a maneuver, but such variations are important only in a few special circumstances (like sweeping the wings during a maneuver).

Admittedly, it is possible to separate dynamic pressure and velocity effects from weight and inertia effects, defining partially dimensionalized derivatives that would be insensitive to velocity and dynamic pressure. Equivalently, you can use nominal mass data when accurate mass data for the vehicle tested are unavailable. We have used nominal mass data as a last resort or to start analysis before accurate mass data were available. This solves only a few of the problems of dimensional derivatives. For most programs, the necessity for accurate nondimensional derivative estimates is a primary factor in the requirements for accurate mass data. We typically require the mass and moments of inertia accurate within 2 or 3 percent. Inertia errors much larger than this would become the dominant error source of the more accurate nondimensional derivatives.

5.2.2 Moment Reference

The second use of mass data is to define the moment reference point of the flight estimates. The equations of Chapters 2 and 3 are valid only in an axis system referenced to the aircraft center of gravity. All estimated moment coefficients are thus referenced to the aircraft center of gravity (the reference point is irrelevant to force coefficients).

The moment coefficients from two flight maneuvers with different centers of gravity are not directly comparable. Similarly, the moment coefficients from a flight maneuver are not directly comparable with predictions unless the reference point for the predictions coincides with the flight center of gravity. To compare moment data from different sources, you must first transform them all to a common reference point.

For most aircraft, there is little variation in the lateral and vertical center-of-gravity positions for various allowed loadings, and you can safely neglect the lateral and vertical reference translations. You will run into lateral trim problems long before lateral center-of-gravity motion causes measurable changes in the stability and control derivatives (except for the nuisance bias parameters). Gainer and Hoffman (1972) gives complete transformation equations in three axes. To apply the transformations, you first need to know the flight center-of-gravity position; if you are neglecting the vertical and lateral transformations, it is sufficient to know the longitudinal component of the flight center-of-graviy position.

For longitudinal translation of the reference point, the most important transformations are

$$C_{m_{\alpha_{ref}}} = C_{m_{\alpha flt}} + \frac{x_{ref} - x_{flt}}{c} C_{N_\alpha} \tag{5.2-10a}$$

$$C_{n_{\beta_{ref}}} = C_{n_{\beta flt}} + \frac{x_{ref} - x_{flt}}{b} C_{Y_\beta} \tag{5.2-10b}$$

where the subscript flt indicates a value referenced to the flight center of gravity and ref indicates a value referenced to any other reference point. The values x_{flt} and x_{ref} are the longitudinal positions of the flight center of gravity and the reference point, respectively, in the same units as b and c; x_{flt} and x_{ref} can be measured from any arbitrary point because they are subtracted, but they must be measured with the positive direction to the rear of the aircraft. This sign convention for center-of-gravity position is universal, so we will use it even though it is opposite to our X axis. (This difference probably arose from conventional axis systems for different disciplines; aircraft construction jigs define X axes positive rearward). The transformation of C_{m_α} and C_{n_β} is crucial to aircraft stability, and you can never neglect it. It is one of the major factors determining the safe range of flight center-of-gravity positions.

Transformations of control derivatives are occasionally important:

$$C_{m_{\delta_{ref}}} = C_{m_{\delta flt}} + \frac{x_{ref} - x_{flt}}{c} C_{N_\delta} \tag{5.2-11a}$$

$$C_{n_{\delta_{ref}}} = C_{n_{\delta flt}} + \frac{x_{ref} - x_{flt}}{b} C_{Y_\delta} \tag{5.2-11b}$$

The equations for transforming damping derivatives are more complicated, but we have never seen flight estimates of damping derivatives accurate enough to justify transforming them (unless the reference point is ridiculous, like outside the vehicle). See Gainer and Hoffman (1972) for these equations. To compare flight and predicted values, either transform the flight values to the reference point of the predictions or transform the predicted values to the flight center of gravity. Both options have small, but sometimes troublesome, problems.

The problem with transforming the predictions to the flight center of gravity is obvious: The flight center of gravity can be different for each maneuver, making presentation difficult. If the flight center-of-gravity positions are essentially coincident for most or all of the test maneuvers, this objection does not apply, and the approach works well. This is often the case for unpowered aircraft. This approach is also good for looking at results from individual maneuvers.

Presentation is simplified if you transform the flight values to the single reference point of the predictions. The problem with this approach is that it can introduce extra bias and scatter in the flight data. To see the source of the extra scatter, look at Equation (5.2-10a). The flight estimate of C_{m_α}, translated to the reference point, consists of two terms: the raw flight estimate of C_{m_α} and a translation term proportional to C_{N_α}. Therefore, the error in the translated C_{m_α} equals the error in the raw C_{m_α} plus the error in the translation term. If the reference point is far from the flight center of gravity and the flight estimates of C_{N_α} are biased or scattered, then the translation can add large bias and scatter to the flight C_{m_α} estimates. We observe this problem, for instance, on the space shuttle orbiter at hypersonic speeds, where the force derivatives have a lot more scatter than the moment derivatives. If the reference point is close to the flight center of gravity or if the flight estimates of C_{N_α} are accurate, then the translation will cause no problems.

There are two compromise approaches to avoid or alleviate these problems. One compromise is to define a reference point in the middle of the flight center-of-gravity range and translate both flight and predicted data to this reference. This minimizes the magnitude of the translation term for the flight data and thus minimizes the errors added by the translation. The cost of this improvement is the possible confusion caused by the use of a reference point different from that used for other purposes. Of course, it may be that the reference point of the predictions is already near the middle of the flight center-of-gravity range, in which case you probably have few translation problems (unless the flight

center-of-gravity range is large). This compromise helps primarily when the reference point for the predictions is something unreasonable, like the tip of the nose.

Another compromise is to use the predicted C_{N_α} or fairings of the C_{N_α} flight estimates to translate the flight estimates to the reference point. This is appropriate when the individual flight estimates of C_{N_α} have a lot of scatter. The disadvantage of this approach is that it confuses matters by mixing estimates from individual flight maneuvers with predicted data or fairings of flight data. You no longer have a simple comparison of predictions with the data from each flight maneuver.

We have concentrated on C_{m_α} because the C_{m_α} translation is usually the most important longitudinal one and thus the one where most problems occur. The same comments apply to all the translations that are big enough to be important. The lateral-directional derivative most affected by translation is C_{n_β}. We have emphasized potential problems because there is not much to say about situations in which there are no problems. You might analyze many test aircraft and never have problems with translating data to the reference point of the predictions (the most common approach). If your flight data are all at one center-of-gravity position, the reference point of the predictions is at the same position, and the flight C_{N_α} estimates are accurate, then this entire section (5.2.2) is somewhat moot. You should be aware, however, of the potential problems so you will recognize them if they do arise.

5.2.3 Sensor Reference

In addition to its role as the reference point for the aerodynamic moments, the aircraft center of gravity is also the reference point for the sensor measurements. The angle-of-attack and angle-of-sideslip vanes and the linear accelerometers sense values that are dependent on the sensor locations. Equation (2.3-2) shows the form of this dependence. The sensor positions in Equation (2.3-2) and the similar equations in Chapter 3 are all in terms of the distance from the sensor to the flight center of gravity.

Whether you model the sensor positions in the observation equations or you correct the sensor data to the center of gravity, you need the same information: position of the sensors relative to the center of gravity. This generally requires knowledge of the sensor position and the center-of-gravity position expressed in the same fixed reference frame. Because the effects of sensor position and center-of-gravity position are so related, we discuss them together here; consistent errors in all the instrument positions have the same effect as an error in the center-of-gravity position.

The sensor position relative to a fixed frame is easy to measure (unless you forget to do it until the sensor package is removed or the aircraft becomes inaccessible to you). You usually know the x component of the center of gravity accurately. The y component of the center of gravity is usually very close to zero and is, in any event, relatively easy to measure.

The z component of the center of gravity can cause serious problems but is often ignored because it is relatively difficult to measure (at least compared with the x and y components) and has only a small effect on the open-loop aircraft stability and control. If you ask what the center of gravity is, you are likely to get only the x component unless you explicitly ask for all three components (and if you ask for all three components, a common response is to question your need for the data and point out how difficult it will be to measure the vertical component).

Fortunately, the vertical component of the center of gravity usually changes little with different loadings, which somewhat simplifies keeping track of it. In many cases, it is safe to use a single value for the vertical component of the center-of-gravity position, independent of loading. You cannot, however, blindly assume that the sensors are at the same height as the center of gravity.

You should always be alert to the possibility of errors in the measured center-of-gravity or sensor positions. Such errors occur regularly enough that you should make it a rule to examine the data specifically looking for indications of them. You should check for such errors as soon as feasible after the start of the test program and after any major changes in sensor positions or the flight center of gravity. You should also check for errors after major changes in the bookkeeping that keeps track of or accounts for these positions. For instance, if you have been getting raw sensor data, and someone upstream of you in the data handling "improves" the data by correcting it to the center-of-gravity position, then you better check for several possible sources of error. Check that they are using the same values for positions, that they are using all three axes (not just x position), and that they are correcting to the flight center-of-gravity position (not the reference position). Also check that they are using the same equations for the correction. It is easy to get the signs wrong; the sign of the term for correcting the data is the opposite of that in the observation equation, and there are often inconsistent sign conventions for such things as vertical position. There seems to be about a 50 percent chance of getting the sign of any given correction term right the first time, no matter how careful you are (well, not quite that bad, but it sometimes seems like it). You need to do these checks soon enough that you can recheck the instrument positions and the center of gravity before something significant changes on the vehicle or it becomes innaccessible to you (for example, it crashes or it completes testing and is transferred elsewhere). In checking the reasonableness of quoted sensor positions, there is no substitute for going to the airplane and looking at the sensor package.

Errors in the measured center-of-gravity or sensor positions result in characteristic errors in the matches of flight and estimated time histories; the primary method of checking for errors in the measured center-of-gravity or sensor positions is to look for the associated characteristics in the time history matches.

To deduce what the time history errors should look like, look at the terms that correct for sensor position in the observation. For the major terms, which are the only ones with measurable effects, the correction term is proportional to one of the states or state derivatives. The error in the time history match should thus be roughly proportional to the same signal. The error is not exactly proportional, even in an idealized case with no noise, because the estimator adjusts the parameter estimates to take out as much of the fit error as possible; the fit is based on these estimated parameters rather than the true values, which are unknown. In spite of the distortion in the fit error introduced by the estimator, the general character of the error remains similar and is recognizable by visual inspection. We look for an error signal that is large where the correction term is large, large and of opposite sign where the correction term is large and of opposite sign, and small where the correction term is small. The sign of the error can be either the same as or the opposite of the correction term, but it must be consistent. Additive bias or long-term error trends can somewhat confuse the issue but seldom distort the characteristic shape beyond recognition.

Figure 5.2-1 illustrates the fit of flight and estimated time history data resulting from an erroneous specification of instrument position (the figure shows only the relevant signals). These data are from the first flight of the 3/8-scale, remotely piloted F-15 aircraft (Iliff et al., 1976). The data and the problem are both real; we made the mistake illustrated here in our first attempts to analyze the data. Not all cases showed the problem as clearly as this one; because this was one of the cases with the most obvious discrepancies, we chose to investigate it in detail. Choosing cases that exhibit the worst symptoms is an important early step in tracing problems.

The match between the measured and estimated roll rate is excellent. The estimated lateral acceleration generally follows the measurement but has a few areas of discrepancy. Some investigators would refer to this match as being a good one (particularly if it were plotted on a less sensitive scale, making the discrepancies less obvious); we have seen much worse fits described as excellent in the published literature. Our general experience is that even small fit errors are grounds for suspecting that something is wrong. We look for some identifying characteristics of the errors as clues to their causes.

In the case of Figure 5.2-1, we noticed (after investigating several unproductive hypotheses) some correlation between the slope of p and the fit errors in lateral acceleration. At about 1.3 sec into the maneuver, p has a large negative slope, and there is a significant negative error in lateral acceleration. The slope of p becomes positive just after 2 sec, where there is a positive lateral acceleration error. Just before 4 sec, the slope of p becomes sharply negative, as does the lateral acceleration error. The correlation is not perfect; between 3 and 4 sec there is a large positive p slope with no corresponding error, and there is some error in lateral acceleration just before 2 sec, where the p slope is small. There is enough correlation, however, to establish a strong suspicion; distortions introduced by the estimator could account for the uncorrelated errors.

Figure 5.2-1 also shows the computed $\dot{p}$. Although there was no measured $\dot{p}$ for this airplane, the computed signal is useful for this kind of debugging. The computed $\dot{p}$ plot shows clearly that a_y looks a lot like $\dot{p}$ (with changes in sign, scale, and bias). The aileron signal (not shown in the figure) also correlates strongly with a_y, but significant side force directly caused by the aileron is physically unlikely (and in any event, we were already estimating side force due to aileron as an unknown). The derivative of p, however, affects a_y through the term for sensor vertical position in Equation (2.3-2k).

Let us check the magnitude of this vertical position correction term. The peak values of $\dot{p}$ are about 1.5 rad/sec^2. Dividing by $g \cong 10$ m/sec^2 gives a peak a_y error of about 0.15 g per meter of vertical position error. Fit errors of up to 0.03 g are evident; these would require about 0.2 meter of vertical position error. This value is at least plausable enough to merit further investigation. The sensed a_y is quite sensitive to even small changes in vertical position because there are large $\dot{p}$ values in typical lateral maneuvers. This conclusion holds for most aircraft.

Motivated by this plausability check, we asked for vertical center-of-gravity and sensor position data. We found that the lateral accelerometer was 15 cm below the center of gravity. Our initial analysis had relied on a statement that the accelerometers were essentially at the same height as the center of gravity. Figure 5.2-2 shows the match that results from reanalyzing the data with the correct value for the vertical position of the lateral accelerometer. The fit of a_y is much improved. There is also a slight improvement in the p fit, but the p fit of Figure 5.2-1 is good enough that the difference is hard to see.

This 15-cm error in the assumed position of the a_y sensor had a significant effect on several of the derivative estimates. The largest effects were on the side force derivatives (which is no great surprise). Figure 5.2-3 shows the flight estimates of side force due to sideslip with both the correct and the incorrect values of accelerometer position. Figure 5.2-3(a) shows data from the initial analysis, which used the incorrect accelerometer position. The dashed line is a fairing of this data. Figure 5.2-3(b) shows data obtained using the correct accelerometer location. The solid line is a fairing of this data, and the dashed line is copied from Figure 5.2-3(a) for reference. Correcting the assumed accelerometer location increased the magnitude of the C_{Y_β} estimates by about 50 percent at low angles of attack. The estimates with the correct accelerometer position agree much better with predictions (not shown).

These consistent errors in the side force derivatives suggest a second diagnostic indicator of sensor position errors (the first indicator is the characteristic time history mismatches). If you have good

agreement with predictions of moment derivatives but see consistent bias differences between predictions and flight estimates of side force derivatives, then be suspicious of an erroneous value of the vertical position of the a_y sensor or the center of gravity. Similar differences in the normal force derivatives can indicate erroneous values of the longitudinal position of the a_n sensor.

Thus, we see that a relatively small error in the vertical position of the a_y accelerometer resulted in moderate discrepancies in the time history match and dramatic errors in the estimated side force derivatives. We learn two lessons from this result. First, it emphasizes our statement that you need accurate measurements of the center-of-gravity and sensor positions in all three axes.

Second, you must be critical of your results. Assume that there are errors in your first analysis of an aircraft (this assumption is probably more accurate than most). Take an actively negative attitude, and look for the errors (it is much less embarassing to find your own errors before other people use the results). Recall that the case shown here was one of those with the most obvious discrepancies; in other cases you might miss the discrepancy unless you actively look for it. Even if you noticed the discrepancy, an optimistic attitude might tempt you to dismiss it as a "pretty good fit;" considering that most of the fits were better, you might rationalize that this maneuver had an instrumentation problem or turbulence. (We have collected a long list of such rationalizations suitable for explaining almost any data.) Discrepancies like those in Figure 5.2-1 merit thorough investigation, and it is risky to dismiss them lightly.

5.2.4 Kinematics

We use the term "kinematics" to describe the basic dynamics of rigid-body motion. This specifically excludes characterizations of the aerodynamic forces acting on the vehicle. The kinematics largely derive from the geometry of the vehicle and the reference axes. For instance, the differential equations for the Euler angles (Equations (2.2-22g) to (2.2-22i)) derive purely from the geometry of the axis transformations and thus are kinematic equations. The kinematics also include such effects as gyroscopic and Coriolis terms, which depend on angular rates and moments of inertia but not on aerodynamics. It is these gyroscopic and Coriolis terms that interest us in this section, because they require knowledge of the moments of inertia.

For most stability and control maneuvers, the gyroscopic and similar terms are small; many of the terms are zero for symmetric airplanes. Most of the other terms are usually negligible because they are proportional to products of small angular rates and also to small cross products of inertia. The two terms most likely to be important are the I_{xz} terms on the left sides of Equations (2.2-22e) and (2.2-22f). For most aircraft, I_{xz} is the only significant cross product of inertia.

The longitudinal analysis of the first flight of the space shuttle (Iliff et al., 1981) illustrates one of the exceptional situations in which several of the normally negligible gyroscopic terms are significant. There were no intentional stability and control maneuvers on the first flight and little longitudinal maneuvering of any kind. There were several lateral bank maneuvers for trajectory control. Kinematic coupling resulted in some longitudinal motion during these bank maneuvers. Small though they were, these longitudinal motions provided the only longitudinal data available in some regimes of the entry. The shuttle had a flight instrumentation package with sufficient resolution to measure these small motions. Figure 5.2-4 shows a time history match of one of these maneuvers. Although the data from these maneuvers were of marginal quality for several reasons, they were adequate to establish rough approximations of aerodynamic derivative values for safety-of-flight purposes. The small gyroscopic effects played a central role in these maneuvers.

Gyroscopic moments from a rotating engine can be significant in some configurations, particularly those with high power-to-weight ratios. The terms proportional to I_{xe} in Equations (2.2-22e) and (2.2-22f) model this effect; for multiengine planes and multispool engines, the equations are slightly more complicated but similar in form. Inclusion of these terms requires knowledge of the moments of inertia of the rotating parts of the engines; you need the moments of inertia only about the axes of rotation. Your accuracy requirements for the engine inertia data can be fairly loose because the engine gyroscopic moment is small enough that an error of a few percent is negligible. The accuracy requirements vary depending on the magnitude of the engine gyroscopic terms, but as a general guideline, engine inertias accurate to 10 percent should suffice for all but the most exceptional cases. In many cases you need only know the order of magnitude so that you can establish that the term is negligible.

5.2.5 Sources

In this section, we briefly list the common sources of aircraft mass data (details can be found in the references).

Our most common source of mass data is ground-based measurements. The weight and the longitudinal and lateral components of the center of gravity can be accurately measured by scales under each wheel. The vertical component can be measured by tilting the aircraft while on the same scales or by suspending the aircraft. Swing tests measure the moments of inertia, using springs with known spring constants. Wolowicz and Yancey (1974) gives the details of these and other ground-based measurement techniques. We restrict the discussion here to a few general comments about using the results of the ground-based tests.

The first consideration in using ground-based tests is the necessity of adjusting the data to reflect any differences between the ground test and flight configurations. Some of the most obvious differences are the fuel and cargo loadings; the ground-based tests might be done with fuel tanks empty, full, or

both. Safety considerations often preclude tests with partially full tanks. (Partially full tanks are actually a greater fire hazard than full tanks because of the fuel/air mixture in the empty portion of the tanks.) You need to extrapolate or interpolate the data to the fuel loadings of each flight maneuver. This process varies in complexity, depending on the accuracy required and the available instrumentation. A typical scheme uses fuel quantity gauges on each tank, recorded either in the instrumentation data stream or on a pilot lap sheet. A more complex scheme might use fuel flow meters on each of several tanks, corrected by comparing integrated fuel flows with postflight measured quantities (fuel flow meters, so corrected, are usually more accurate than fuel quantity meters). This scheme can be further complicated if fuel is transferred between tanks during flight. In some cases you can get by with no fuel instrumentation at all, perhaps by starting a flight with full tanks and using a predicted fuel flow rate.

Other adjustments to the ground test data include adjustments for cargo, passengers, stores, equipment changes, and configuration. The adjustments are generally straightforward. These are mostly items with accurately known weights and locations. The greatest difficulty is simply keeping track of exactly what the flight loadings are. It is easy to forget to record something trivial like how many pilots were on a particular flight. Landing gear position is a common configuration difference between ground-based tests and flight. Some ground-based tests are practical only with the landing gear extended. The adjustment for landing gear is small enough that you can approximate it, knowing the total landing gear weight, without special tests. (The largest effect is on the vertical center of gravity.) You can also approximate the effects of minor configuration changes like clipping wing tips.

As an example of the kind of data you might use, Table 5.2-1 shows some mass data tables used for a remotely piloted oblique-wing vehicle (Maine, 1978). Table 5.2-1(a) shows the empty weight data as a function of wing skew angle. With the wing skewed, this vehicle has a nonzero I_{xy}, but I_{yz} is still zero and thus not shown. I_{xp} is the moment of inertia of the propeller, the most significant portion of the rotating mass of the propulsion system. Table 5.2-1(b) shows fuel loading data as a function of time from start of the takeoff roll (fuel consumption during warmup is assumed negligible). This table assumes a constant known fuel flow rate and uses the known positions of the two fuel tanks. The Δ values in Table 5.2-1(b) add to the empty weight data of Table 5.2-1(a); there is a small amount of unusable fuel in the tanks even when the "all fuel expended" condition occurs, which explains the nonzero Δ at that condition. All the inertia data in these tables are referenced to a nominal center of gravity of 29.3 percent of the reference chord. The vertical center of gravity is not significantly affected by skew or fuel, and the lateral center of gravity is on the vehicle centerline.

The final adjustment to ground-based measurements is translating the moments of inertia to the flight center of gravity. Our equations of motion require moments of inertia about the center of gravity for each maneuver, not about some arbitrary reference point. The translation is simple (Wolowicz and Yancey, 1974), but it is easy to accidentally omit.

We have mentioned being "careful" about measuring mass characteristics. It is difficult to explicitly state what you do to be "careful" in all regards, but the handling of reference points provides a specific and pertinent example. You should always explicitly indicate the reference point of all data that are dependent on a reference point; moments of inertia are in this class of data. Assuming that everyone understands what reference point you are using constitutes careless data handling and invites errors. If we receive inertia data without an explicit indication of its reference point, we take it as a warning of carelessness, and we feel compelled to review the processing of the inertia test data, starting from its raw form. Errors arise easily when, for example, one group assumes that everyone knows they are using the wind-tunnel reference point, a second group assumes that the data are referenced to the flight center of gravity, a third group forgets that a reference point is important, and the data are really referenced to the pivot point of the test fixture. We can broadly generalize this principle: Write out what you mean explicitly rather than assuming that others will know. The worst that can result is that you present redundant data, a small price to pay for getting the results right.

Writing out assumptions also helps insure that your own analysis is internally consistent. The hypothetical group that forgets the relevance of the reference point is in fact too large to ignore. Such negligence can result in equations with errors or internal inconsistencies; it is this possibility that makes us dubious when we receive data without a labeled reference point. Subsequent assurances that "we always use the wind-tunnel reference point" do little to assuage our doubts (particularly if the source cannot tell us exactly where the wind-tunnel reference point is).

A second good source of mass data is manufacturer's records of component buildup. If the manufacturer keeps meticulous records of the weight, position, and mass distribution of each component in the aircraft, these records can give total aircraft moment-of-inertia data that are more accurate than those from a swing test. As a flight data analyst, you will have little choice as to whether such records are kept. To be worthwhile, these records must have been kept from the start and integrated with the production change procedures. For a completed airplane, do not consider reconstructing the mass distribution component by component. On the other hand, if the manufacturer has maintained a serious weight control and monitoring program from the initial design stages, this is probably the most accurate source of mass distribution data available.

The use of manufacturer component buildup data is much like the use of ground-based test data. You need to adjust the data to the exact configuration and loading of each flight maneuver and then translate the data to the flight center of gravity. If the manufacturer is not specific about what configuration, equipment, loading and reference point the data apply to, then the data probably are not worthwhile.

You can roughly estimate inertias from simple data like aircraft weight, wing span, general configuration, and inertias of similar aircraft. Such estimates are good primarily for reasonability checks.

As a last resort, there are several ways of estimating mass characteristics from flight data, but for flight test purposes, it is better to use other sources for mass data. Methods of estimating mass characteristics from flight data span a wide range of complexity. Some of the simplest methods are the most useful in some situations. Unfortunately, many of the simple methods presume knowledge of the aircraft aerodynamics and are thus inappropriate for flight test, where their use amounts to circular reasoning. For instance, you can trivially deduce the aircraft weight from the dynamic pressure and the coefficient of lift. Unfortunately, in flight test you cannot assume that you know the coefficient of lift for a given flight condition; you are more likely to be interested in deducing the coefficient of lift from the weight and dynamic pressure. Similarly, if you know the pitching moment characteristics, you can deduce the longitudinal center-of-gravity position from the control surface positions required to trim at a given flight condition. The more knowledge you are willing to assume about the aircraft aerodynamics, the more readily you can deduce mass characteristics from flight data, but methods based on knowledge of the aerodynamics have only minimal utility in flight test.

A class of methods more appropriate to flight test requirements involve applying known forces or moments to the vehicle. Any time you apply a known force or moment, you can deduce information about mass characteristics. For instance, suppose you drop a large known weight from the aircraft in level flight (the weight having been stored in some way that does not affect the aerodynamics). One way (not the only way, and probably not even the best) to deduce the aircraft weight from this drop is to maintain the same configuration and flight condition (except for a_n) as before the drop. The aircraft weight after the drop is

$$W_a = W_d/(a_n-1) \tag{5.2-12}$$

where W_d is the weight dropped and a_n is the normal acceleration after the drop. This equation does ignore several small effects, like aerodynamic changes due to the changes in normal acceleration.

You are unlikely to use methods like Equation (5.2-12) to determine aircraft weight because you will get better accuracy with less difficulty by simply weighing the aircraft on the ground. Flight measurement of mass characteristics is a last resort, and you will use it only for quantities, like moments of inertia and vertical center-of-gravity position, difficult to measure with available ground equipment.

There are numerous methods of applying known moments and forces. We already mentioned dropping known weights. Similar principles apply to moving a known weight within the vehicle (perhaps pumping a known amount of fuel from one tank to another). You can externally apply known forces and moments with rocket engines (although you must be careful to avoid plume impingement and flow interference). Poulter (1972) uses drag chutes instrumented with strain gauges.

It is possible (but seldom advisable) to estimate center-of-gravity position simultaneously with the aerodynamic coefficients. The offsets of the instruments from the center of gravity are simply included as additional unknown parameters to be estimated.

5.3 ATMOSPHERIC DATA

Useful atmospheric data include static pressure, density, and temperature as functions of altitude. These data are needed to obtain true altitude by comparing onboard pressure measurements with the atmospheric data tables. They also provide a possible substitute for onboard static pressure and temperature measurements if the altitude is independently known (perhaps from radar). Atmospheric data are seldom critical, but they can make the analysis easier and provide redundant data for cross checking.

Atomospheric data can be obtained from several sources. One of the most common sources is balloon flights. For rough values, you can sometimes use a standard atmosphere, corrected based on ground-level conditions.

TABLE 5.2-1. – OBLIQUE-WING AIRCRAFT MASS DATA

(a) Empty weight data

	Wing skew angle, deg			
	0	15	30	45
I_X, kg-m^2	314	302	262	211
I_Y, kg-m^2	602	622	648	718
I_Z, kg-m^2	824	824	824	824
I_{XZ}, kg-m^2	69	69	69	69
I_{XY}, kg-m^2	0	-54	-94	-108
I_{Xp}, kg-m^2	0.241	0.241	0.241	0.241

(b) Fuel loading data

Time, min	Event	Total mass,[a] kg	Center of gravity, percent c	ΔI_Y and ΔI_Z,[b] kg-m^2	ΔI_{XZ},[b] kg-m^2
0	Start of takeoff roll	414	28.0	54	3
36	Forward fuel expended	400	31.5	12	3
78	All fuel expended	383	33.2	3	0

[a]Assumes a fuel flow rate of 0.39 kg/min.

[b]The values in these columns are added to the empty weight data in Table 5.2-1(a) to obtain the data adjusted for fuel loading.

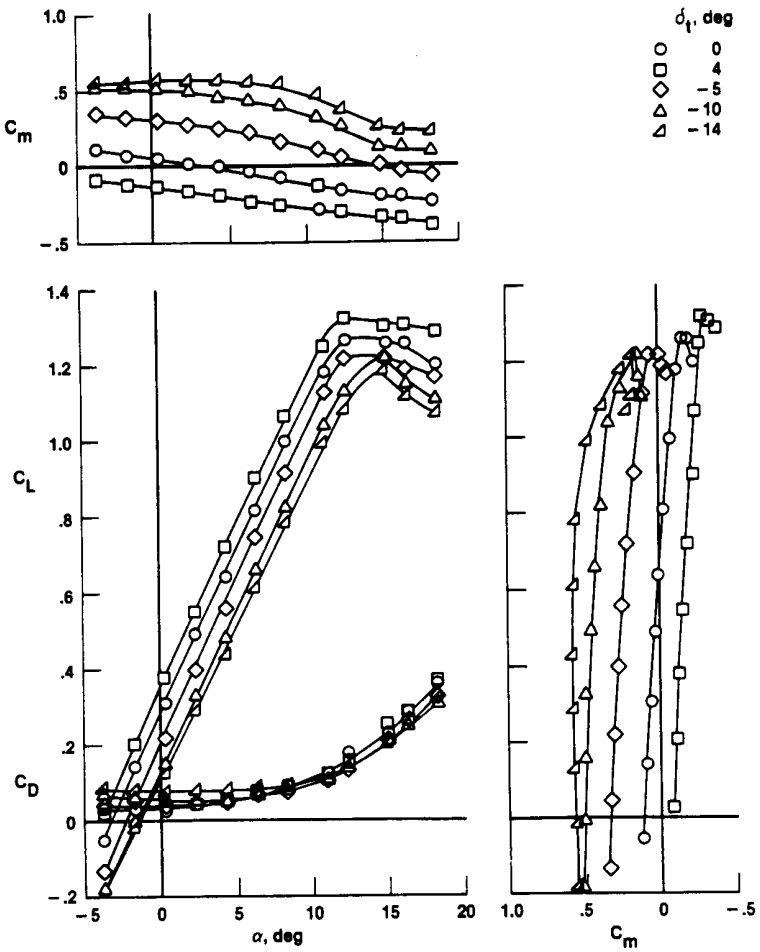

Figure 5.1-1. Longitudinal coefficient data from wind-tunnel tests (Fink and Freeman, 1969).

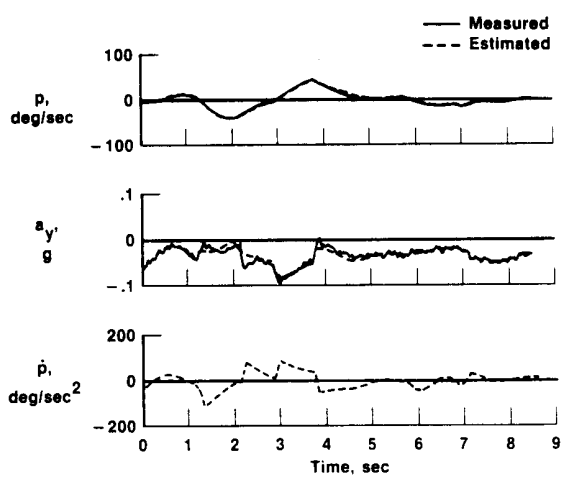

Figure 5.2-1. Time history match with incorrect sensor position.

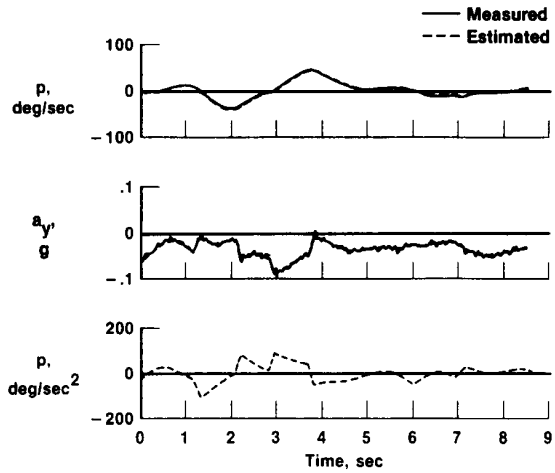

Figure 5.2-2. Time history match with correct sensor position.

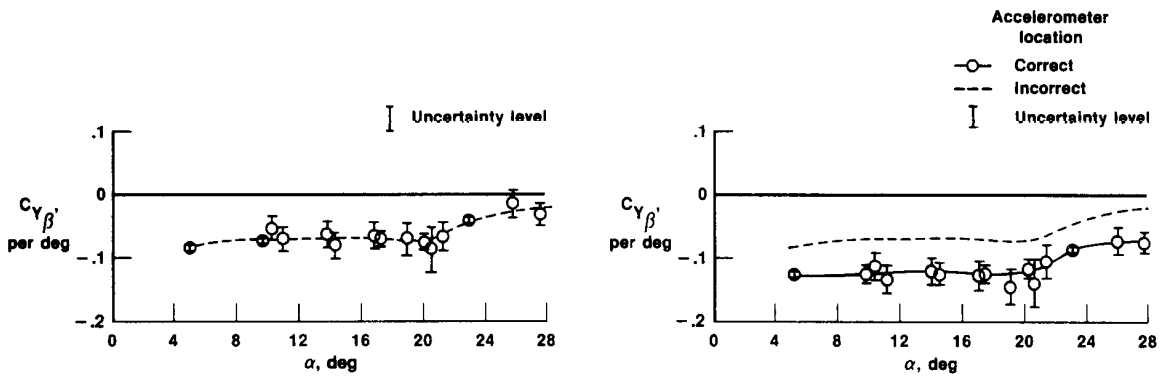

(a) Incorrect accelerometer location. *(b) Correct accelerometer location.*

Figure 5.2-3. C_{Y_β} estimates with incorrect and correct sensor positions.

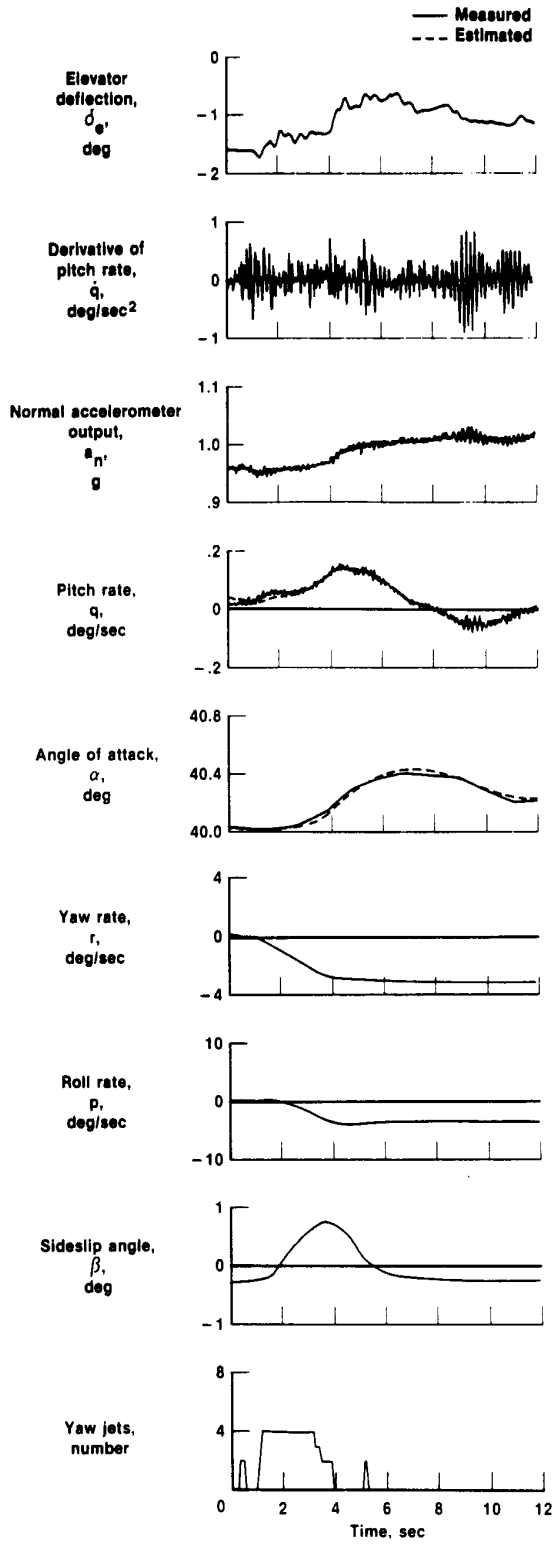

Figure 5.2-4. Fit of small longitudinal maneuver with gyroscopic effects.

6.0 FLIGHT TEST MANEUVERS

The flight test maneuvers are the part of stability and control testing that is most visible to people outside the specialty area. The entire test team and interested outsiders generally watch the test flights. During the second space shuttle entry, for instance, the stability and control maneuvers drew significant media attention (there were no intentional test maneuvers on the first flight). Mundane details such as instrumenting the vehicle and analyzing the data are secondary, mentioned only when explaining the purpose of the maneuvers. These external viewpoints distort the modern flight test process; flight test no longer consists solely of flying the airplane in various conditions to see if it breaks. We must acknowledge, though, that in spite of all the other work involved, flying the airplane is the basic element of flight testing.

This chapter discusses the kinds of flight maneuvers you use for stability and control derivative estimation. The discussion ranges from flight scheduling to design of individual maneuvers. A complete flight test program will include numerous other maneuvers for testing things such as performance, handling qualities, static and dynamic structural characteristics, system operation, and mission suitability. A major part of flight test planning is the integration of the various maneuver requirements and flight restrictions arising from different disciplines. In some cases, a single maneuver can meet the requirements of more than one discipline (Mulder et al., 1979).

6.1 THE FLIGHT ENVELOPE

The flight envelope is the set of all conditions at which the aircraft can safely fly. The flight envelope includes combinations of Mach number, altitude, velocity, dynamic pressure, angle of attack, normal acceleration, weight, center of gravity, flap position, wing sweep, engine power settings, and other parameters.

6.1.1 Envelope Coverage

No single test maneuver can possibly explore the complete flight envelope of an aircraft. The first step in planning flight test maneuvers is determining the flight envelope your tests must cover. This decision, more than anything else, determines the duration of the flight program. For some purposes, testing at a single flight condition is adequate (for instance, if your purpose is just to try a new identification algorithm rather than to learn anything new about the airplane). A few maneuvers in a single flight may be adequate for testing at a single flight condition.

For flight testing a new airplane design, you will need to test at a matrix of aerodynamic conditions covering the entire flight envelope. For a high-performance airplane with a large flight envelope, this testing usually spans hundreds of maneuvers, dozens of flights, and months of time. Some of the parameters defining the flight envelope do not directly affect aerodynamics and thus can be omitted from the test matrix. For instance, weight is always a factor in defining the safe flight envelope, but it has little to do with aerodynamics. Structural deformations related to weight affect the aerodynamics in principle, but the effects are too small to measure in most cases.

Requirements may vary in different parts of the flight envelope, depending on operational requirements and predicted problems. If small errors in the predicted aerodynamics would result in unacceptable characteristics (or if they are unacceptable as predicted) at a normal operating condition, where full aircraft capability is required, then you will need extensive data near the problem flight condition. If the predicted aerodynamics indicate docile behavior at a flight condition encountered only in an emergency, where the only requirement is to come back in one piece, then a cursory flight check for gross errors in the predictions is adequate. You should do at least a cursory flight check in all areas of the aerodynamic flight envelope, no matter how benign the predicted behavior; one of the things you are looking for in flight test is errors in the predictions. A corollary of Murphy's law states that if you omit all flight testing at a condition predicted to be benign, then the predictions will prove to be wrong at that condition.

The form of the flight envelope to be covered also influences how you will have to organize and present results. When the aerodynamics are a function of many parameters, a coherent presentation of the meters. Considerations of how to present the results can influence exactly what flight conditions you request for test maneuvers.

6.1.2 Envelope Expansion

Testing a new airplane design is significantly different from testing a previously flown configuration (we have participated in both types of test programs). A major factor in testing a new aircraft design is expansion of the flight envelope. It is imprudent to take the airplane to the predicted limit of its performance on the first flight. Instead, you restrict the first flight to the most benign conditions possible. From this point, you gradually expand the demonstrated and tested flight envelope to the limits of the airplane's capabilities. This expansion of the flight envelope intrinsically involves extrapolating data outside of the previously tested envelope to determine whether the first test flights into new regimes will be safe. To minimize the effects of possible extrapolation errors, the envelope expansion proceeds in small steps, requiring a large number of maneuvers and flights.

The envelope expansion process places special demands and constraints on the test maneuvers. You can do maneuvers only in the cleared portions of the flight envelope. Near the extremes and in newly explored parts of the envelope, you may be restricted in the magnitude and type of maneuvers allowed until there is more confidence in those areas.

There is a quandary in these restrictions in that you need good quality maneuvers to gain the confidence that you need before you are allowed to do the good quality maneuvers. The solution involves designing the best maneuvers possible within the constraints, getting the most information out of the available flight data, and then relaxing the constraints to the extent consistent with flight safety. In critical flight regimes, this confidence building can be a slow process. After five space shuttle flights through some regimes there were still constraints on the allowable maneuvers in those regimes.

Time is a major constraint in envelope expansion because results from some maneuvers are needed to clear the airplane for subsequent flights. In a fast-paced envelope expansion program, you will be extremely conscious of the tradeoff between timeliness and thoroughness. To meet schedule demands, you will be forced to put out a preliminary data analysis that is adequate for safety clearance. After the envelope expansion phase, when you have time for a more thorough analysis, you can publish refined results.

The constraints on testing a previously flown configuration or a minor modification of a previously flown configuration are much less stringent. We are often involved in such tests to obtain data for control system studies, simulators, or other purposes. After a few initial flights to test the data system, you can get hundreds of maneuvers, covering the entire flight envelope, in one or two flights. Furthermore, the test flight conditions can be spaced farther apart than in an envelope expansion program because interpolation is more reliable than extrapolation.

6.1.3 Scheduling

We leave to management texts the job of discussing critical paths, program evaluation and review technique (PERT), and such management tools for scheduling. On any project large enough to use such formal tools, the flight data analyst will have little influence on planning the schedule. This does not mean that scheduling is an issue you can ignore. Indeed, scheduling can be a crucial factor in whether you get usable data or not, and you should be aware of the potential problems so that you are better equipped to avoid them. The less influence you have on the schedule, the more attention you have to pay to ensuring that you get usable data (see Section 1.5 regarding success-oriented programs).

In the ideal flight test program (from the data analyst's point of view), you fly one or two checkout flights and then ground the airplane until you have done enough analysis to tell whether there are data system problems. Realize that every flight made with a bad data system might have to be reflown. Flights made with "small" data system problems could require several times the expected analysis effort to find and fix the problems with confidence (our definition of a small problem is one that is fixable). For these reasons, you want to minimize the number of flights with an unvalidated data system. You can achieve this either by slowing down the flight schedule or by speeding up the data system validation (it will probably be easier to influence the latter). If you are forced into a rush flight schedule, you must place high priority on rapid validation of the data system; cut those corners necessary to process and analyze some data as rapidly as possible. The analysis need not be detailed enough to publish, even as preliminary data; it just needs to identify whether there are data problems.

Even after the data system is checked out, the data from any particular maneuver can be bad because of poor maneuver execution, data dropouts, sensor failures, turbulence, aircraft characteristics significantly different than predicted, or other reasons. Therefore, you should continue to allow for repeat maneuvers. The safest policy is to fly every maneuver at least twice and critical maneuvers three or four times. Even then you should allow contingency time for other maneuvers in case analysis reveals generic problems (and the analysis must be far enough along to reveal such problems while you still have the option of flying more maneuvers). You can then afford to lose some data. If you are fortunate enough to have good data for all the maneuvers, the duplicate data will help give confidence in the results.

If you complete the entire flight program before analyzing any data, and you perform no repeat maneuvers, then you have a miniscule chance of getting good data and a somewhat larger chance of getting useless data; but the most likely outcome is that you will spend many times the anticipated time in analyzing the data, and you will get inconclusive results. You might be convinced that the results are probably good, but you will be unlikely to stake your career on them.

6.2 SMALL-PERTURBATION MANEUVERS

Section 6.1 discusses general issues relating to the number of maneuvers and the selection of flight conditions. We now begin consideration of the design of individual maneuvers. We concentrate first on the design of small-perturbation maneuvers. A small-perturbation maneuver is a maneuver to obtain data at a single flight condition. The aircraft motions are constrained to relatively small changes from the reference condition. This section discusses influences on the use and design of small-perturbation maneuvers. The issue of identifiability, which is an important influence on the maneuver design, merits special emphasis and is covered in Section 6.3.

6.2.1 Reasons for Small-Perturbation Maneuvers

There are several reasons for our emphasis on small-perturbation maneuvers. The most important is that small-perturbation maneuvers are naturally suited to locally linearized models of the aerodynamics. Large maneuvers exceed the range of validity of locally linearized models and thus necessitate the use of nonlinear aerodynamic models. Nonlinear models are much more difficult to work with than linear models, for reasons we discuss elsewhere (Maine and Iliff, 1984; and others). Therefore, except where there are compelling arguments for nonlinear models, we prefer to use linear aerodynamic models and thus small-perturbation maneuvers. Conversely, nonlinear aerodynamic models require large maneuvers to identify the nonlinearities; this point is moot if the requirement for a nonlinear model arises from inability to stabilize the airplane well enough to do a small-perturbation maneuver.

The second reason for emphasizing small-perturbation maneuvers is the difficulty of designing adequate large-amplitude maneuvers. Even if you have tools to handle the nonlinear estimation problem, the design of adequate large maneuvers is itself a formidable task. You will find sufficient challenge in the design of good small-perturbation maneuvers.

The third reason for emphasizing small-perturbation maneuvers is sometimes so important as to make all other arguments irrelevant: Safety considerations do not allow large amplitude maneuvers in new aircraft designs or untested flight regimes. No matter how convincing your arguments, you may find that small maneuvers are all that you are allowed; you must then design maneuvers to get the best data possible within the constraints of flight safety. If you are working with a proven airplane or on a high-risk unmanned project, you will have considerably more freedom in maneuver design.

6.2.2 Signal-to-Noise Ratio

The most obvious question about small-perturbation maneuvers is how small they should be. The answer to this simple question varies widely from case to case and is a compromise between conflicting pressures. In this section we examine signal-to-noise ratio, one of the conflicting factors influencing the choice of maneuver size; in Section 6.2.3 we examine other factors.

The signal-to-noise ratio is the amplitude of a signal divided by the amplitude of the noise. The most precise definitions would use root-mean-square amplitude, but we typically use eyeball estimates of peak amplitudes, which are adequate for the purposes here. Each of the aircraft sensors will have a different signal-to-noise ratio for a maneuver.

The definition and evaluation of the signal amplitude are relatively straightforward and need no elaboration. A useful definition of the noise amplitude is more oblique. For our current purposes, the noise is anything not accounted for by the deterministic part of the system model. This noise includes, among other things, effects of sensor resolution and accuracy, quantization, atmospheric disturbances, vibration (in the flight environment, not in the lab), and unmodeled nonlinearities. The relative sizes of the various effects vary greatly.

It is important to consider the frequency content of the noise. Coherent high-frequency noise, usually from structural vibration, is seldom a real problem. Such high-frequency noise is easy to filter out. Aliasing of the high-frequency noise can cause problems, as discussed in Section 7.3, but these problems are avoidable.

Figure 6.2-1 (duplicated from Section 7.3) shows some data contaminated by high-frequency structural noise. The apparent signal-to-noise ratios of q and particularly of a_n are low. The noise amplitude on a_n is larger than the signal. Figure 6.2-2 (also duplicated from Section 7.3) shows the same data after filtering. The noise amplitude is reduced by over an order of magnitude; it is barely perceptible on a_n and is imperceptible on q. The apparently poor signal-to-noise ratio of Figure 6.2-1 is misleading because the clean signals shown in Figure 6.2-2 can be obtained from the data. Section 7.3 discusses these data in more detail.

Only the noise in the same general frequency range as the system response constitutes a problem. Such noise is difficult to distinguish from system response and thus causes estimation errors. High-frequency noise is readily distinguishable from system response and does not cause significant estimation errors when properly handled. Therefore, to get a useful measure of the signal-to-noise ratio, you should consider only the noise near the system frequencies. This avoids the misleading values illustrated by Figure 6.2-1.

The signal-to-noise ratio is the primary factor favoring large maneuvers. If the noise is independent of the maneuver size and if the system is perfectly linear, then larger maneuvers result in larger signal-to-noise ratios, and thus better parameter estimates; the variance of the estimates is inversely proportional to the square of the amplitude of the maneuver, all other things being equal. Of course, neither of these assumptions is strictly true, but this idealization shows the general mechanism by which larger maneuvers give better estimates. The assumptions break down severely at large amplitudes, but signal-to-noise considerations place a definite lower bound, dependent on the noise level of the data, on the usable maneuver amplitudes. As a rough quantitative guide, a signal-to-noise ratio of 100 is about the best you can realistically expect. The most important signals should have a signal-to-noise ratio of at least 10 to get good results.

Figure 6.2-3 shows the time history of a tiny longitudinal maneuver during the first space shuttle entry (Iliff et al., 1981). To minimize the risks of this first entry, there were no intentional stability and control maneuvers (the extreme case of a limitation imposed by safety considerations). The airplane was flown as smoothly as possible throughout the entry; transients the size of those in Figure 6.2-3 are unavoidable with the existing shuttle control system. This maneuver would be too small to analyze with typical aircraft instrumentation; the 0.1-deg/sec peak pitch rate of the maneuver is only a few times the typical pulse code modulation (PCM) resolution for a pitch rate gyro channel. However, the shuttle has a high-resolution data package, the aerodynamic coefficient identification package (ACIP), specifically designed for stability and control analysis.

Even with this high-resolution data system, the maneuver shown in Figure 6.2-3 is about the lower limit of what we can reasonably analyze on the shuttle. The limitation arises from the instrument noise and vibration level (as the package is currently installed), which is significantly larger than the resolution. The signal-to-noise ratio of the pitch rate and angle of attack is about 10; that of a_n is difficult to estimate but is probably about 1. The signal size in angle of attack is the 0.5° motion of the signal during the maneuver, not the 40° total value. The dynamic signal in a_n is difficult to see; the

steady increase during the maneuver is due mostly to increasing dynamic pressure and does not constitute a response to the dynamics of the maneuver. We got reasonable (not excellent, but reasonable) pitch coefficient estimates from maneuvers of this amplitude but marginal to poor normal force coefficient estimates; such qualities are about what we expect from examining the signal-to-noise ratio of Figure 6.2-3.

6.2.3 Other Sizing Considerations

Small-amplitude nonlinearities can define lower bounds on acceptable maneuver size. We normally think more of large-amplitude nonlinearities, but there are some dead-band and hysteresis effects that are apparent only for small motions. For instance, a control surface might have little effectiveness until it deflects far enough to get out of the wing boundary layer or a small area of separated flow. There is speculation, as yet unverified, that a similar small-amplitude nonlinearity could explain the space shuttle's small oscillations between Mach 1 and Mach 2 and the apparently contradictory well-damped response to larger inputs in the same regime.

Several factors favor small maneuvers. The first such factor is the requirement of linearity. As the maneuver magnitude increases, the locally linearized model of the aerodynamics loses its validity. The estimates gradually become poorer representations of the airplane, eventually passing an ill-defined border between acceptability and unacceptability. For extremely large maneuvers, there is little we can do except abandon either the maneuver or the linear model. For maneuvers near the border of acceptability, there are several approaches to improving the quality of the estimates.

A closely related factor favoring small maneuvers is the necessity to maintain a nearly constant flight condition. The locally linearized model applies only at a specific flight condition. If the maneuver is large enough that the flight condition changes significantly, the linear model may become inadequate. For example, you must control Mach number precisely for maneuvers at transonic conditions; the aerodynamic coefficients at Mach 0.95 could be radically different from those at Mach 0.98. The linear model is more flexible about changes in some flight condition parameters than in others. We can tolerate large changes in dynamic pressure if we linearize using point-by-point measured values of dynamic pressure (see Section 3.7); this adds significantly to the computer costs but not to the complexity of the data analysis.

A maneuver can be quite large in some senses, while still meeting the requirements of a small-perturbation maneuver. Figure 6.2-4 shows a maneuver designed for estimating $Cm_{\dot{\alpha}}$. Maine and Iliff (1979) discuss the details of the design of this maneuver. For now, let us note that the maneuver includes a 360° aileron roll, with large changes in altitude, attitude, acceleration, and dynamic pressure. However, the angle of attack, elevator position, and sideslip angle (not shown) stay within relatively small limits. This is a small-perturbation maneuver, and it was successfully analyzed with a linear model.

The final factor limiting the size of the maneuvers is flight safety. The maneuver shown in Figure 6.2-4 would be unacceptable in an envelope expansion program or on a large transport plane.

These considerations place lower and upper bounds on the acceptable maneuver amplitude. For most aircraft, the range between the lower and upper bounds will be large, the best maneuver amplitudes being those near the middle of the acceptable range.

This discussion has been more descriptive than quantitative. Table 6.2-1 gives a few specific numbers for reasonable maneuver amplitudes. We must emphasize, however, the severe limitations of these rough numbers. First the numbers are oriented towards high-performance airplanes such as fighters; this is the class of airplane with which we have the most experience (though we do have experience on a wide variety of aircraft). For a large transport, you will generally use about the same displacement angles, but smaller rates than these numbers indicate. Second, even for a given airplane, the maneuver amplitudes will vary with flight condition. Finally, the magnitudes shown in the table for different signals will often be inconsistent; a 2° angle-of-attack change might imply a 1-g change in normal acceleration for a high-speed condition.

6.2.4 Design Constraints

There are several practical constraints on the design of stability and control test maneuvers. This section briefly lists the kinds of constraints we have most commonly encountered. Some of the constraints are precisely defined; such constraints include control surface position limits and actuator rate limits. The hardware is incapable of exceeding these limits.

Other constraints relate to the control system. It might be impossible to independently define the position of each control surface without redesigning the control system. For instance, a single lateral stick command often results in a blend of inboard aileron, outboard aileron, spoiler, and rudder motions (or whatever combinations of surfaces the particular airplane has); in this case a maneuver with only outboard aileron motion is not possbile. The control system often has feedback loops that are impossible or unsafe to turn off. Such feedback loops present special problems, discussed in Section 6.3.4 You might also encounter filters, electronic or mechanical, in the control system that limit your ability to make sharp inputs.

Combinations of the aerodynamics and the control system can prevent stabilization at extreme flight conditions. This can happen either because there is insufficient control power to trim at the desired condition or because the airplane, though having enough raw power, is too poorly behaved to stabilize at the condition. Extreme off-trim conditions may be impossible to reach safely, even in a transient maneuver.

Some constraints are vaguely defined. Many maneuver magnitude constraints, whether due to safety or linearity considerations, have some flexibility. A maneuver with a brief transient beyond the stated boundaries might be acceptable, while another maneuver that closely approaches several safety boundaries could be judged as too risky for the expected return. In general, there are degrees of acceptability rather than simple definite boundaries.

The requirement to maintain a constant flight condition falls in the flexible class. Section 6.2.3 discusses the distinction between flight condition parameters that can vary widely and those that must be precisely controlled. The exact amount of acceptable deviation is seldom defined. This requirement is often a major factor in maneuver design. Figure 6.2-5 shows a maneuver dominated by the requirements of maintaining Mach number constant at 0.9 and angle of attack constant at 12° during a lateral-directional maneuver. The airplane had insufficient thrust to maintain these conditions in level flight.

In more conventional flight regimes (level cruise flight) you often choose the order and sign of inputs to minimize changes in the flight condition. If the first part of a lateral-directional maneuver starts the aircraft rolling to the right, you might choose the sign of the latter part of the maneuver to counter the right roll. Another approach is to delay until the end of a maneuver those inputs likely to result in the largest upsets of the flight condition. This allows you to choose the time segment actually used in the analysis so as to exclude the large flight condition changes.

One good approach to minimize the flight condition changes during a maneuver is to start from a non-steady condition. For instance, if the maneuver will cause the aircraft to accelerate, start from a decelerating condition. Figure 6.2-4 shows an example of applying this idea, which in this case reflects pilot common knowledge: It is best to pull the nose up before starting a roll. The dynamic pressure change in this maneuver would have been much larger were the maneuver initiated from level flight.

6.2.5 Pilot Involvement

Unless you have the capability of programming automatic inputs, the maneuvers must be such that a pilot can fly them with the required precision. This constraint is best addressed by discussing the proposed maneuvers with the pilot early in the planning stage. Kleingeld (1974) demonstrates precise execution of a demanding maneuver; extensive simulator training aided this task. The maneuver shown in Figure 6.2-5 illustrates extremely precise control in spite of demanding physical conditions (over 4 g) and unusual attitudes.

A knowledgeable test pilot can often suggest alternative strategies for obtaining the data you need. The suggestions sometimes can be as simple as restating the maneuver description in different terms. For instance, an impractical maneuver requirement of starting an aileron pulse 1.8 ± 0.1 sec after a rudder pulse takes on a more reasonable appearance when restated to require the aileron pulse at the time of peak positive sideslip angle. Besides being easier on the pilot, the latter statement is probably a more robust statement of your requirements, considering possible differences between your predictions and the actual aircraft response. You may find that the pilot has an excellent instinctive feel for things that sound complicated, like what kinds of inputs will excite the Dutch-roll mode (a subject likely to be of much interest to you; see Section 6.3.2).

The maneuvers shown in Figures 6.2-4 and 6.2-5 were designed in close consultation with the test pilot, who had a good understanding of the engineering objectives and made major contributions towards designing a maneuver to meet the objectives. If you have good test pilots, take advantage of their skills to help you design good maneuvers, rather than giving them yes/no choices. Flying test maneuvers is, after all, their area of expertise. Even if you have automatically programmed inputs, the pilots are likely to be helpful in designing the inputs.

A side benefit of actively involving the pilots in maneuver design is that they are then better equipped to adjust to unexpected situations. You might have more stringent safety limitations than expected because of in-flight failures or aircraft characteristics different than expected; a pilot familiar with your objectives might be able to modify the maneuver to meet the new constraints instead of abandoning the maneuver as impossible. The necessary modifications are sometimes so simple as to be obvious if the pilot has a simple understanding of the objectives. Suppose, for instance, you requested that a maneuver be initiated at 12° angle of attack, but the airplane is unexpectedly limited to a maximum of 10° angle of attack. If the pilot has no more information than that, you will probably lose a maneuver. If the pilot understands that your objectives just require the highest practical angle of attack, estimated as 12°, the pilot might get you usable data by starting the maneuver at a 9° or 10° angle of attack.

You can encounter unexpected opportunities, as well as unexpected limitations. Even in the most tightly packed flight programs, unforeseen limitations encountered in flight can force cancellation of many of the activities planned for a flight, leaving idle time. If the limitations do not affect the stability and control maneuvers (for instance, if only the flutter instrumentation fails on a flight planned to gather flutter data), you have a good opportunity to gather extra data. Much of the cost of a test flight is incurred getting the airplane in the air; the savings from ending a flight early are usually negligible compared to the cost of the flight. Therefore, you want to get as much data as possible out of each flight as long as the airplane is safe to fly. A pilot with a good understanding of your requirements can often improvise useful maneuvers on the spot instead of landing early for lack of better things to do. Such improvised maneuvers need not be ideal, just useful.

6.3 IDENTIFIABILITY

Identifiability is the central issue in the design of maneuvers for parameter estimation. The issues discussed in Section 6.2 place important constraints on the maneuver design, but they serve a secondary role as modifiers to the primary goal of identifiability. The objective of the stability and control test maneuver is to provide good identifiability within the constraints imposed.

The unknown parameters of a system are said to be identifiable if it is possible to identify their values based on measurements of the system inputs and outputs. Although it is common to talk loosely of the system being identifiable, the definition depends inherently on the parameterization of the system. Almost any system can be made unidentifiable by allowing enough of its parameters to be unknown; conversely, a system can often be made identifiable by specifying known values for enough of its parameters. Therefore, it is more reasonable to talk about the parameters of the system being identifiable than about the system itself being identifiable.

As we use the term, identifiability is defined for a specific input (thus the relevance to this chapter on maneuver design). It is trivial to demonstrate systems that have unidentifiable parameters for all possible inputs; the most trivial examples are systems with no outputs. Therefore, input design is obviously not the only factor in identifiability. The system structure and parameterization are the primary factors, and they define the limits of what is achievable by maneuver design. In this chapter we focus primarily on maneuver design. Naturally, the subjects of maneuver design and identifiability are closely related; you cannot definitively settle one without considering the other.

We leave formal discussion of identifiability to other sources. Maine and Iliff (1984, chapter 11) discusses several aspects of identifiability. Plaetschke and Schulz (1979) compares several proposed methods of automatically designing maneuvers to "optimize" identifiability. You can design maneuvers with good identifiability by following a few general principles (discussed in Sections 6.3.1 to 6.3.4). Our emphasis is on criteria that are easy to evaluate in practical situations; in some cases the answers are obvious by inspection, with the appropriate insight. Most practically encountered identifiability problems are simple, often trivially simple. Avoid burying your head so deeply in a morass of pedagogical detail that you miss what is obvious to the unsophisticated.

6.3.1 Independent Inputs

You must have independent inputs on every control surface for which you intend to estimate derivatives. Part of this requirement is obvious at first glance: You cannot estimate rudder derivatives if the rudder does not move. The requirement that the inputs be independent, although simple to explain, is a major source of confusion and error. If the motion of one surface is directly proportional to the motion of another, then there is no way to distinguish the effects of the two individual surfaces. Aileron-to-rudder interconnects give a classic example of this problem. Suppose the pilot moves the lateral stick in an airplane with an aileron-to-rudder interconnect. Not only does the rudder motion not help the data, it actually destroys the aileron effectiveness information. You can estimate neither aileron nor rudder derivatives from such a maneuver, even though both surfaces move, because the motions of the two surfaces are not independent. The best you can do is estimate an equivalent combined effectiveness, ignoring the question of how much of the effect comes from each surface.

There are many circumstances in which surfaces move together, as in the aileron-to-rudder interconnect example. In some cases, the combined equivalent derivative is all you need. In large airplanes you sometimes find that a single control surface is physically split into two or more segments, each driven by a separate actuator. In normal operation you do not care how much of the effectiveness comes from each segment, because all the segments move together within the accuracy of the actuators. (In any event, superpositioning individual segment effects may not work well, because of interference effects of the adjacent segments.) Therefore, except in failure-mode testing, you do not need maneuvers with independent inputs on each segment.

Multiple surfaces moving together for roll control are common. These surfaces can include inboard ailerons near the wing root, outboard ailerons near the wing tip, differential horizontal stabilizers, asymmetric spoilers, and others. In this case, the ratios of surface motions often depend on flight condition; the outboard ailerons, for instance, are often held fixed during high-speed cruise flight. The ratios can also depend on amplitude; the spoilers sometimes have a dead band so that they deflect only for large stick motions. Furthermore, it may be desirable to examine the effects of control system changes that involve different ratios of surface motions. You usually want independent inputs on each of these control surfaces; otherwise you must estimate combined equivalent derivatives, which have limited utility.

It is acceptable for surfaces to move together during part of a maneuver as long as there is independent motion elsewhere in the maneuver. The independence requirement applies to the maneuver as a whole, not to portions of the maneuver. Figure 6.3-1, for instance, illustrates a reasonable aileron and rudder input. The second pulse provides rudder information, and the first pulse provides combined rudder-aileron information. Using both pulses, we can reasonably expect to estimate both rudder and aileron derivatives. Thus, as long as the pilot does both a lateral stick pulse and a rudder pedal pulse in each maneuver (a good idea anyway, as discussed in Section 6.3.2), the aileron-to-rudder interconnect does not degrade the results.

You cannot get acceptable independence from actuator nonlinearities, noise, small lags, dead bands, or other minor differences between the motions of different surfaces. A likely effect of such minor differences is that instead of failing with a divide-by-zero error, the minimization algorithm will diverge or give ridiculous answers (often wrong by several orders of magnitude). Useful independent inputs must have large independent components.

Indeed, the worst inputs are those with fairly small differences. If the inputs have large independent components, then you can estimate independent derivatives. If two inputs are dependent or so nearly dependent that you can neglect the difference, then you can at least estimate combined equivalent derivatives, which have some utility, albeit limited. It is usually safe to neglect differences due to such small problems as noise, actuator errors, and quantization. Your worst problems occur when two inputs are different enough that you cannot neglect the difference, yet not different enough to allow good independent estimates.

Nonlinear control laws are a common source of such troublesome differences. On the space shuttle, for instance, the reaction control jets are inherently nonlinear, each jet being either on or off. There is a blend of rudder and yaw-jet control in some flight regimes. Although the rudder motions and jet pulses have characteristically different shapes, they tend to operate in the same direction at the same time. This type of data is extremely difficult to analyze. Neither independent derivatives nor combined equivalent derivatives work well. In fact, there is no approach that reliably works well, except to get different maneuvers with greater independence. The results are sensitive to small errors in modeling, measurements, and predictions. You can get wide ranges of results from minor changes in maneuver times, weightings, and other details. Therefore, the results are likely to reflect more the analyst's personal judgments than objective deductions from the data. With data like these, you are likely to spend many months getting questionable results; you might be well advised to abandon the effort at the start and declare the data unusable.

6.3.2 Modal Excitation

Your inputs must significantly excite all the modes of the model you are analyzing and should minimize excitation of unmodeled modes. This is probably the single most important principle of input design for parameter estimation; it applies to any system (aerospace, economic, biological, or other). If you can devise an input that excites the appropriate modes, you are over the biggest obstacle to designing a good input; the rest of the design process consists mostly of modifying the input to meet constraints.

This principle is more a restatement of the problem than a solution. It does not indicate how to design an input to excite the appropriate modes. The restatement does present the problem in terms more understandable to people familiar with the dynamics of the system. The test pilot, for instance, can probably suggest how to excite the longitudinal short-period mode without exciting the phugoid.

The design of inputs for good modal excitation separates into two issues: selecting appropriate control surfaces or other control devices to use and specifying the shape of the input signals as a function of time.

Controllability is the theoretical basis for the selection of which control surface to use. To get excitation of a mode, you must select control surfaces that make the mode controllable. A classic example of a controllability problem is a structural shaker positioned at a nodal point. Such a shaker cannot excite the mode that has a node at the shaker.

For aircraft stability and control analysis, the surface selection problem is relatively straight forward. You are seldom in a position to specify the design and placement of control surfaces to be used in testing. The decision reduces to which of the existing control surfaces to use. The selection options are further reduced by the considerations of Section 6.3.1, often to the point of leaving no decision to make.

In many cases you will have a choice between designing several maneuvers, each using one of the control surfaces, or designing a single maneuver using all the control surfaces. If more than two surfaces are involved, compromise choices are also possible. Although other constraints can affect this decision, adequate modal excitation is usually easier to ensure when you use several different surfaces in the same maneuver. (However be sure to make the inputs independent, as discussed in Section 6.3.1.) Section 6.4.2 gives a specific example of the advantages of maneuvers that use multiple control surfaces.

We can give little universally useful advice about how to specify the shape of the input signals for achieving good modal excitation. For the specific case of the aircraft stability and control models that we normally use, Section 6.4 gives examples of inputs that we have used with good results. For more general situations and where there is less previous experience to draw from, the frequency content ideas of Section 6.3.3 are helpful.

6.3.3 Frequency Content

The system modes are best excited by frequencies near the system natural frequencies. Input frequencies much higher than the system natural frequencies result in negligible response (or excite higher frequency unmodeled modes). Very low input frequencies result in essentially static data. Therefore, good inputs should have most of their energy in frequencies near the system natural frequencies.

Taking this principle to the extreme suggests inputs consisting of sine waves at the system natural frequencies. Such narrow-band inputs, however, cause several problems. Such design is overly sensitive to errors in the predictions of the system natural frequencies. More serious is the fact that narrow-band inputs tend to degrade identifiability. You cannot accurately characterize a system based on data generated by input at a single frequency.

An ideal input would have a fairly broad spectrum, covering the range of system natural frequencies and decreasing outside of that range. A frequency sweep is an obvious signal form meeting these require-

ments. Frequency sweep inputs are often used in structural testing. Figure 6.3-2 shows a frequency sweep maneuver on a 3/8-scale F-15 aircraft model (Iliff et al., 1976). We obtained good results from this and similar maneuvers.

Koehler and Wilhelm (1977) presents an input design, which they call a 3211 input, based on these frequency content ideas. Their design is relatively easy for a pilot to fly, compared with something like a frequency sweep, which is best performed with an automatic maneuver system. Figure 6.3-3 shows a 3211 input from Plaetschke and Schulz (1979). The input is a series of four contiguous steps with alternating signs, lasting for 3, 2, 1, and 1 time units. The length of the time unit is adjustable to center the frequency band of the input around the system natural frequencies. We have tried 3211 inputs on several aircraft and have obtained good results. Plaetschke et al. (1983) compares flight data results from 3211 inputs with those from several different inputs that were designed based on optimality criteria.

6.3.4 Feedback Systems

Automatic feedback systems are often a source of identifiability problems. We discuss two kinds of identifiability problems caused by feedback, along with several approaches to the solutions of these problems.

The easiest solution to feedback problems is to turn the feedback system off during the stability and control maneuvers. You need only turn off the feedbacks in the axes of the maneuver. Lateral-directional feedbacks do not degrade longitudinal maneuvers. The lateral-directional feedbacks often improve longitudinal maneuvers by helping to stabilize unwanted lateral-directional oscillations; conversely, longitudinal feedbacks often improve lateral-directional maneuvers.

If turning the system off is impractical (for instance, if the open-loop airplane is unstable or if there is no mechanism for turning the feedbacks off), then reducing the feedback gains can alleviate some of the problems. If gain reduction is impractical or inadequate, you need to look more specifically at the nature of the problems.

The first problem is that feedback systems often make it difficult to excite the dynamics of the system. With some high-gain feedback systems, the aircraft motion decays without oscillation almost immediately after the input stops. Such deadbeat response is the aim of the feedback system, but it gives little data on which to base parameter estimates. There is high control surface activity, but it is of small amplitude, highly correlated, and concentrated at high frequencies. Any feedback system will result in closed-loop response characteristics different from the aircraft open-loop characteristics. This difference is not necessarily a problem in itself; it becomes a problem only when the closed-loop response is deadbeat.

If you cannot turn the feedback system off or reduce the gains, the next best solution to the problem of deadbeat response is to use persistently exciting inputs. You should also pay particular attention to getting a relatively broad frequency spectrum in the input. The usual pulse-type inputs often give insufficient response for analysis. We suggest frequency-sweep inputs as a good choice for this situation if you have the capability to do them. A series of 3211 inputs might also work well (Marchand, 1977). In any event, the input should persist for almost the entire planned duration of the maneuver because the aircraft response will stop soon after the input stops.

The second identifiability problem caused by feedback systems is linear dependence (Koehler and Wilhelm, 1979). Consider a simple example in which yaw rate is fed back to the rudder to augment the yaw damping. The rudder position is then directly proportional to the yaw rate during aileron maneuvers. In this event, there is no way to distinguish rudder effects from yaw rate effects. This is similar to the dependence problems mentioned in Section 6.3.1, except that the dependence discussed there is between two control surfaces. As the feedback system becomes more complicated, the problem becomes more complicated, but it retains the same basic character. If a control surface position is a direct function of several vehicle state variables, then you will be unable to distinguish the control surface effects from combinations of the state variable effects. As discussed in Section 6.3.1, lags and other small differences usually make such dependence problems more difficult to resolve instead of easier.

The solution to dependence problems caused by feedback is simple in principle and is the same as that suggested in Section 6.3.1: Add an independent input signal to each control surface that moves during a maneuver. It does not then matter whether the control surface position also has a component due to feedback. If it is impractical to add an independent input signal (for instance, if the control surface is dedicated solely to feedback functions and this is not easy to change), then you will probably be forced to assume that some of the derivatives are known from other sources or to use some form of equivalent or combined derivatives.

There is a common misconception that any analysis of maneuvers with a feedback system on will give equivalent closed-loop derivatives. This misconception is based on earlier hand-computation methods of analysis (which are still useful for some purposes). There are no such equivalent derivatives in our model; we normally get estimates of the bare-airframe open-loop derivatives whether the feedback system is on or off (Koehler and Marchand, 1979). We are forced to use a model with equivalent derivatives only when we cannot get adequate maneuvers by the methods discussed previously or when control surface measurements are missing.

The reason that our analysis gives open-loop derivative estimates is that actual control surface positions are inputs to the model. The effect of feedback is to change the input to the model, not to change the model's inner structure. If we used pilot stick position as the input, then the feedback loops would

be internal to the model, and we would anticipate estimating equivalent closed-loop derivatives. This is one of the reasons (Section 8.2 discusses others) that we use control surface position measurements in preference to measurements of pilot stick position.

6.4 SAMPLE MANEUVERS

In this section we show sample small-perturbation maneuvers designed for aircraft stability and control analysis. We give several empirical rules used to design these maneuvers. We do not claim that the maneuvers or rules of this section are optimum, but they have given adequate results in a wide variety of circumstances.

The maneuvers are based primarily on adaptations of standard piloting tasks and flight test maneuvers previously used for other purposes, rather than on mathematical derivations of good input signal shapes. The design of the maneuvers is largely a refinement task; given that some simple standard maneuver comes reasonably close to our requirements, we can place restrictions on the maneuver or slightly modify it to better achieve our purposes. This design philosophy is the reason for the predominance of pulse-type inputs.

6.4.1 Longitudinal Maneuvers

Longitudinal inputs are usually easy to design. Our usual longitudinal models include only the short-period longitudinal mode. This mode can be excited by a simple pulse or doublet on the elevator or other longitudinal control surface. The exact shape of the pulse is unimportant. For best excitation, pulse duration should approximately equal natural period of the short-period mode. The constraints mentioned in Section 6.2 determine the best pulse size and whether singlets or doublets are best. Large longitudinal pulses can result in significant flight condition changes.

Figures 6.4-1 to 6.4-6 show typical longitudinal test maneuvers from several aircraft: a JetStar (Brenner et al., 1978), an F-111A (Brenner et al., 1978), a Beech-99 (Tanner and Montgomery, 1979), a T-37 (Maine, 1981b), a 3/8-scale F-15 model (Iliff and Maine, 1975), and a Piper PA-30 (Brenner et al., 1978). Other longitudinal maneuvers are illustrated throughout this document.

6.4.2 Lateral-Directional Maneuvers

Our usual lateral-directional models include the Dutch roll, roll, and spiral modes. The best way to ensure adequate excitation of all modes is to have both aileron (or other primary roll control) and rudder input during the maneuver. It is certainly possible to excite all the modes with aileron input only or with rudder input only, but these approaches require careful input design and execution, varying from aircraft to aircraft. It is our general experience that a simple rudder doublet, combined with an aileron doublet in the same maneuver, is easy to execute and gives consistently good excitation.

If you must use data with the aileron and rudder inputs in separate maneuvers, the best way to avoid problems is to pair aileron and rudder maneuvers from similar flight conditions (assuming that you have such matching maneuvers). You can analyze the set of two (or more) maneuvers together, obtaining a single set of estimates that applies to both maneuvers. We call this approach double-maneuver (or multiple-maneuver) analysis.

Figures 6.4-7 and 6.4-8 show separate rudder and aileron maneuvers from a Piper PA-30 aircraft. Figure 6.4-9 shows estimates of C_{n_p} from several such maneuvers on the PA-30. The data exhibit large scatter, with both positive- and negative-valued estimates, although there is significant clumping of points with small Cramér-Rao bounds near the center of the scatter band.

Figure 6.4-10 shows the same data segregated into rudder and aileron maneuvers. The aileron maneuvers are obviously superior for estimating C_{n_p}; they give less scatter and smaller Cramér-Rao bounds. Foster (1977) documents a similar conclusion. The advantage of using the aileron maneuver for estimating p derivatives is fairly clear from looking at the maneuver time histories shown in Figures 6.4-7 and 6.4-8. The aileron maneuver gives large-amplitude roll rates with shapes characteristically different from the other signals; the roll rate from the rudder maneuver is smaller and is similar in shape to β. The rudder maneuvers, however, are better for estimating some other derivatives; the rudder input gives far better excitation of the Dutch roll mode.

Figure 6.4-11 shows the C_{n_p} estimates from multiple-maneuver analysis of the same data. Every case analyzed included one or more rudder maneuvers and one or more aileron maneuvers, paired by flight condition. This analysis combines the advantages of both the rudder and the aileron maneuvers, giving good estimates of all the derivatives.

Although you can use multiple-maneuver analysis in this way to pair aileron and rudder maneuvers, you will have fewer problems if you include both aileron and rudder inputs in each maneuver flown. This avoids the problems of pairing maneuvers, small flight condition mismatches, and additional initial conditions (you must consider the initial condition of each part of a multiple maneuver).

The order and timing of the aileron and rudder pulses are largely influenced by the requirement to minimize flight condition changes. There have been several studies on varying the number, order, and timing of the pulses (Cannaday and Suit, 1977). For pulse-type maneuvers, most of these studies agree that there is little noticeable effect on the identifiability of the derivatives as long as there are both aileron and rudder inputs.

Figures 6.4-12 to 6.4-15 show some typical combined rudder-aileron inputs from several aircraft: an F-8C (Steers and Iliff, 1975), an oblique-wing airplane (Maine and Iliff, 1981b), an HL-10 lifting body (Brenner et al., 1978), and an F-111A (Iliff and Maine, 1975). Other lateral-directional maneuvers are illustrated throughout this document.

TABLE 6.2-1. – TYPICAL MANEUVER AMPLITUDES FOR FIGHTER-CLASS AIRCRAFT

Signal	Peak magnitude
α	2°
β	1°
p	30 deg/sec
q	10 deg/sec
r	5 deg/sec
a_n	0.25 g
a_y	0.10 g

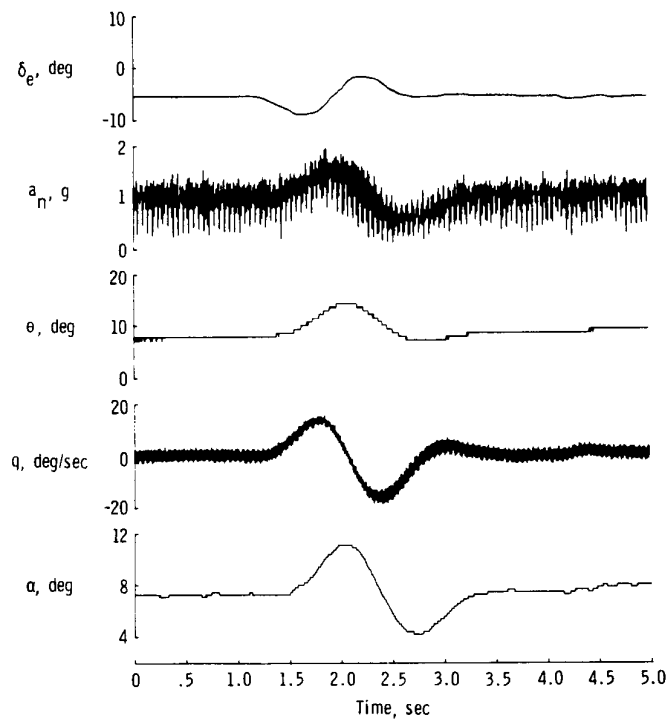

Figure 6.2-1. Oblique-wing 200-sample/sec data.

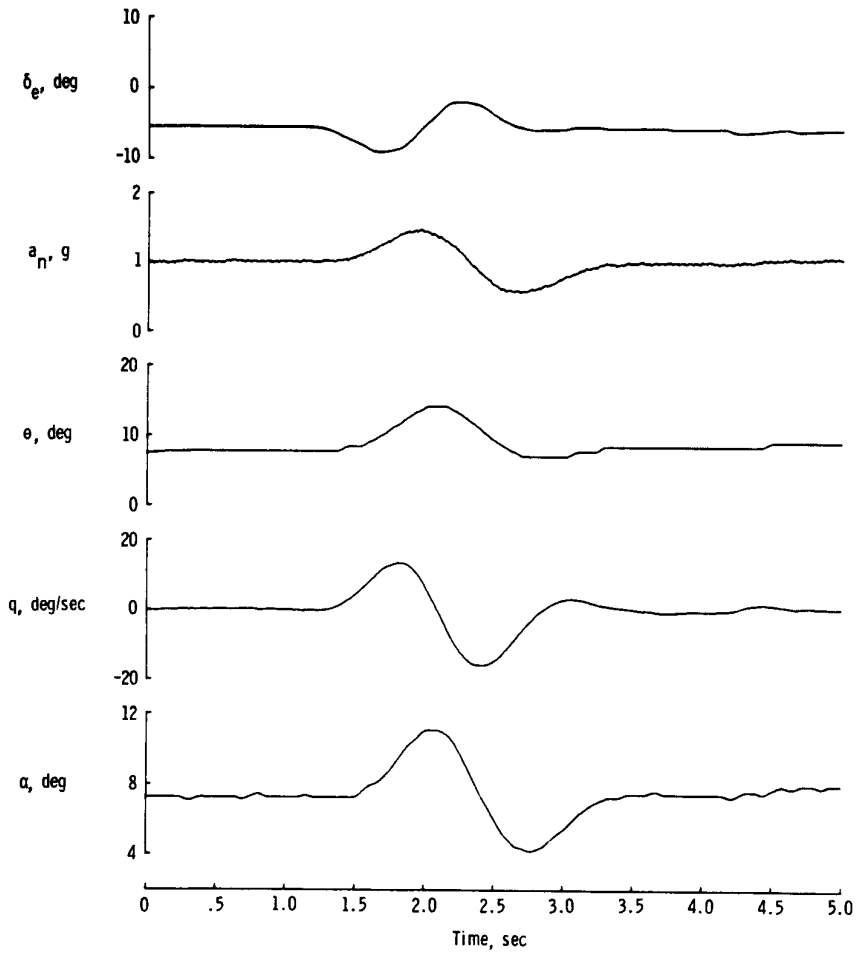

Figure 6.2-2. Time history of filtered data.

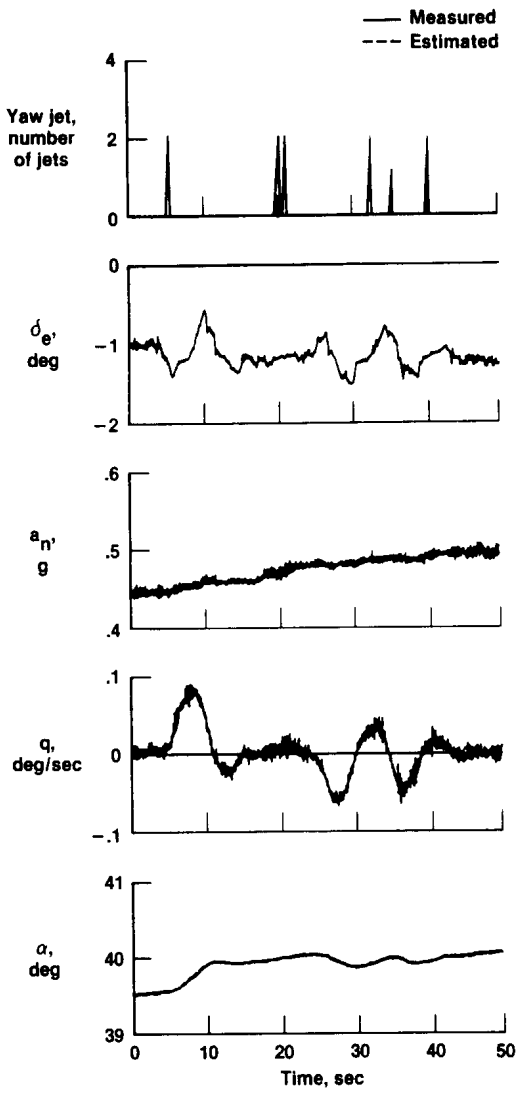

Figure 6.2-3. Tiny longitudinal maneuver of the space shuttle.

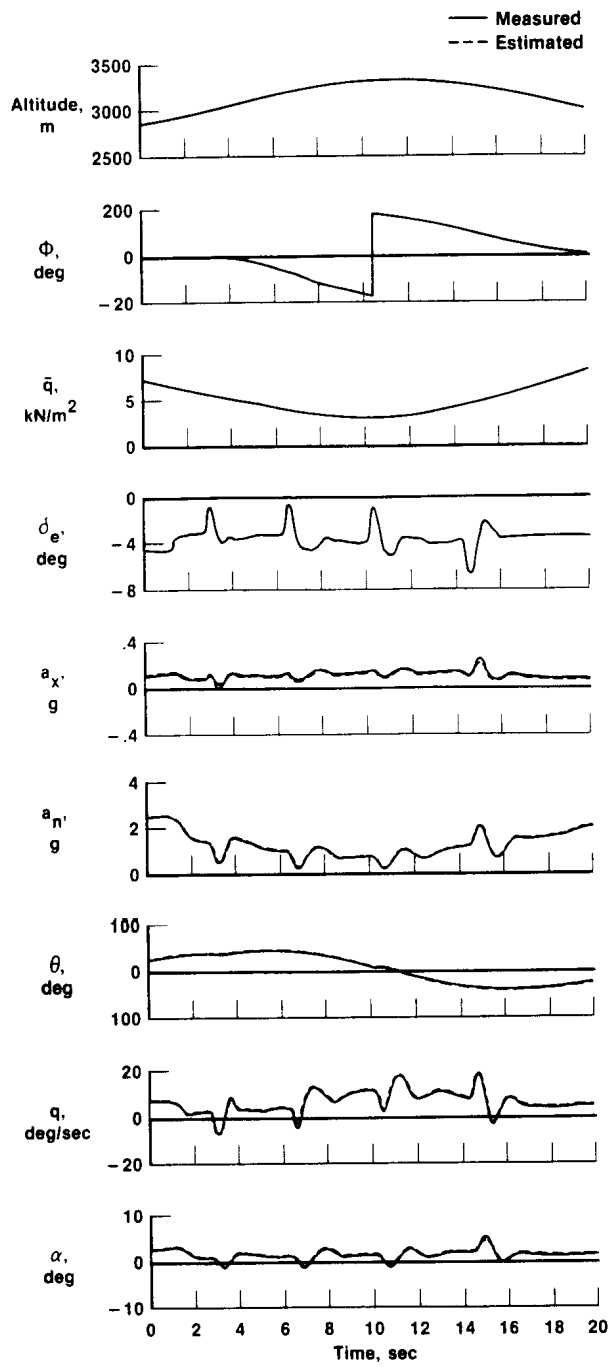

Figure 6.2-4. Small-perturbation maneuver for estimating $C_{m_{\dot{\alpha}}}$.

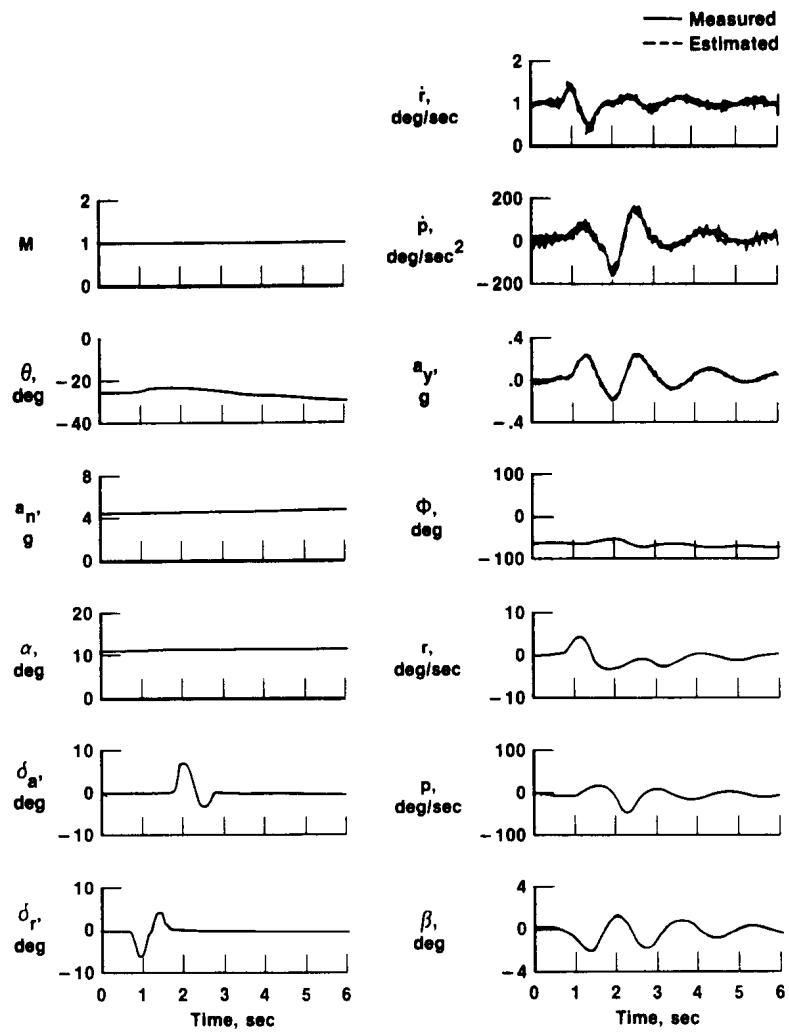

Figure 6.2-5. Maneuver at Mach 0.9, α = 12°.

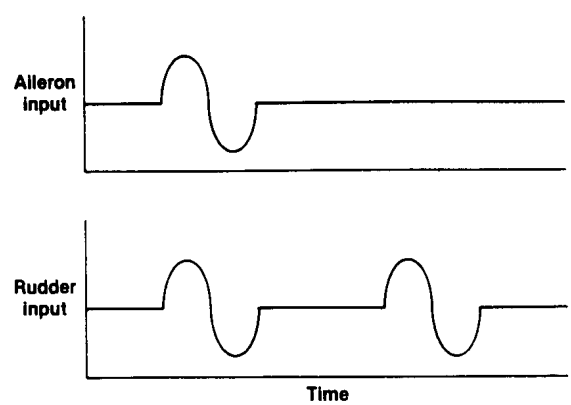

Figure 6.3-1. Combined aileron and rudder input.

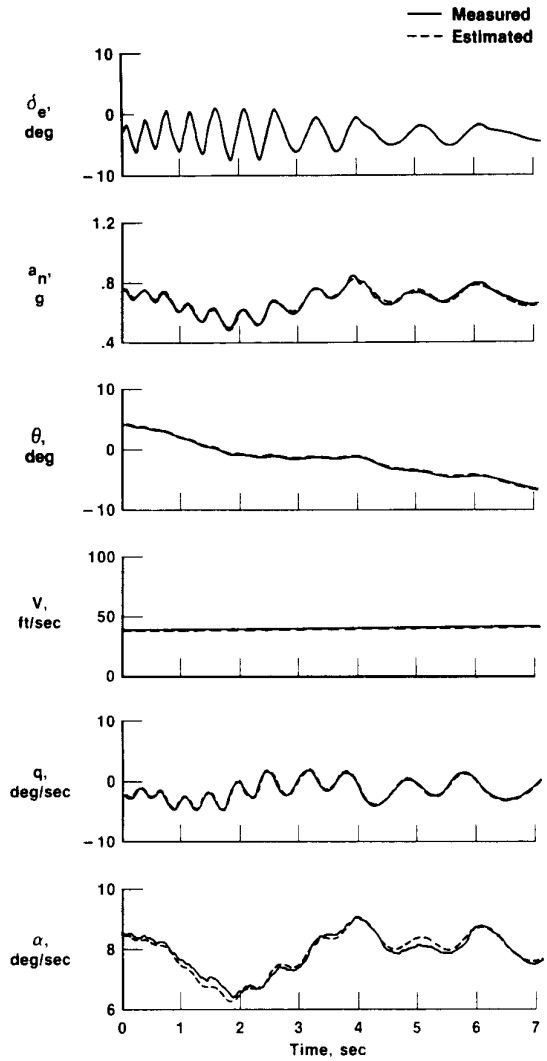

Figure 6.3-2. Frequency-sweep input.

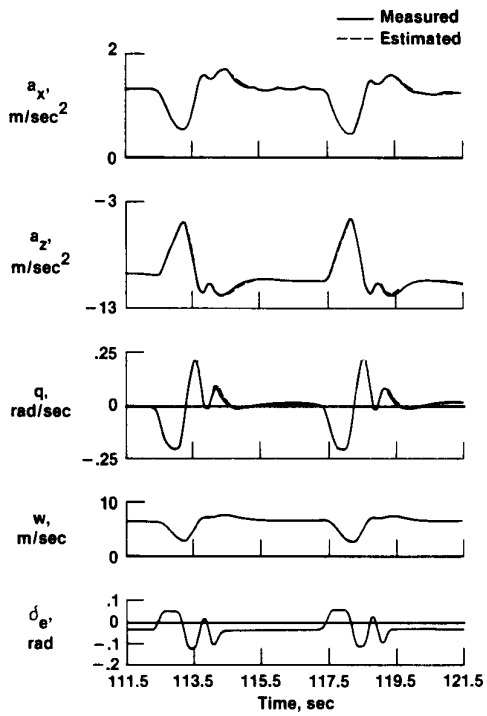

Figure 6.3-3. A 3211 input (Plaetschke and Schultz, 1979).

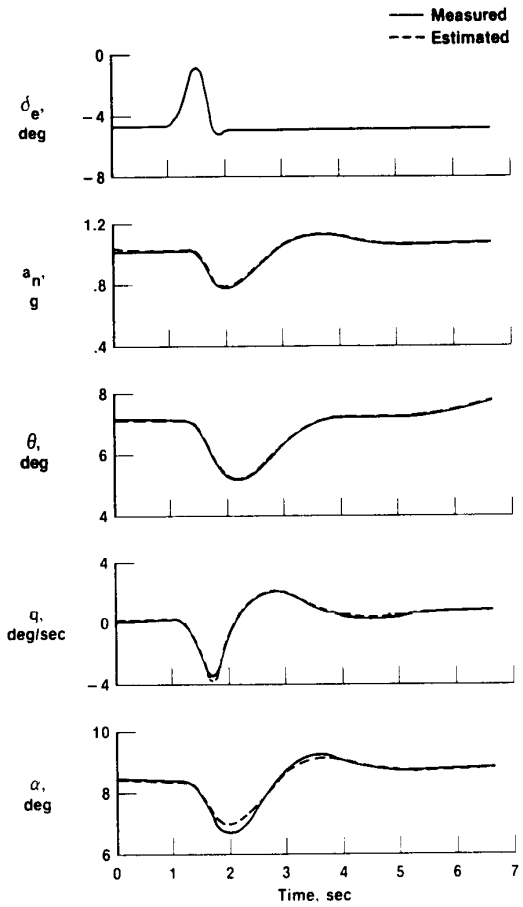

Figure 6.4-1. Longitudinal maneuver of a JetStar aircraft.

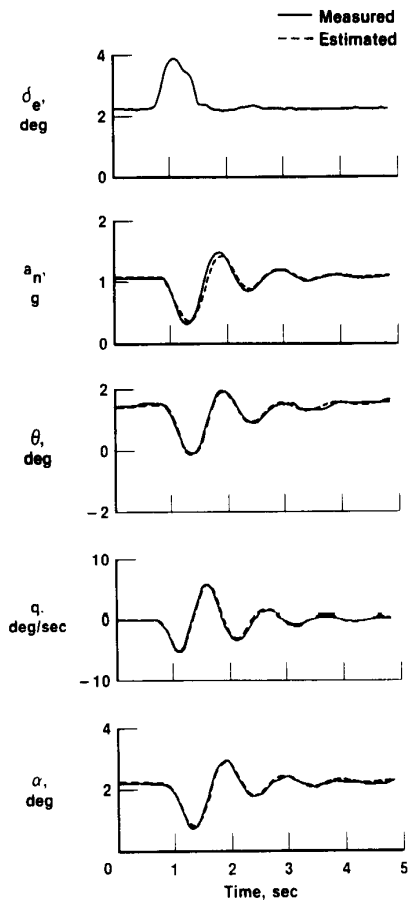

Figure 6.4-2. Longitudinal maneuver of an F-111A aircraft.

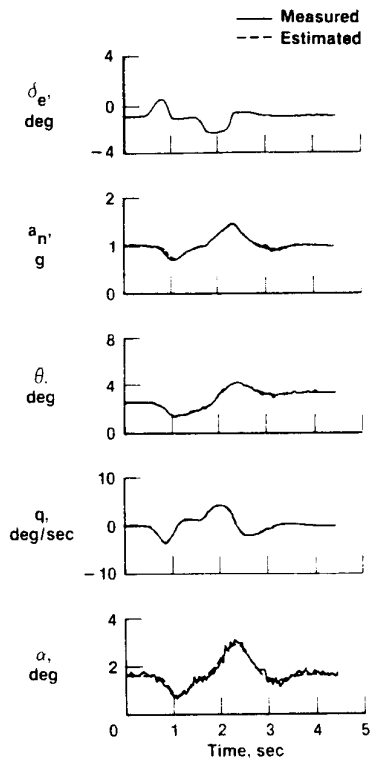

Figure 6.4-3. Longitudinal maneuver of a Beech-99 aircraft.

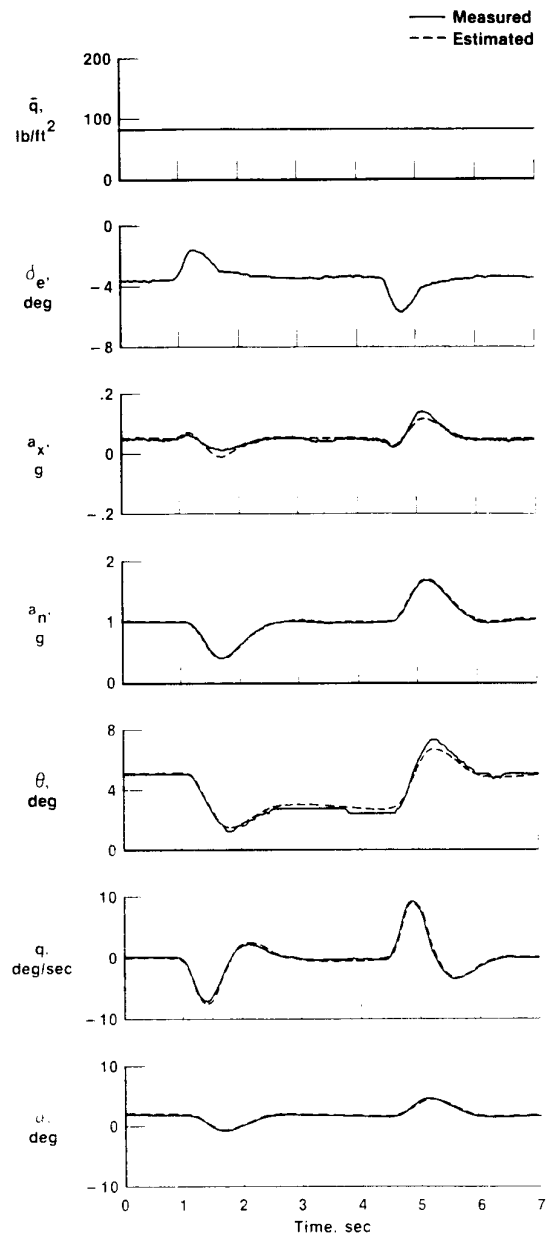

Figure 6.4-4. Longitudinal maneuver of a T-37 aircraft.

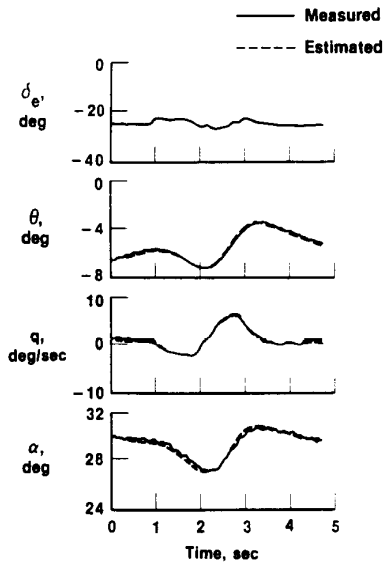

Figure 6.4-5. Longitudinal maneuver of a 3/8-scale F-15 model.

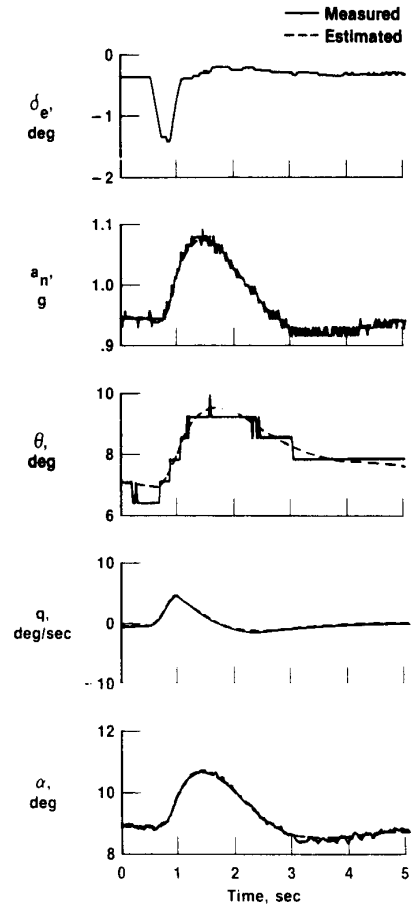

Figure 6.4-6. Longitudinal maneuver of a Piper PA-30 aircraft.

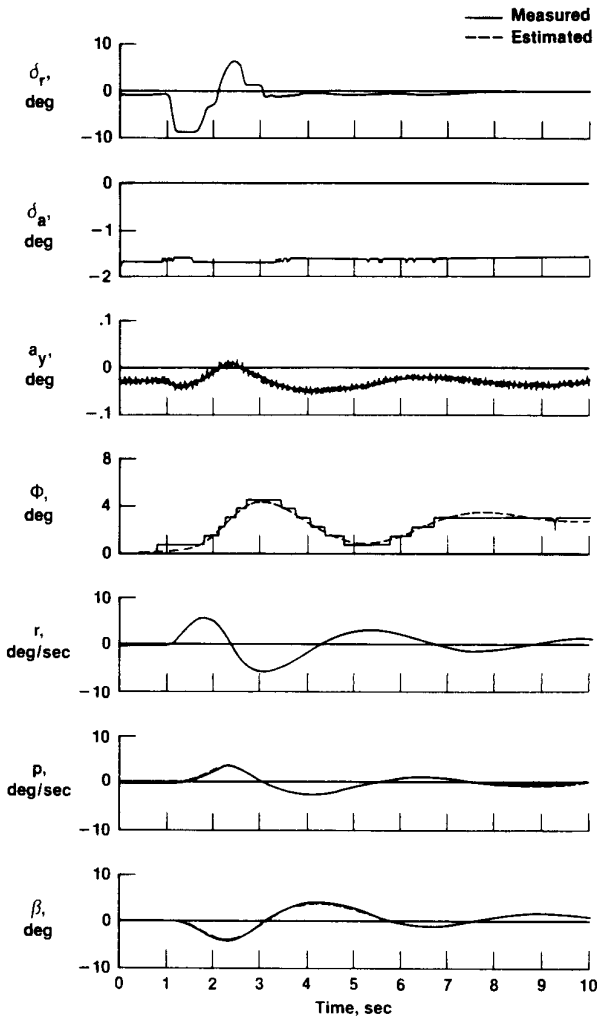

Figure 6.4-7. Rudder maneuver on a Piper PA-30 aircraft.

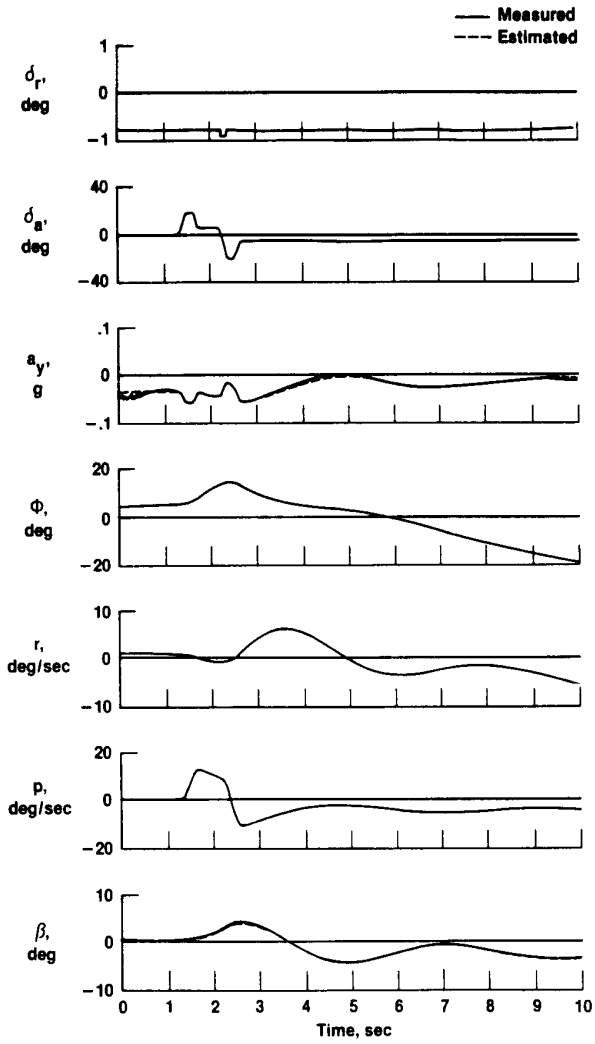

Figure 6.4-8. Aileron maneuver on a Piper PA-30 aircraft.

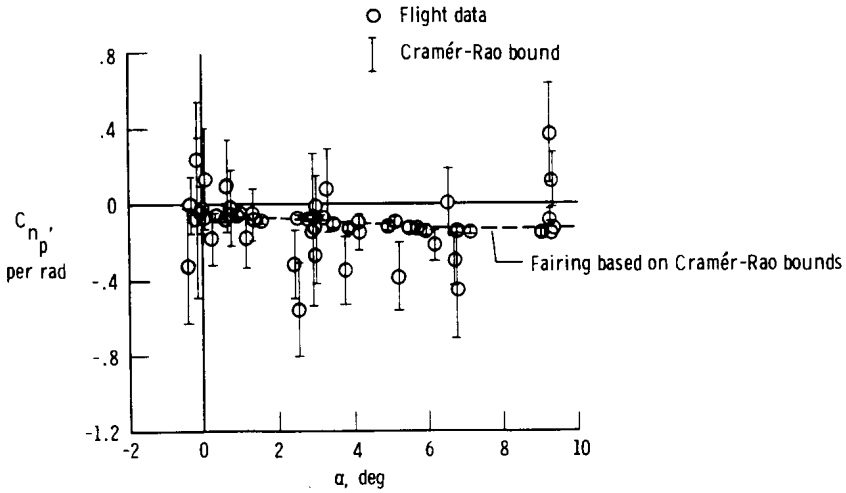

Figure 6.4-9. Estimates of C_{n_p} on a Piper PA-30 aircraft.

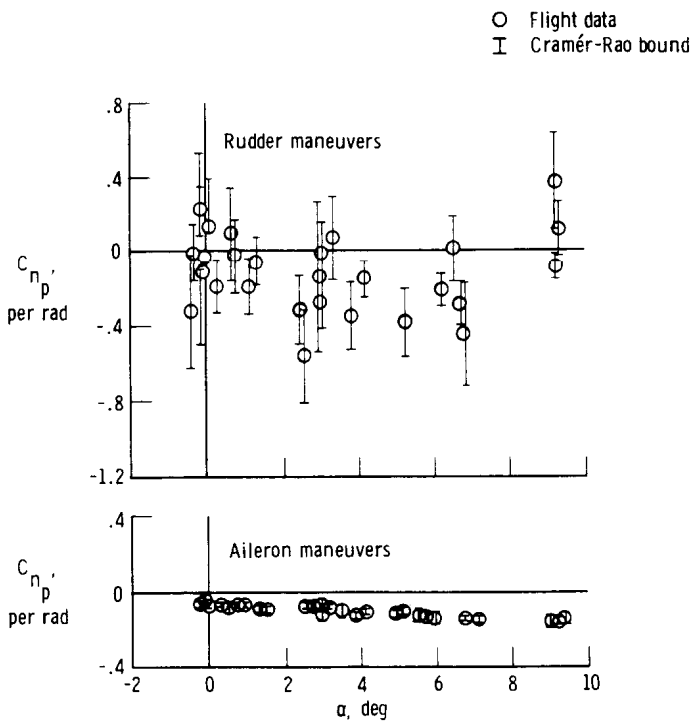

Figure 6.4-10. Estimates of C_{n_p}, segregated by maneuver type.

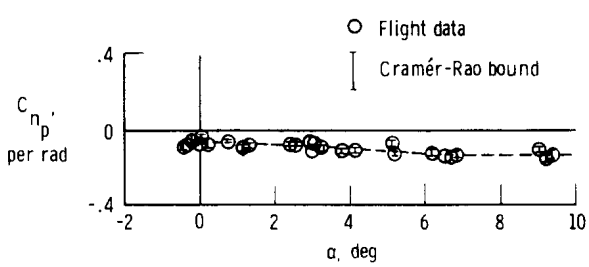

Figure 6.4-11. Estimates of C_{n_p} from multiple-maneuver analysis.

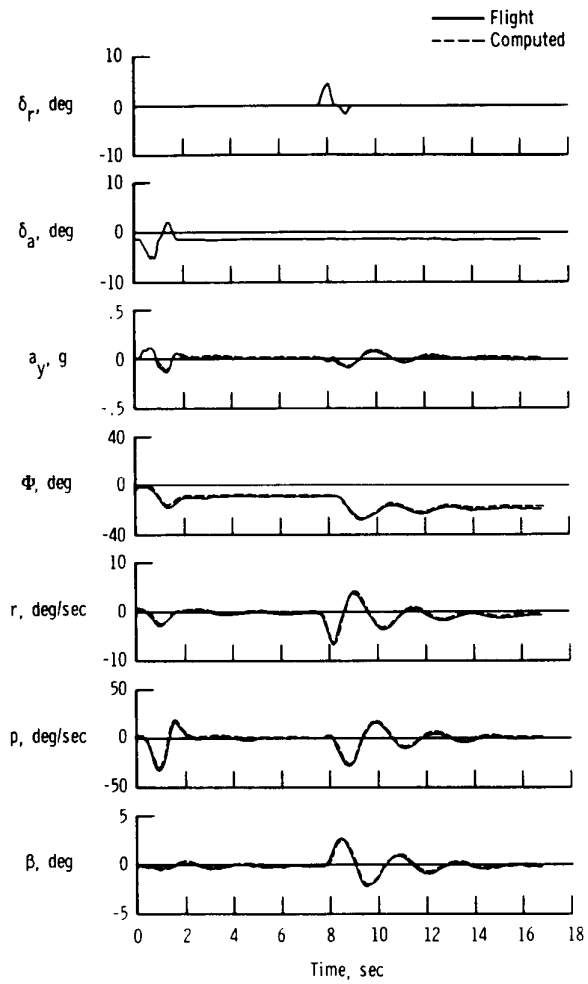

Figure 6.4-12. Lateral-directional maneuver of an F-8C aircraft.

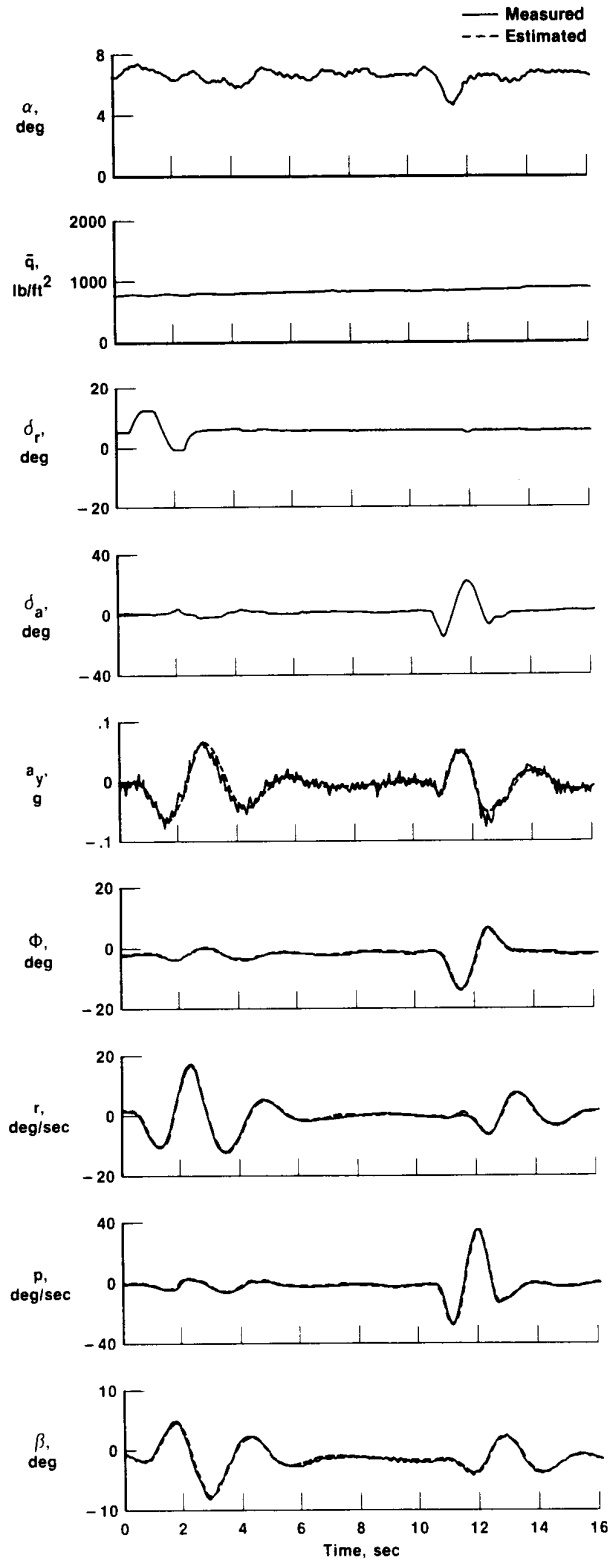

Figure 6.4-13. Lateral-directional maneuver of an oblique-wing aircraft.

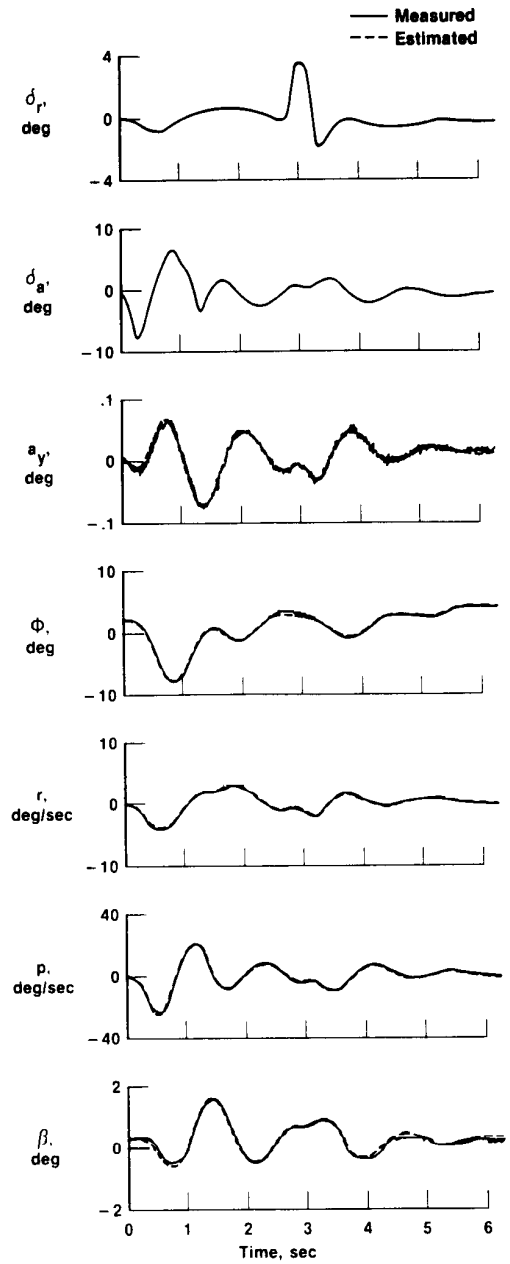

Figure 6.4-14. Lateral-directional maneuver of an HL-10 lifting body.

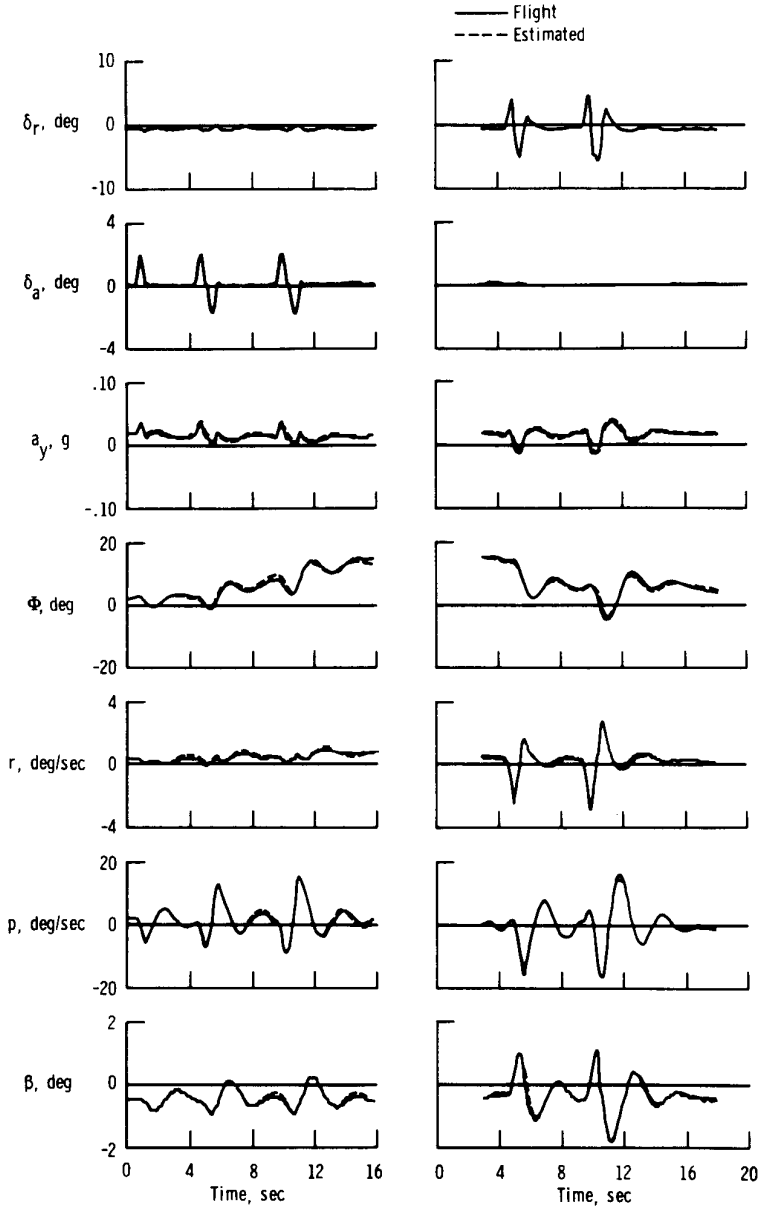

Figure 6.4-15. Lateral-directional double maneuver of an F-111A aircraft.

7.0 DATA ACQUISITION SYSTEM

The data acquisition system is an important part of stability and control flight testing. To do a good job, the data analyst must have a detailed understanding of how the data get from the sensor to the analysis program. This chapter discusses the classes of data acquisition systems and the problems that arise in data acquisition.

If you are unaware of details like how the data were filtered, digitized, time tagged, and recorded, then you cannot truly claim to understand exactly what the data tell about the airplane. Analyzing the gross characteristics of the aircraft requires few of these details, but the more information you try to extract from the test data, the more details you need to know about the data. If you have nothing to work from but a data tape and a list of the signals on the tape (straits we have seen many people in, and occasionally have occupied ourselves) avoid making grandiose promises or placing great confidence in your results.

You need to take a systems approach to analyzing the data system; that is, you need to look at the entire system, from the input to the output. The system includes sensors, recording system, postflight processing, and the connections between these parts. The connections and interactions between the components can be as important to the overall system performance as the individual components.

There is strong synergism among the negative effects of data system problems. Two problems, neither of which alone would unacceptably degrade results, can result in unusable data when they occur in combination. This is an area where simulation is often overused and misapplied; a little experience with real flight data is far more valuable. We often see simulation used to pick the minimum usable sample rate, resolution, and other specifications; this almost invariably results in severe analysis difficulties.

Indeed, we prefer to avoid the whole idea of the minimum acceptable data system. A minimum acceptable data system too often becomes unacceptable in practice, due to overreliance on misleading simulations, neglect of the synergism of data system problems, and unforeseen circumstances. The process of getting usable data from a minimum acceptable data system is typically lengthy and labor intensive, when it is possible at all. It is our observation that questions about minimum acceptable data quality portend a project destined to have unusable data; these questions often arise in connection with success-oriented philosophies (see Section 1.5), which ignore the fact that although it is theoretically possible to get by with the proposed system, the odds are against it. We are much more encouraged when a project starts with discussions of what kind of data system would do the best job, followed by a determination of how close to this ideal we can get.

Because the subject is so large, we devote two chapters to the data system. This chapter discusses the recording system and the connection of the sensors to the recording system. Some of the issues raised are also pertinent to other parts of the data system. Chapter 8 covers the instrumentation, which is the front end of the data system.

7.1 TYPES OF RECORDING SYSTEMS

There are several different means of recording flight data for analysis. The characteristics of the data recording system strongly influence the use of the data for stability and control derivative estimation. This section lists and describes the characteristics of some of the recording methods. The list is incomplete, reflecting the most common methods. The discussion is brief and general; this document is not the place for a detailed treatment of data recording methods.

The most common type of recording system used at Ames-Dryden is pulse code modulation (PCM). Figure 7.1-1 illustrates the major components of a simplified PCM system. PCM systems are inherently digital. The analog signal from each sensor first undergoes signal conditioning as required; this typically involves presample filtering and scaling to the voltage range of the analog-to-digital converter. The analog-to-digital converter samples each conditioned signal at regular intervals and outputs a digital scaled integer value. The units of this scaled integer are called counts.

The counts values are stored in digital form on an onboard tape or telemetered to the ground for similar tape storage (usually we store both onboard and telemetered data for redundancy). Postflight computer programs read the PCM tapes and convert the counts values into appropriate engineering units for analysis.

We find PCM systems to be a reliable means of recording high-quality data. The digital recording format has the advantage of not being subject to distortion. You may lose bits of data if there are transmission or media problems, but such dropouts are easy to detect. You do not find your data distorted by such subtle problems as temperature-dependent nonlinearities or tape drive speed, wow, or flutter. Of course, the sensor itself can have temperature or other environmental dependencies, but once the data are digitized you can be confident that it is subject to no further distortions. We recommend PCM systems as the primary data recording method for all serious test programs. Many of the discussions in this and later chapters emphasize the types of data problems common to PCM systems.

Our second most common form of data recording is at the other extreme of complexity – handwritten notes. These are often called pilot lap notes, although they can be made by other crew members or ground support personnel. The frequency response and time correlation of the hand-recorded data make their use obviously impractical for dynamic signals. In some cases, hand recording is practical but inconvenient compared with automatic recording systems, which do not misread dials or forget to record the value and which can automatically supply data to analysis programs. However, handwritten notes are sometimes the only practical means of recording numerous parameters. Basically anything that remains essentially

constant during a maneuver is suitable for hand recording; typical hand-recorded items include cargo loading, fuel weights, configuration, control system modes, altitude, airspeed, air temperature, engine speeds, and subjective turbulence level.

The most important point about hand-recorded data is to remember to record it. Do not confuse complexity with importance. Something like permanently recording the number and seating of crew members and their approximate weights may seem so trivial as to be unworthy of mention on the flight plans and data requests; it will not seem so trivial six months later when you have to throw out all the data from the flight because you cannot accurately reconstruct the weight and center of gravity.

Analog recording systems that we have used include frequency modulation, oscillograph charts, pulse-amplitude modulation, and pulse-duration modulation. Some of these methods have been used at Ames-Dryden in the past and are still in use at some installations. We generally regard these systems as inferior to PCM systems because of the noise and distortion inherent in analog systems.

Data are occasionally recorded by video cameras. We have, on occasion, acquired dynamic data by reading cockpit guages from video recordings. This is similar in many ways to hand-recorded data, except that the video method allows higher frequency data and better time correlation. This type of data gathering is labor intensive and (luckily) rare in stability and control analysis.

7.2 TIME TAGS

All practical methods for estimating aircraft stability and control derivatives assume exact knowledge of the relative time of each measurement. There are always, however, some small errors in our real knowledge of the times. Our real requirement is that any such errors be small enough to have negligible effect on the derivative estimates. This section investigates the effects of time errors to determine how much error we can tolerate.

We use the phrase "time tag" to refer to the information about the time of each measurement. The time tag can have several forms. In many cases the time tag is actually data recorded with each measurement or group of measurements. A simple example of such a recorded time tag is the time read from a watch written on hand-recorded notes. PCM or other automatic data systems often record time as a data channel. Video records are sometimes time tagged with an image of a digital clock on each frame.

Implicit time-tagging methods contrast with these explicit methods. A common form of implicit time tagging is to assume that the measurements come at a known regular sample rate and that the time of each measurement is thus a known increment after the time of the preceding measurement. Implicit time tagging can be as accurate as, if not more accurate than, explicit time tagging, but it provides no good way to merge data from two independent data streams.

We discuss aerodynamic coefficient identification package (ACIP) on the space shuttle to illustrate data time tagging. We present this simplified picture of the ACIP system only as one example of how data can be time tagged; different data systems use different time-tagging methods. The ACIP system uses a mixture of explicit and implicit time tagging. Figure 7.2-1 shows the layout of the ACIP PCM system for the initial flights. The system has a constant bit rate of 64,000 bits/sec. There are 8 bits/word, giving a word rate of 8000 words/sec, and there are 46 words per data frame, giving a frame rate of 173.9 frames/sec. There are 64 frames per data cycle, giving a cycle rate of about 2.7 data cycles per second.

Words 1 to 3 of each frame are used for synchronization and do not concern us here. The 19 high-sample-rate dynamic parameters are in words 9 to 46 of each frame. The counts values for each of these parameters occupy two consecutive 8-bit words. These parameters are recorded every frame, so their sample rate is the same as the frame rate of 173.9 samples/sec.

Word 8 contains the low-sample-rate data, mostly temperatures and other environmental data. The counts values for these parameters need only one 8-bit word. Word 8 contains a different one of these parameters for each of the 64 frames of the data cycle; thus there is room for 64 parameters in this word, each sampled at the cycle rate of 2.7 samples/sec. This method of recording several low-sample-rate parameters in the same word is called subcommutation. Word 4 contains the frame number (0 to 63), which distinguishes among the parameters subcommutated in word 8. Each of the subcommutated parameters is assigned a particular frame number.

Words 5 to 7 of each frame contain time data from a system clock. The exact form of the data is irrelevant here. These data constitute an explicit time tag for each frame of data. The time recorded is that at the start of the frame. The data signals are digitized one at a time as they are recorded, not at the start of the frame. To determine the sample time of each measurement, you must add the word number divided by 8000 to the time recorded at the start of the frame (minus a small time for the analog-to-digital conversion). Thus, the time of each measurement is the sum of the explicit time tag for the frame plus the increment implicit in the word position.

The absolute times of the data are irrelevant. It does not matter whether you express the time in Greenwich Mean Time, elapsed time from takeoff, or any other convenient time base. What does matter is that you know the relative times of any two measurements. If you merge two or more independent data sources, however, they must be referenced to the same time base; this requires system clocks synchronized to some agreed-upon standard.

It is important to distinguish between requirements on the times of the measurements and requirements on our knowledge of these times. The emphasis of this section is on the knowledge. We are not overly concerned, for instance, with requiring all of the signals to be sampled simultaneously. Almost all

analysis programs assume such simultaneous sampling; generalization is possible in principle, but it is awkward and time consuming. As long as we know the time of each measurement, we can interpolate all the signals to give the effect of simultaneous sampling. Interpolation does introduce errors, which could be avoided by using a system with simultaneous sampling, but such errors are usually negligible. The problems resulting from misinformation about the times of the measurements are far worse.

Errors in time tags can come from several sources. Any physical sensor has some time lag inherent in its response. The signal conditioning and analog-to-digital conversion introduce further lags. None of these lags are problems intrinsically, provided that you remember to account for them and that you have accurate values for them.

Most time-tagging errors result from carelessness and poor understanding of what the time tags mean. A typical error would be to use the ACIP time as the time for every measurement in the frame, as though the sampling were simultaneous, the instruments had instantaneous response, and there were no presample filters. The only way to ensure that you have avoided time errors is to understand exactly how the data get from the sensor to you, what the timing is for each step, and how the time-tag information is obtained.

We cannot give a universal formula for how to interpret time tags, because time-tagging methods differ so much. Some systems do sample all signals simultaneously or nearly simultaneously (Renz et al., 1981); such sample-and-hold techniques give the most easily interpreted time tags. At the other extreme, the space shuttle data system includes signals (not in the ACIP system) for which the data are digitized and stored in a buffer at 6.25 samples/sec but time tagged and recorded at 5 samples/sec. We eventually gave up trying to accurately time tag some of the shuttle measurements for which the timing depends on a flight computer that can be interrupted by higher priority tasks (like controlling the vehicle).

Figure 7.2-2 (from Steers and Iliff (1975)), illustrates some of the effects of time-tagging errors. These data are all based on one maneuver, with the data artificially time shifted to create timing errors. Each point on the figure represents an estimate of L_β using data with either p, β, or δa shifted by the indicated amount. A positive shift is defined such that the shifted signal lags the other signals. The figure shows little effect of β shifts on the L_β estimates for this maneuver, but shifts in p and δa have large effects. A positive shift of 0.1 sec in δa or a negative shift of 0.1 sec in p results in about a 50-percent change in the L_β estimate. Time errors of 0.1 sec or more have occurred in several programs in our experience. An intuitive explanation for the sensitivity to such shifts is that if δa has a positive shift or p has a negative shift, then the p motion appears to precede the δa input. Such apparently autonomous p motion looks like an instability, resulting in large effects on the derivative estimates. Steers and Iliff (1975) shows data similar to those in Figure 7.2-2 for other derivatives and other airplanes.

The conclusion from Figure 7.2-2 and similar data is that some of the derivative estimates are extremely sensitive to time shifts in some of the signals. The more complicated the model, the more sensitive it will be to time shifts. Our empirical rule for aircraft stability and control analysis is that time errors of more than 10 msec may cause problems and should be rectified. Errors of less than 10 msec are probably unimportant, and errors of less than 1 msec are certainly unimportant. We tend to ask for data accurate to 1 msec and settle for accuracy of 10 msec. For signals like altitude and airspeed, which have little dynamic content, the timing requirements are less stringent; accuracy of 20 to 50 msec is usually adequate for these signals.

7.3 ALIASING AND PREFILTERING

Digital recording systems such as PCM sample the continuous-time signals at discrete time points. Regardless of the recording method, you must digitize the data at some point in the processing in order to analyze the data on a digital computer. This digitization raises obvious questions of sample rate and resolution, dealt with in Sections 7.4 and 7.5. This section addresses signal aliasing, a subtler and often misunderstood aspect of the digitization. Every engineer working with digitized data should be aware of the general nature of aliasing because it can distort data beyond recognition if you are careless. This section briefly outlines the concept of aliasing. Oppenheim and Shafer (1975) gives a more thorough treatment.

All systems of which we are aware sample the data for each signal at regularly spaced time points, that is, at a constant sample rate. We henceforth assume such regular sampling; otherwise many of our analysis tools would be inapplicable.

The Nyquist sampling theorem (Oppenheim and Shafer, 1975) is perhaps the most basic principle of sampled signal theory and is the source of our current concern. This theorem states that a signal with a regular sample interval of Δ sec (a sample rate of $1/\Delta$ samples/sec) can contain no information at a frequency higher than $1/(2\Delta)$ Hz. This is a loose statement of the theorem but is adequate for our purposes. This limit frequency is called the Nyquist frequency or the half-sample frequency (because it is half the sample rate).

A continuous-time signal can, in principle, contain data at any frequency. Realistically, the data in any physical signal are restricted to a finite bandwidth, but the bandwidth is likely to be in the kilohertz range or higher, several orders of magnitude above aircraft stability and control frequencies.

At first glance, the Nyquist limit does not seem like a great problem to us. Frequency limits of 12.5 or 25 Hz (25 or 50 sample/sec data) are high enough to include virtually all useful aircraft sta-

bility and control data. The higher frequency data in the continuous-time signal is a combination of structural resonance, ac power frequencies, engine vibration, thermal noise, and other nuisance data. It is tempting to conclude that sampling actually improves the data by removing the unwanted junk above the Nyquist frequency; unfortunately, this optimistic conclusion is completely false.

When you digitize the continuous-time signal, the data above the Nyquist frequency do not innocuously disappear. Instead, the high-frequency data shifts to an apparent lower frequency. This frequency shift is called aliasing or frequency folding. If the Nyquist frequency is N, a continuous-time signal of frequency f aliases to

$$f_{alias} = \begin{cases} \mathrm{mod}(f,2N) & \mathrm{mod}(f,2N) < N \\ 2N - \mathrm{mod}(f,2N) & \mathrm{mod}(f,2N) > N \end{cases} \tag{7.3-1}$$

Figure 7.3-1 gives a simple illustration of aliasing of a constant-frequency signal. The solid line is a presumed continuous-time signal. The crosses mark sampling times. The frequency of the signal is slightly less than twice the Nyquist frequency; thus Equation (7.3-1) gives an alias frequency of slightly greater than zero. The dashed line is the alias, the low-frequency apparent signal of the sampled data. This alias signal is obviously a poor representation of the true continuous-time signal. Note that if the sample rate were decreased slightly, the alias signal would be a constant (zero frequency).

The effect of aliasing is that all the high-frequency noise, rather than conveniently disappearing, shifts to lower frequencies and contaminates the stability and control data. After you sample the data, there is no way to separate the original low-frequency component from the contaminating aliased data.

The same principles of aliasing apply to thinning digital data to lower sample rates. Figure 7.3-2 shows a time history of 200-sample/sec data from the oblique-wing vehicle (Maine, 1978). Figure 7.3-3 shows the power spectral density (PSD) of the normal acceleration signal. The large hump below 1 Hz in the PSD contains the useful stability and control data. The general noise level is about three orders of magnitude in spectral density below the stability and control data, but there are strong signals at 17.7, 60, and 80 Hz. The peaks at 17.7 and 80 Hz are from structural vibration, and the 60-Hz peak is from the propeller rotating at 3600 rpm. The 17.7-Hz signal is evident in Figure 7.3-2.

Figure 7.3-4 shows the PSD of a 25-sample/sec thinned normal acceleration signal. The structural and propeller response data are folded down and spread over the spectrum between 0 and 12.5 Hz, with the result that the spectrum is nearly white. The response peaks no longer stand out clearly, and the low-frequency data are significantly contaminated.

The only effective solution to aliasing is presample filtering, often called antialias filtering. Filter the continuous-time data to effectively remove the high-frequency components before sampling (and before thinning). This filtering must be done before the sampling or thinning step; there is no way to undo the aliasing after the data are sampled or thinned.

Figure 7.3-5 shows the PSD of the 200-sample/sec normal acceleration after a prethinning filter consisting of a 17.7-Hz notch and a third-order 20-Hz low-pass filter. This filter reduces the structural and propeller response to at or below the general noise level, low enough that their aliases will not significantly degrade the stability and control data. Figure 7.3-6 shows a time history of the filtered data. Thinning these data to 25 samples/sec gives a good signal for stability and control analysis.

Figure 7.3-7 shows the roll rate PSDs from the same data before and after filtering. Some of the same peaks appear in the unfiltered roll rate as in the normal acceleration. These peaks are almost imperceptible in the filtered roll rate. Figure 7.3-8 shows the angle-of-attack PSDs before and after filtering. None of the peaks observed in the other two signals are in the angle-of-attack spectrum. There is, however, evidence of a nose boom bending mode at about 4.5 Hz. We verified that this was a boom mode by hand plucking the boom and visually estimating the response frequency (this is a prime example of the kind of test that is trivial if the airplane is still around in the flight test configuration). The filter does not attenuate this mode, but the frequency and amplitude are such that the boom mode does not create problems. The boom mode is perceptible in Figures 7.3-2 and 7.3-6, barely larger than the resolution (look particularly near the end of the time history in Figure 7.3-6 and notice that the apparent bit noise is fairly coherent at that frequency).

Any data system design process should include careful consideration of presample filter requirements. We have occasionally seen good data from systems with no presample filtering other than that from stray capacitance and the sensors' inherent response characteristics, but do not count on such fortuitous circumstances. As a general rule, a low-pass filter at 40 percent of the Nyquist frequency is reasonable. For systems with high sample rates, a first-order filter is usually adequate; higher order filters may be necessary in low-sample-rate systems. You should, however, consider individual requirements in detail, rather than universally applying these rules. We do not cover details of filter and specification design in this document. See Williams (1981) for such a discussion.

As a user of a data system, you should know what presample filtering is done to the data. Presample filters are a major and often overlooked source of time lags, which need to be corrected in time tagging the data (see Section 7.2).

You may have to ask very specific questions to get information about the presample filtering in a data system (and it may take a while to find the right person to ask in a large project). The people responsible for the data system might take presample filtering so much for granted that they forget to

tell you about it. More than once we have explicitly asked whether there were any filters on a signal and been told that there were none, only to discover later that the response meant no filters other than, of course, the presample filter. We now habitually ask specifically about presample antialias filters. If the response is a blank stare or the equivalent, you have not yet found the right person to ask. (If the person responsible for data system design has not heard of antialiasing, you have serious problems.)

A systems approach is crucial to the design and description of antialias filters (or filters in general). Always look at the filter in context, not as an isolated component. For instance, the break frequency of a passive filter depends on the source impedance loading. The quoted break frequencies are at zero load; the actual break frequency under load can be significantly lower. The characteristics under actual loading conditions are what really interest you.

7.4 SAMPLE RATE

The sample rate of the digitized data is an important data system specification. This section discusses the considerations affecting the choice of sample rates. The discussion is closely related to that of Section 7.3 on presample filters.

The higher the sample rate, the better the digital data represent the continuous-time signal. The improvement eventually reaches a point of diminishing returns where it fails to justify the extra data processing and analysis time. There is certainly little point in sampling much faster than the sensor can respond. Data system bandwidth places upper bounds on practically achievable sample rates.

We have observed no benefit from analyzing aircraft stability and control data at more than 25 or 50 samples/sec. However, we often find it prudent to request the digitized data at 100 or 200 samples/sec; we then digitally filter and thin the data to 25 or 50 samples/sec for analysis. Although it requires higher bandwidth, more processing time, and extra processing steps, there are several advantages to obtaining the digitized data at higher sample rates than needed for analysis.

The first advantage is the flexibility of digital filtering. Analog presample filters must be designed before flight and are difficult to change. If you digitize data at 25 samples/sec, say, the Nyquist frequency is 12.5 Hz (see Section 7.3). To filter out noise above the Nyquist frequency, you need a presample low-pass filter with a break frequency near 5 Hz, based on our criterion of 40 percent of the Nyquist frequency. The filter roll-off must be steep enough to eliminate any strong structural resonances in the 20 to 30 Hz area, which folds down into the critical 0 to 5 Hz area after sampling. Such a high-order low-frequency filter would introduce significant distortion in the stability and control data.

If you digitize data at 100 samples/sec, the analog presample filter would have a break frequency around 20 Hz, which is high enough to avoid significant distortion. Before thinning the data to 25 samples/sec, you can examine the spectrum of the digitized data and design digital notch filters for any troublesome low-frequency structural resonances. With these resonances removed, you can probably get by with a prethinning low-pass filter with a break frequency around 10 Hz. This combination of a notch and a low-pass filter introduces less distortion than the stronger low-pass filter otherwise required. In principle, you could design an analog notch and low-pass to achieve the same effect, allowing you to sample at 25-samples/sec, omitting the intermediate 100-sample/sec data. The difference is that the analog notch must be designed based on predictions and cannot be changed after the flight.

We could summarize this argument by saying that any information lost or contaminated in analog prefiltering and digitization is irretrievable. The only way to fix mistakes is to revise the system and fly again. Hardware revisions and extra test flights are much more expensive than a little extra data processing, so it is best to digitize at a higher sample rate than you really need, leaving some room for error. Regard the extra processing time as insurance.

The second advantage of digitizing at 100 or 200 samples/sec is that it lessens the potential problems with time skews caused by nonsimultaneous sampling. This advantage is not as important as the filtering issue, because known time skews can be handled by interpolation. With the higher sampling rate, the skews may be small enough to safely ignore, saving a step in processing and thus partially compensating for the extra filtering step. More important, any errors in evaluating and correcting the skews will likely have smaller effects at the higher sample rates because the terms are smaller.

The lower limits on acceptable sample rates are determined by the necessity to accurately represent the continuous signals. Sample rates around 10 samples/sec have proven adequate on several aircraft (Brenner et al., 1978). When using low-sample-rate data, consider several factors. First, the motion of the aircraft control surfaces can have frequency content well above the aircraft natural response frequencies. The sample rate must be adequate for the control position measurements as well as the response measurements. This consideration reduces the differences between sample rates required for large and small aircraft. The sample rate requirements for stability and control analysis of the space shuttle, for instance, are dominated by considerations of the reaction control jets. Although the vehicle natural frequencies are low during the early part of entry, the reaction control jets can fire in 80 msec bursts. Sample rates lower than 25 samples/sec would unacceptably distort the jet signals.

Second, low sample rates accentuate the effects of other problems such as time skews, noise, low resolution, and modeling error. If you choose the lowest sample rate that works well on simulated data, we can almost guarantee poor results from flight.

In summary, we usually ask for sample rates of 100 to 200 samples/sec, digitally filter the data, and thin it to 25 to 50 samples/sec for analysis. You can usually get by with digitization at 25 or 50

samples/sec, provided you are careful about filtering; this approach leaves less room for error. Sample rates as low as 10 samples/sec have occasionally proven adequate, but we recommend against such marginal rates.

7.5 RESOLUTION

Another issue raised by digitization is the resolution or quantization level of the digitized signal. The digitized signal is an integer number of counts. The resolution of this signal is exactly one count, which corresponds to some number of engineering units that depends on the sensor calibration.

If the resolution magnitude is much smaller than the noise level of the digitized data, then you can ignore the issue. This is the case, for instance, for the ACIP system on the space shuttle, discussed in Sections 6.2.2 and 7.2. The ACIP lateral acceleration, in specific, uses a 14-bit digital signal to cover the ± 0.5-g range. The resolution is thus $1/2^{14} \cong 0.00006$ g. Structural vibration causes a noise level of about 0.01 g, so we can safely ignore any resolution problems.

The stairstep appearance of the α and Θ signals in Figure 7.3-2 is typical of data where the resolution is coarser than or comparable to the noise level. Note that digital filtering (but not analog filtering) can disguise this characteristic appearance, as in Figure 7.3-6.

The discussion of signal-to-noise levels in Section 6.2.2 applies to the issue of establishing acceptable resolution limits. A resolution of about 1/10 of the maneuver size is about the poorest you can ever accept. In this case, anticipate the immediately obvious potential for derivative errors on the order of 10 percent; actual errors in some derivative estimates might be much larger. Resolution of 1/100 of the maneuver size is fairly good. Exact criteria vary from system to system and even signal to signal. You can get by with much lower resolution in relatively unimportant signals (like bank angle) than in critical signals (like roll rate or aileron position).

We have most commonly used 10-bit PCM systems, giving a resolution of 1/1024 of full range. The small-perturbation stability and control maneuvers typically cover only a small part of the full instrument range, so the resolution is not 1/1024 of the maneuver size. The worst resolution is often on Euler angles, which are scaled over 360°, giving about 1/3° resolution; this can be near our cited minimum of 1/10 of the maneuver size. The Euler angles are among the least critical signals, so this low resolution is acceptable.

We have successfully used 8- and 9-bit PCM systems on some programs. Such systems require specific tailoring of the data ranges to achieve acceptable resolution for stability and control data. This tailoring often involves duplicating several signals on two or more channels with different ranges, as described in the following.

To provide data on large maneuvers or emergencies, scale a coarse channel for the full flight envelope. You want a full-range data channel in a flight test of a new vehicle, even if you intend to fly in only a small portion of the range. If a loss of control or other unexpected emergency results in flight at extreme conditions, you must be able to study the data thoroughly to find out what happened, how to prevent it from happening again, and how any emergency recovery procedure worked. Full-range data channels for such purposes are usually considered a mandatory safety requirement. This coarse-resolution channel may have unacceptable resolution for small-perturbation stability and control maneuvers; for example, an Euler angle scaled over 360° in an 8-bit PCM system has a resolution of about 1.4°, not worth using for most small maneuvers.

A second channel records the same signal, scaled to a smaller range to obtain better resolution. Choose the range of this high-resolution channel to be slightly larger than the small-perturbation maneuvers, allowing margin for error. If you scale an 8-bit channel for twice the expected maneuver size (a healthy margin for error), the resolution is 1/128 of the maneuver size.

Wallace-O-Leonard sensors use another method of combining data from two channels to get better resolution than that obtainable from a single channel. These sensors use a multiturn potentiometer, each turn corresponding to a different calibration. A second coarser channel determines which calibrations to use for the fine channel. Section 8.7 presents some data from a Wallace-O-Leonard sensor.

Using one of these channel duplication schemes, you can get good resolution from even an 8-bit PCM system. Note that range tailoring requires either that you give up the full range data or that you use extra data channels. We occasionally find channel duplication necessary even on 10-bit PCM systems for high-performance aircraft. A highly maneuverable fighter is likely to need a range of about -5 to +10 g on normal acceleration. With a 10-bit system, this gives a resolution of 0.015 g, which is perhaps acceptable but not as good as we like.

In several recent programs we have used PCM systems ranging from 11 to 16 bits of resolution. Such high-resolution systems simplify data system design considerably. You seldom need to make design compromises or kludges like duplicating channels to improve resolution. Note that duplicating a 10-bit channel results in the use of 20 bits to effectively achieve resolution comparable to an 11- or 12-bit channel. A true 11- or 12-bit system achieves this resolution with much more efficient use of the data bandwidth; furthermore, it eliminates the analyst's problem of whether to use one of the duplicate signals or some combination of them.

We used four maneuvers from a Piper PA-30 aircraft to study the effects of reduced resolution. Figure 7.5-1 shows a fit of one of the maneuvers. The data were sampled at 200 samples/sec and thinned to 50 samples/sec for analysis. These data came from a 9-bit PCM system. Table 7.5-1 gives the resolu-

tions of the signals in the figure. Most of the signals have resolutions in the neighborhood of 1/50 to 1/100 of the maneuver size. The worst resolution is that of a_y at about 1/20 of the maneuver size, and the best is that of δ_a at about 1/200 of the maneuver size. We simulated the effect of reduced resolution data systems by quantizing the data of Figure 7.5-1 at levels of 2, 4, 8, 16, and 32 times poorer resolution. Figure 7.5-2 shows the data of Figure 7.5-1 with the resolution reduced by a factor of 16. The a_y signal is obviously worthless. There is recognizable signal content in β, p, r, and ϕ, although the resolution is from 1/3 to 1/6 of the maneuver size, worse than our cited minimum acceptable level of 1/10 of the maneuver size. The resolution on δ_a is barely within the acceptable range.

Figure 7.5-3 shows the estimates of three derivatives as a function of the resolution reduction factor (to aid visualization, the symbols are slightly offset from the exact factors of 1, 2, 4, 8, 16, and 32). Errors resulting from the poor resolution are significant at a resolution reduction factor of 8 and larger at a factor of 16, corresponding to Figure 7.5-2. The data at a resolution reduction factor of 32 are completely unacceptable. These results roughly verify our cited minimum resolution level of 1/10 of the maneuver size.

Do not expect results in most cases to be as good as those shown in this example. This example represents the most benign situation in several regards. First, the best resolution is on the control surface, which is one of the most resolution-critical signals. Second, the data quality is excellent, except for the resolution: The noise level is low, the sample rate is high, the sensors are of good quality, and the system was carefully calibrated. Third, the maneuver is mild, and the aircraft is quite linear in this flight regime.

TABLE 7.5-1. – PIPER PA-30
DATA RESOLUTION

Signal	Resolution
β, deg	0.1
p, deg/sec	0.25
r, deg/sec	0.15
ϕ, deg	0.75
a_y, g	0.003
δ_a, deg	0.1

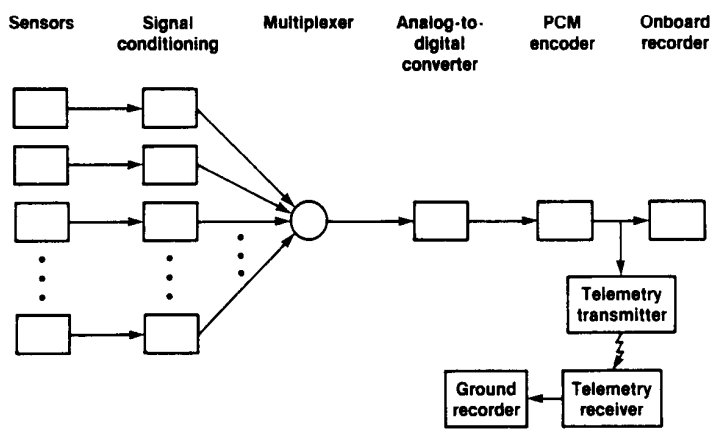

Figure 7.1-1. Simplified PCM system.

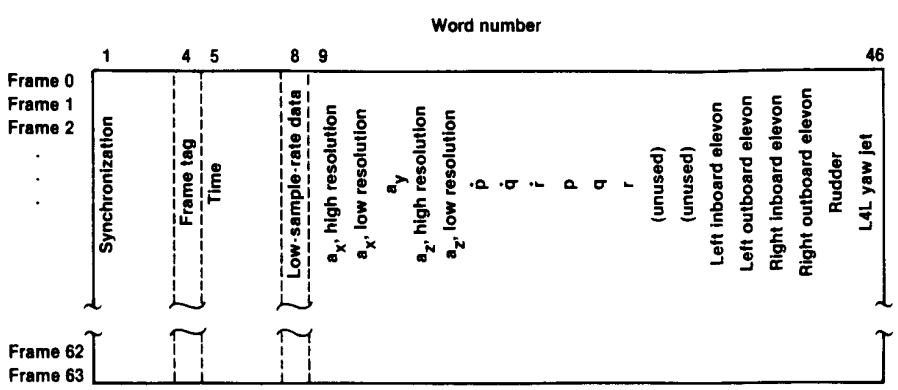

Figure 7.2-1. ACIP PCM system layout.

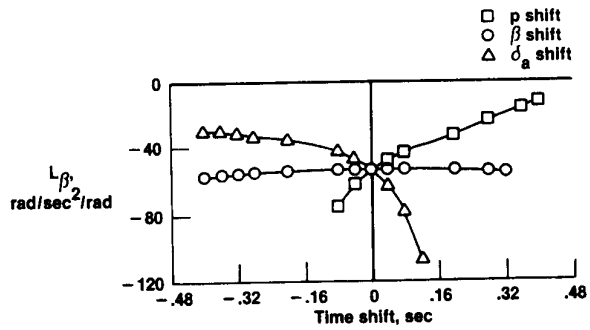

Figure 7.2-2. Effects of time shifts.

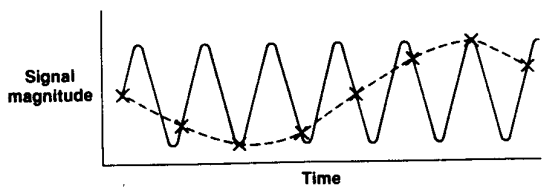

Figure 7.3-1. Example of aliasing.

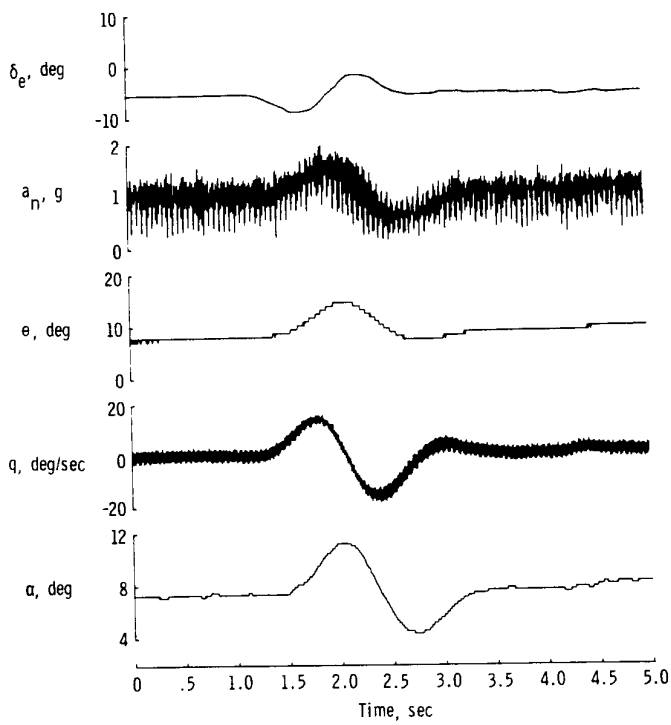

Figure 7.3-2. Oblique-wing 200-sample/sec data.

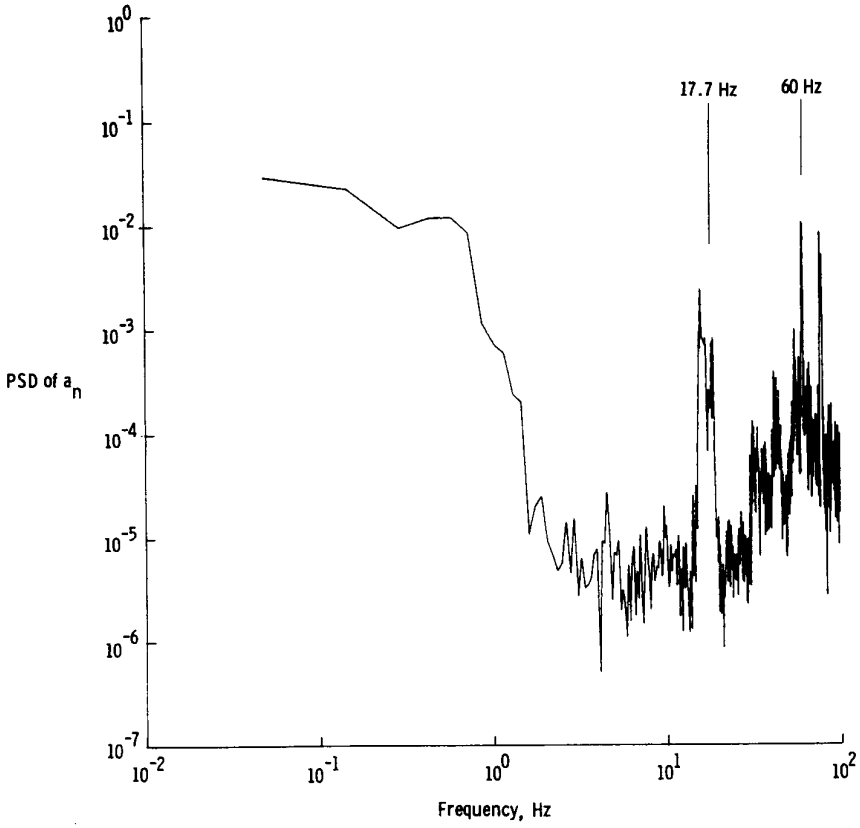

Figure 7.3-3. Power spectral density of 200-sample/sec normal acceleration.

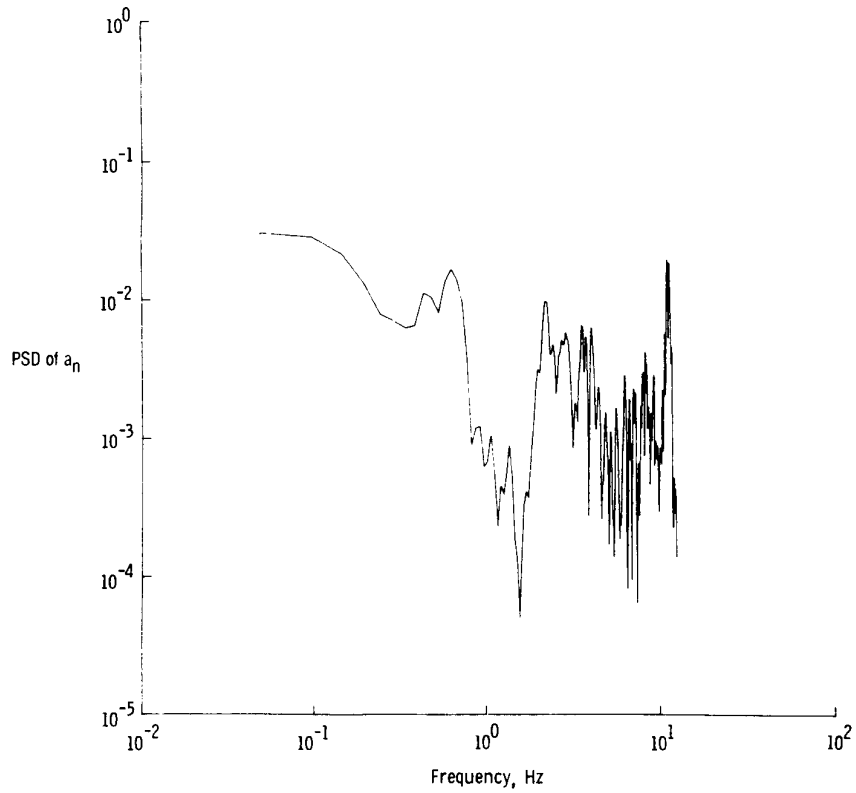

Figure 7.3-4. Power spectral density of 25-sample/sec thinned normal acceleration.

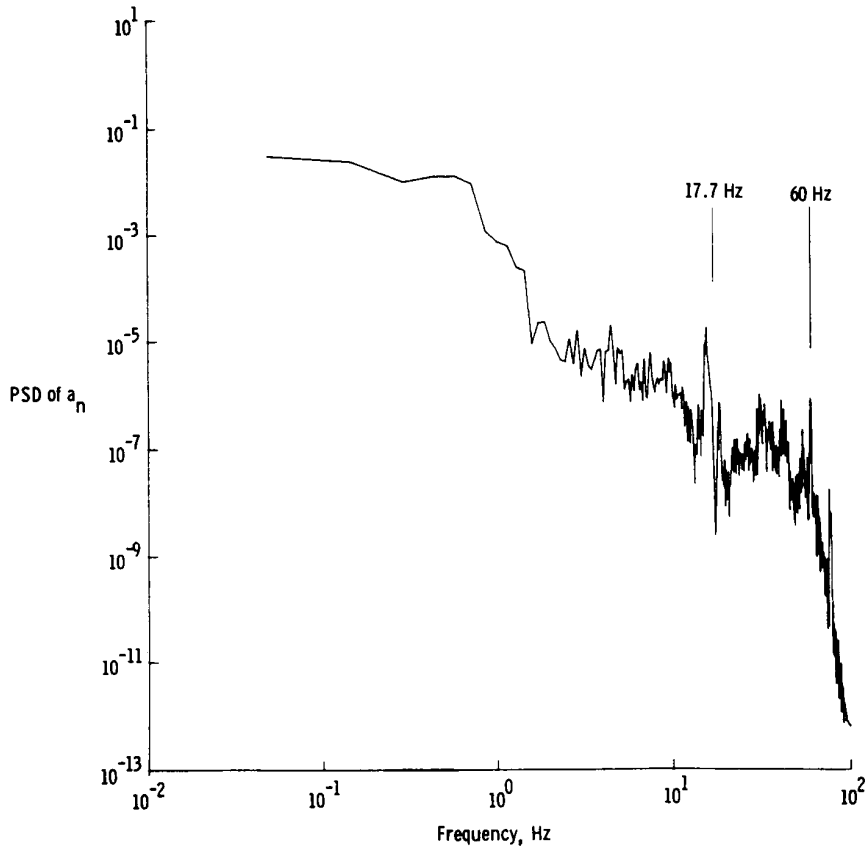

Figure 7.3-5. Power spectral density of filtered normal acceleration.

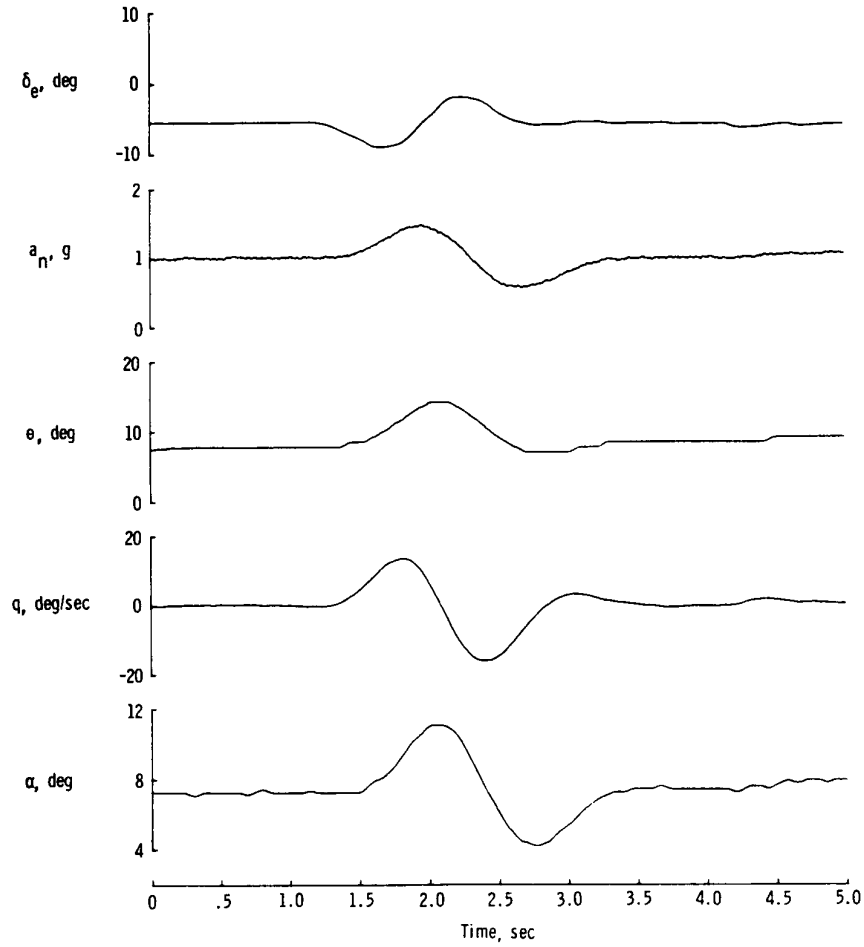

Figure 7.3-6. Time history of filtered data.

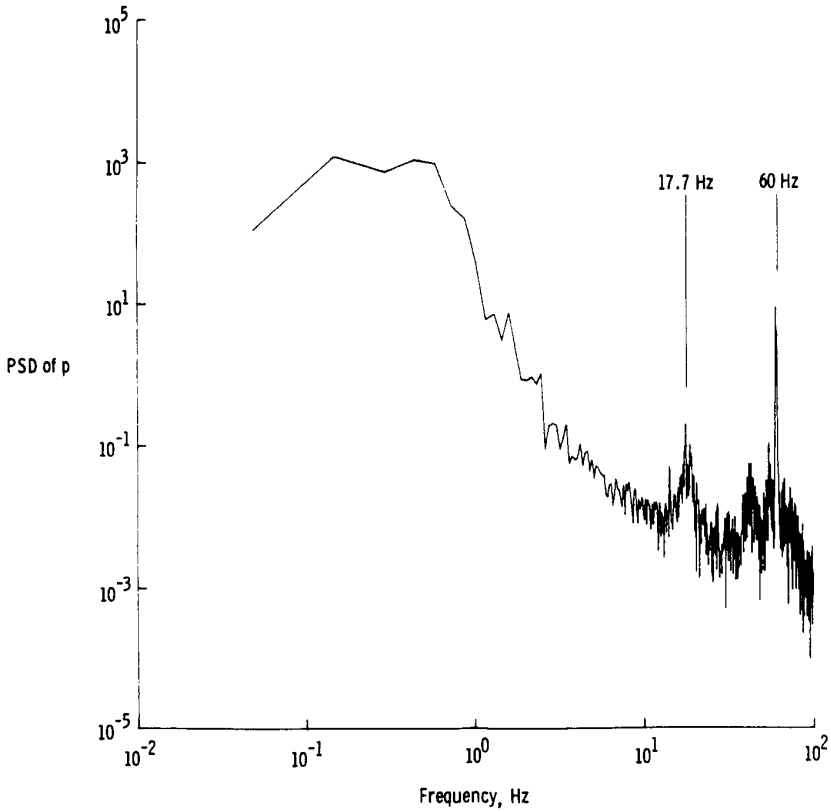

(a) Unfiltered roll rate.

Figure 7.3-7. Power spectral densities of unfiltered and filtered roll rates.

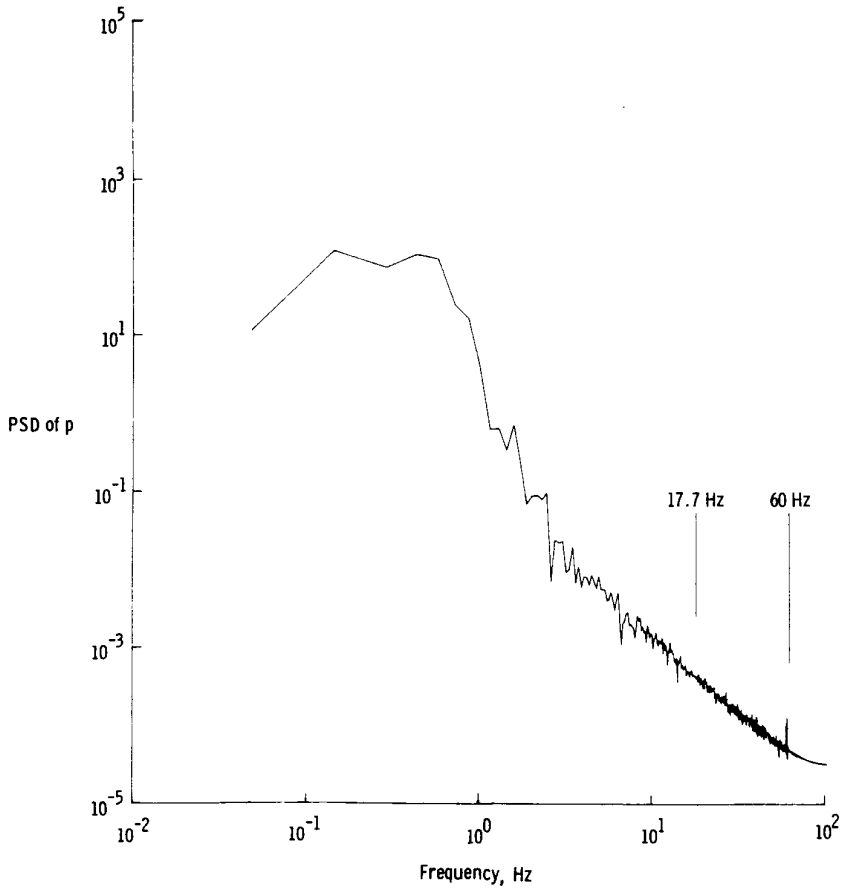

(b) Filtered roll rate.

Figure 7.3-7. Concluded.

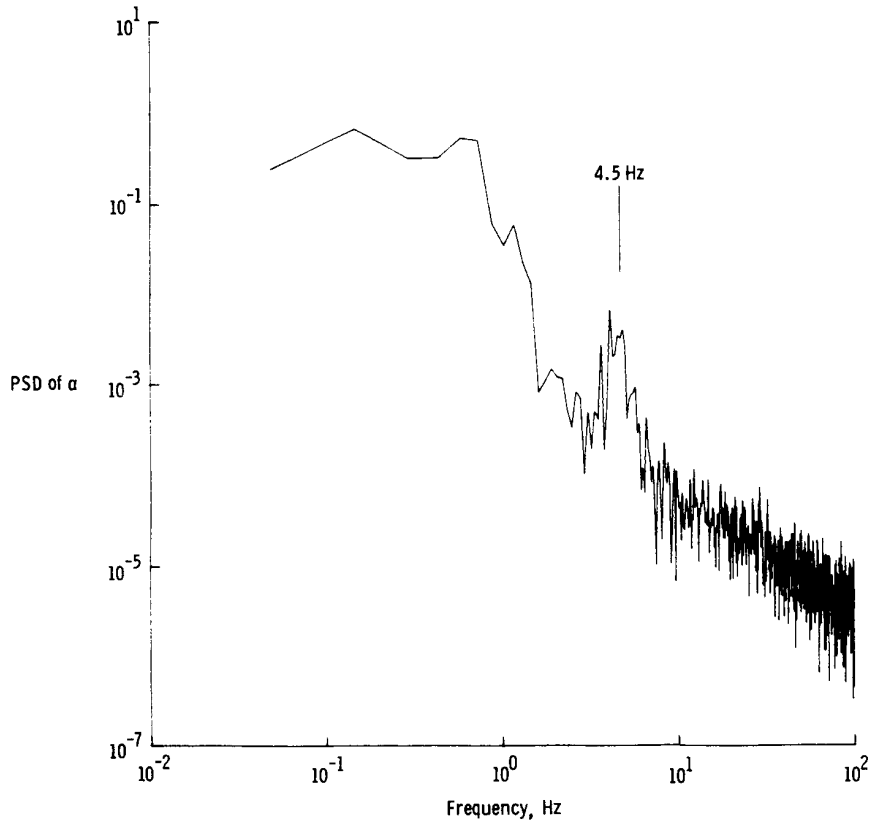

(a) Unfiltered angle of attack.

Figure 7.3-8. Power spectral densities of unfiltered and filtered angle of attack.

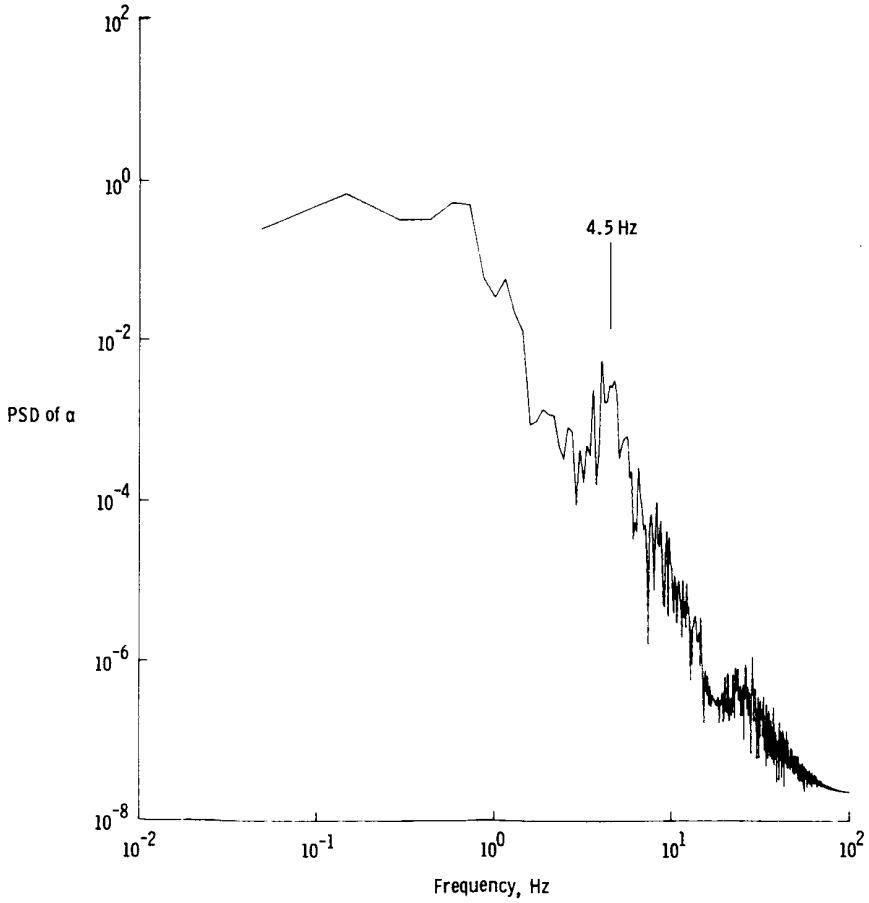

(b) Filtered angle of attack.

Figure 7.3-8. Concluded.

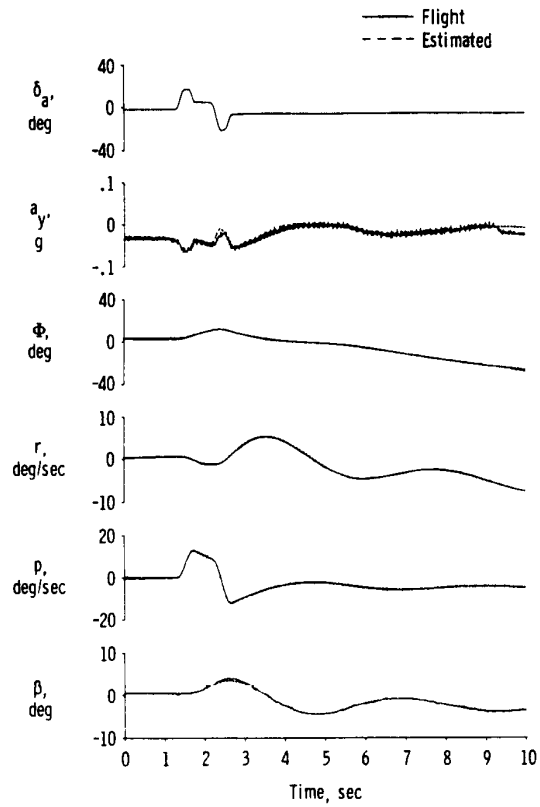

Figure 7.5-1. Piper PA-30 time history.

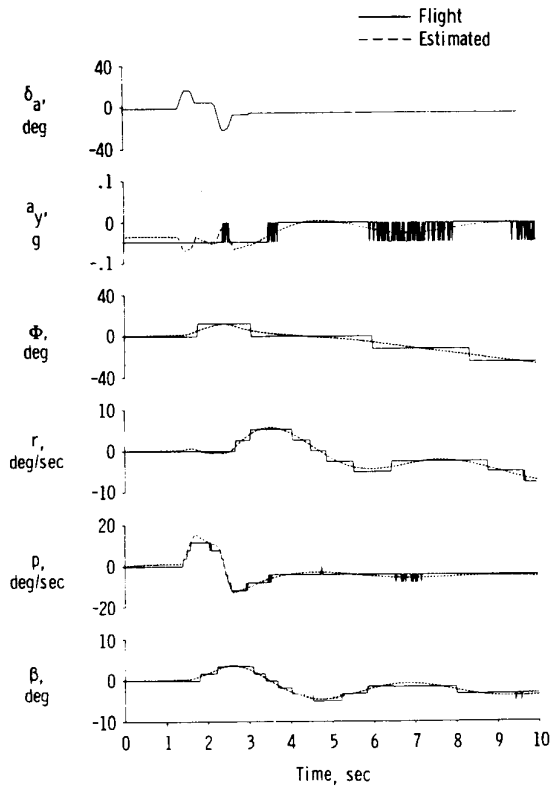

Figure 7.5-2. Resolution reduced by a factor of 16.

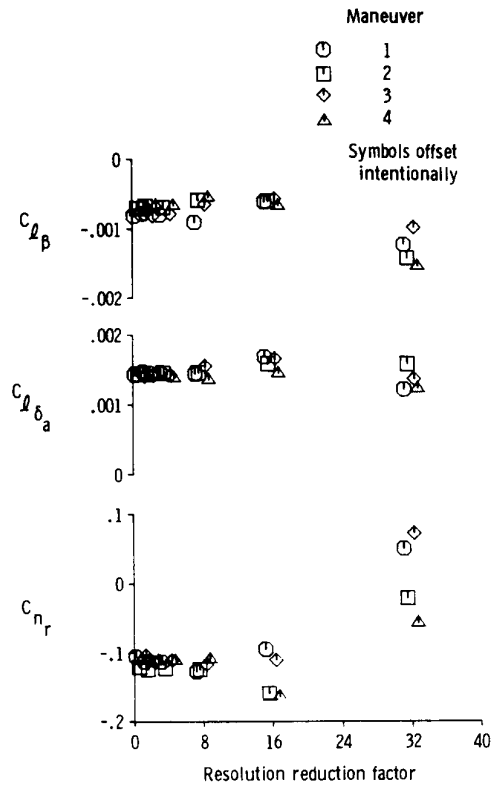

Figure 7.5-3. Derivative estimates as a function of resolution.

8.0 INSTRUMENTATION

In this chapter we examine the instrumentation transducers. We list the signals useful for aircraft stability and control analysis and discuss the uses and relative importance of the signals. The signals range in importance from virtually indispensable to convenient for record keeping. We discuss the kinds of transducers used for each signal and the general characteristics of the transducers.

The virtually indispensable signals are time, control positions, angular rates, and linear accelerations. Although it is theoretically possible to obtain estimates without some of these signals, we consider these to be a minimum practical set of measurements (a_x is optional in some cases). Section 7.2 discusses the measurement of time; problems related to time tagging are quite different from problems with the other measurements, which this chapter covers.

Other useful signals discussed in this chapter are flow angles, Euler angles, angular accelerations, air data, engine parameters, configuration parameters, and fuel and loading data. This list is roughly in order of decreasing importance, although there is variation from case to case; you really should have measurements of flow angles and Euler angles, whereas configuration and loading data are just for bookkeeping convenience.

8.1 GENERAL TRANSDUCER CHARACTERISTICS

This section discusses some broad classes of transducer characteristics pertinent to almost all transducers. Other characteristics may also be pertinent to specific transducers.

Transducer range, resolution, accuracy, and dynamic response are basic issues. The transducer itself has a rated range and resolution, which is often different from the range and resolution of the digitized signal. Of course, the digitized signal can never improve on the transducer characteristics, but it is common for the digitized signal to have a smaller range or poorer resolution than the transducer.

Accuracy is defined in so many different ways that we seldom find accuracy specifications useful. In any event, you need more details about the characteristics of the transducer errors than just a single number bounding the amplitude. Regard accuracy specifications as little more than a summary of detailed data. For many of the signals used in aircraft stability and control analysis, resolution is more important than absolute accuracy; do not confuse the two. It is important only to accurately measure the change in some signals. For other signals, unbiased values are critical (Marchand, 1974).

Linearity is a desirable characteristic in any transducer. Many transducers are reasonably linear except at the extremes of their ranges; thus, you can get smaller deviations from linearity if you do not need the full range. Linearity (although helpful) is not an absolute requirement; it just makes it easier to calibrate the instrumentation and process the data.

Repeatability, the characteristic of always giving the same output for identical inputs at different times, is a critical transducer specification. Two broad categories of repeatability problems are environmental dependencies and drift (however, we make no claim that these categories include all repeatability problems). The magnitudes of the transducer errors due to problems in either of these categories depend largely on how much effort you are willing to invest.

Environmental factors, such as temperature, can cause changes in transducer output. Other environmental factors likely to affect transducers include pressure, vibration, acceleration, magnetic fields, and power supply voltage and frequency; the following discussion of temperature effects applies equally to these other environmental factors. Of course, if the transducer is supposed to measure temperature, then temperature effects do not constitute repeatability problems.

You can minimize environmental dependence problems by controlling the environment. It is particularly common to put sensors and electronics in temperature-controlled enclosures. Vibration effects can be minimized by an appropriate choice of mounting locations and methods. The best approach to environmental dependence problems is a combination of environmental control and the use of high-quality sensors having minimal environmental dependence.

You can lessen repeatability problems due to temperature by measuring, or otherwise deducing, the local temperature and correcting the data for the temperature effects. For this approach to work at all, you need specific, quantitive, verifiable data on the temperature effects. This involves extensive and costly testing; better quality transducers or temperature-controlled environments are usually cheaper.

We generally recommend against environmental corrections for stability and control purposes because we have seldom (if ever) seen acceptable data on which to base such corrections to flight-control-grade sensors (inertial-grade sensors have more consistent characteristics). These corrections require considerable time and effort, which are justifiable only if they make a definite improvement in the derivative estimates. Note in this regard that the largest temperature effect is often a bias, which is only of minor importance for much of our data.

There seems to be a common philosophy that the more correction terms you add to the data, the better results you will get. Our experience is diametrically opposed to this philosophy; we emphasize the virtues of simplicity. Findlay (1981) shows that a major effort in making corrections to the space shuttle data appeared to make the data worse instead of better. This is the almost inevitable result of using correction terms (Bendix Corp., 1980) for which the uncertainty is larger than the value. These correction terms, normally necessary for inertial navigation, were inappropriate for application to these

flight-control-grade sensors. If you cannot convincingly demonstrate that a given correction will improve the data, then do not make the correction.

We use the term "drift" to label changes not directly related to environmental factors. Nearly neutral stability in the sensor can cause drift; the directional gyro is a classic case of drift. Long-term drift can also result from aging of the sensor components.

You can minimize errors from drift by recalibrating the system regularly. Slow drift from sensor aging over months and years is easy to account for and may even be irrelevant to a short test program. Sensors that drift significantly during a flight are troublesome. You need a procedure to calibrate them in flight; you need to include this procedure in the flight planning; and you will likely need to figure out how to handle data that are not immediately preceded by a recalibration. For instance, depending on the stability of the gyros and the duration of the test flights, you might need to erect the gyros several times in a flight, and you should always allow for the possibility of small bias errors.

Crosstalk is the response of a transducer or the data system to signals in addition to the intended signal. We use the term in a broad sense to include any phenomenon having the same end effect, be the phenomenon electrical crosstalk, mechanical sensor cross-axis sensitivity, misalignment, or something else. A simple example of crosstalk is a pitch rate gyro that is slightly misaligned and therefore responds also to roll and yaw rates. Even a perfectly aligned gyro has some degree of crosstalk inherent in its design. You can correct the data for crosstalk if you know the amount present. Treat such corrections with the same precautions as the previously mentioned environmental corrections; it is better not to correct for the crosstalk than to apply the corrections with questionable data.

Many important transducer characteristics are not evident in static calibrations. The most basic dynamic specifications are in terms of the transducer frequency response, the Bode plot of the transducer gain and lag as a function of frequency. The bandwidth required for stability and control data is low enough that transducer gain characteristics are seldom a problem; the only cases where we recall needing to account for transducer gain as a function of frequency involved pressure sensors. Transducer lag is more often important because of the sensitivity of the derivative estimates to small lags. A constant time lag (linear phase lag) is an adequate approximation to the frequency response of most transducers in the frequency ranges pertinent to aircraft stability and control analysis.

Other transducer characteristics include dead bands, hysteresis loops, "stickiness," and rate limits. You can seldom do much to correct data for these problems, even if you have quantitative information about them. The main thing to do is to collect appropriate data to ensure that these problems are small enough not to cause difficulties; otherwise you need better transducers.

Some transducers have characteristic glitches passing through specific signal levels. The solution to such glitches is usually to smooth through them with a spike remover. This can be a laborious process if there are many glitches because simple automated spike-detection algorithms are not adequate for the shapes of some glitches (they can have slow onsets, unlike spikes).

Finally, any transducer has a certain random noise level. The noise statistics are useful to know, but we caution against using them as a measure of the noise in the system model. The "noise" in our system models is a combination of many things, including modeling error and the transducer errors mentioned previously; the actual transducer noise is typically a small portion of the total.

8.2 CONTROL POSITIONS

You must have measurements of any control surfaces or other control devices that move during a maneuver. The terminology of our discussion emphasizes control surfaces, which constitute the large majority of control devices, but the principles apply in general. Measurements of control surfaces that remain unchanged during a maneuver are useful, if only to verify that the surface remained fixed, but are not required. Some control surfaces may also count as configuration parameters (Section 8.9), which are useful to record for bookkeeping purposes.

Derivative estimates are extremely sensitive to time-tagging errors in the control position measurements, so your selection of measurement strategies should strongly consider possible timing errors.

The best control position measurements are direct measurements of the surface hinge angles because the less well defined the relationship between the measurement and the actual surface position, the more work you need to do to evaluate the surface position and the more error you are likely to have. Actuator push-rod positions are often reasonable measurements, although not as good as direct surface measurements; the surface position is a nonlinear function of actuator push-rod position, and there is a finite amount of play in the rod attachment.

Actuator command measurements are one step further removed from the actual surface position and thus one step poorer. To use such measurements, you must model the actuator response. This response always has some lag, which should not be ignored.

Control stick (or wheel or pedal) position measurements are the worst alternative, usually suitable only for rough work. Even in the simplest mechanical system, there are dead bands, hysteresis loops, backlash, cable stretch, and nonlinearities in the connection of the control stick to the surface (Foster, 1977). It is impractical to accurately model all these effects; they just constitute error sources. For a more complicated airplane, you must simulate a complete control system, including filters, limits, feedback, sampling effects, actuators, flight condition schedules, and other logic, in order to determine the surface position from the control stick motions. Direct control surface position measurements, or something closely approximating direct measurements, are the key to keeping such control system models out of stability and control derivative estimation. This in turn is one of the simplifications that keeps our system models tractable. If you need to estimate control system parameters, the control stick positions are useful, but we try to avoid this complication, as discussed in Chapter 1.

Control position transducers (CPTs) are simple, reliable, mechanical devices. They have negligible lag and few problems of any kind. Only once have we had strong suspicion of a problem with a control position transducer, and that suspicion is unverified.

8.3 ANGULAR RATES

The angular rates are the most important response variable measurements. You can conceivably get reasonable estimates of moment derivatives with only control surface and angular rate measurements. Conversely, the loss of even a single angular rate measurement is a serious problem that significantly diminishes the odds of getting acceptable results. Angular rate measurements (or reconstructions of such measurements) are mandatory for equation-error estimation methods. Output-error and filter-error methods can in principle work without angular rate measurements, but expect severely degraded results.

Fortunately, rate gyros, the transducers used to measure angular rates, produce excellent-quality signals for stability and control analysis. A typical instrumentation-quality single-axis rate gyro is an inexpensive, lightweight, self-contained unit in an approximately 1-in cube. The gyros have good frequency response, nearly linear calibrations, good repeatability, and virtually imperceptible noise level.

Rate gyros are also quite reliable. We have learned to take rate gyro data seriously, even when it seems anomalous at first glance. Figure 8.3-1 shows a time history, from a scaled F-15 model (Iliff et al., 1976), that we were convinced was conclusive evidence of a rate gyro failure. The sharp peaks in the measured roll rate signal at 1 and 3.5 sec look unreasonable. They imply extremely large rolling moments with no apparent cause. There is no evidence of such a large event in the other measurements; the figure shows a reasonable fit with a linear aerodynamic model. Armed with this evidence, we requested an examination of the roll rate gyro. When static calibration tests looked good, we requested extensive dynamic testing. The dynamic tests also showed excellent gyro behavior.

Still convinced that something was wrong with the gyro, we tried to gather more evidence by differentiating the Euler angle measurements and using Equation (2.2-14) to get an independently derived roll rate signal. Although we were unsure whether the Euler angle measurements were good enough for this task, it was worth trying. To our great embarrassment (considering the effort spent checking our assertion that something must have been wrong with the rate gyro) this independently derived roll rate clearly verified the rate gyro measurement. The rate gyro was good. (We later found that powerful vortices were shed from the forebody at the flight condition of this maneuver. These vortices impinging on the vertical tail caused sudden large moments far in excess of what would otherwise be considered reasonable magnitudes.)

The only significant problem with rate gyro measurements is alignment. Section 5.2.1 shows the sensitivity of some derivative estimates, particularly $C_{n_{\delta a}}$ and $C_{\ell_{\delta r}}$, to errors in the measurement of I_{xz}. Errors in gyro alignment have closely related effects. Suppose a maneuver has a peak roll rate of 50 deg/sec and a peak yaw rate of 5 deg/sec (reasonable values for an aileron pulse on a fighter), and suppose that the yaw rate gyro is misaligned by 1° in the roll-sensitive direction. The yaw rate gyro will sense about 2 percent of the roll motion. The misalignment thus causes yaw rate measurement errors of up to 1 deg/sec or 20 percent of the peak yaw rate; this error level is completely unacceptable.

Some data from the first space shuttle entry were even more sensitive to alignment errors. We attempted to analyze tiny longitudinal motions (see Section 6.2.2), with peak pitch rates of about 0.1 deg/sec, during bank reversals with yaw and roll rates of about 5 deg/sec; in these circumstances, a 1° misalignment in the pitch rate gyro would generate errors of 100 percent in the sensed pitch rate. We generally obtained results of marginal quality from these maneuvers; misalignment may have been one of the contributing factors. We tried to estimate the misalignment and obtained some inconclusive results suggesting alignment errors as large as 0.5°.

Our standard criterion is to request alignment within about 0.1° (0.2 percent crosstalk), which is probably near the practical limit. In-flight airframe deformations make it difficult to define alignment with greater precision, regardless of the precision of the equipment. Since even a perfectly aligned gyro has some degree of crosstalk inherent in its design, it is useful to include crosstalk tests of the assembled gyro package in your calibration procedure.

Although the alignment of rate gyros is important, their position in the aircraft is not. The angular rates of a rigid body are independent of the reference point. The main consideration in selecting gyro location is a good rigid attachment point.

In inertial navigation units, rate gyros are routinely corrected for effects like g-sensitive bias, dynamic misalignment, g-sensitive drift, anisoelastic drift, and temperature. For stability and control analysis, we have found these corrections to be imperceptible (usually less than the resolution); they are more likely to cause problems through misapplication than to help. Stability and control analysis requirements are drastically different than inertial navigation requirements; even the best rate gyros we have used are almost useless for inertial navigation.

Strap-down laser gyros show promise of being useful both for inertial navigation and for stability and control purposes. As yet we have no first-hand experience with laser gyros. Second-hand reports and specification sheets are encouraging. Laser gyros are quite expensive when compared with flight-control-grade rate gyros, but not when compared with inertial system costs.

8.4 LINEAR ACCELERATIONS

We consider linear accelerometer data to be mandatory for acceptable stability and control derivative estimates. You can sometimes get by without a_x, provided that you do not need C_A derivative estimates, but a_n and a_y are crucial. Virtually all the force derivative information is in the accelerometer data, and the accelerometer data also help the moment derivative estimates.

We loosely refer to linear accelerometers as measuring acceleration, but this is not strictly true; accelerometers ignore the component of acceleration due to gravity. See Section 2.1.6 for further discussion.

Acceleration is a second derivative quantity. Because derivative operations amplify high-frequency signals, accelerometer measurements intrinsically contain significant high-frequency structural data. For rigid-body stability and control analysis, these structural data constitute unwanted noise. The solution to the structural noise problem is in two parts: First, you must select a good attachment point and method for the accelerometer package. The attachment point should be structurally strong and rigid and should be in an area of the vehicle where the major structural modes have small amplitudes. The best locations are usually on major structural members near the center of gravity; wingtips are probably the worst locations. (Stability and control accelerometer requirements are drastically different from those of structural dynamics in this and other factors; the odds of meeting both sets of requirements with the same accelerometer are negligible, so do not even try). The attachment should be rigid, with no free play. Do not use shock isolation attachments as a way of minimizing structural noise; this may be a good idea for some applications, but not for ours. Shock isolation attachments introduce poorly quantified lags.

Second, you must filter out the structural noise. Presample filters are virtually mandatory. You often get the best results by sampling at a relatively high rate (with, of course, a presample filter), digitally filtering the high-rate data, and then thinning to the sample rate chosen for analysis. Sections 7.3 and 7.4 discuss these ideas. The shock isolation attachments discussed previously amount to mechanical filters; digital or analog electronic filters have more precisely measurable characteristics and can be better tuned to your specific requirements. Whatever filtering method you use, be sure to account for the filter lags in your time tagging.

Instrumentation-quality linear accelerometers are small units, about the same size and weight as rate gyros. They are reliable, and the signal quality is generally good for stability and control analysis, provided that you take adequate care in mounting and filtering.

Like rate gyros, accelerometers are subject to alignment errors. It is often advisable to mount the gyros and accelerometers on a common baseplate so that you can run combined alignment tests on the assembly and so that you have only a single assembly to align relative to the airplane.

Unlike rate gyros, accelerometers are sensitive to mounting location. The ideal location is at the airplane's center of gravity, but that would be an unreasonable requirement, particularly when the center-of-gravity location changes during flight. Your primary consideration in selecting mounting locations should be using rigid attachment points where the major structural mode amplitudes are small. Mounting locations far from the center of gravity are prone to bending problems. Good mounting locations will often be near the center of gravity, but this is not an intrinsic part of the requirements. Either you can include the accelerometer position in your measurement model or you can correct the accelerometer measurements to the center of gravity. Section 3.5 discusses these alternatives. For either approach, you need accurate measurements of the positions of the accelerometers relative to the center of gravity. Section 5.2.3 discusses the problems of erroneous accelerometer position data.

The comments in Section 8.3 about the inadvisability of using miniscule correction terms are also applicable to accelerometers. Instrumentation-quality accelerometers have excellent linearity; a simple linear calibration, with a bias and a slope, gives data as accurate as typical data system resolution allows. Static alignment corrections are reasonable if you have precise data to base them on. Either correcting or modeling the transducer positions is mandatory. We generally advise against any other

corrections; if other corrections are significant to stability and control analysis, you probably need better quality accelerometers.

8.5 FLOW ANGLES

We highly recommend that you obtain measurements of angle of attack and angle of sideslip, collectively called flow angles. Flow angles are not on our list of mandatory measurements, and we have occasionally obtained good results without using flow-angle measurements. Nonetheless, flow-angle measurements are of major benefit in obtaining good derivative estimates and should be part of any serious stability and control tests. Without flow-angle measurements, you are largely restricted to small-perturbation maneuvers from well-stabilized flight; your analysis task is also more difficult. Flow-angle measurements, or reconstructions of such measurements, are mandatory for equation-error estimation methods.

8.5.1 Uses

We describe six different uses for angle-of-attack measurements, citing the uses roughly in order of importance. Only the first of these uses is relevant to angle of sideslip; the remainder apply only to angle of attack.

The first use, and probably the most important, is as observations. Flow-angle observations greatly contribute to obtaining good α and β derivative estimates; this should be no surprise. (The surprising fact is that they are not mandatory for output-error or filter-error methods.) Accurate estimates of α and β derivatives indirectly help in the estimation of the other derivatives.

The second use of flow-angle measurements, primarily applicable to angle of attack, is in the evaluation of flight condition. Angle of attack is an important flight condition parameter. Predictions are commonly expressed as functions of angle of attack. You need to know the angle of attack to select the appropriate predicted values both for starting and for comparing with the derivative estimates from flight. Even if no predictions are available, you will want to present your estimates as functions of angle of attack. Results presented without any indication of the flight condition to which they apply are virtually useless. If you do not have an angle-of-attack measurement or reconstruction, you have two other options: You can compute angle of attack from other data such as pitch attitude and flightpath angle in steady flight, or you can plot data as a function of something related, such as lift coefficient.

The third use of angle-of-attack measurements is in the kinematic term ($p \sin \alpha - r \cos \alpha$) in the β state equation. This term is relevant whether you have a β observation or not. The importance of the α measurement to this term is a debatable issue. The lateral dynamics are quite sensitive to the value of α used in the ($p \sin \alpha$) part of this term; this implies that an accurate measurement of α is crucial. Our experience is that this term is so sensitive that we cannot rely on the α measurement being accurate enough. Therefore, we normally estimate a bias in the ($\sin \alpha$) term as one of the unknowns in each lateral-directional case. This means we do not really need a measured α for this purpose.

Estimating this bias is, in fact, one way to estimate α without a direct measurement. Figure 8.5-1 shows estimates of $\sin \alpha$ obtained by this means from a 3/8-scale F-15 model (Iliff et al., 1976). The estimates are plotted as functions of measured angle of attack over a range of -20° to 55°. The solid line is the curve of perfect agreement. This figure shows that the $\sin \alpha$ estimates from the β state equation are adequate to give a rough approximation of the angle of attack, but there is a scatter band of several degrees. For this vehicle, the scatter is larger outside the range of -15° to +30°.

If the angle of attack varies significantly during a lateral-directional maneuver, you need a measured α to track the time variation, although you can allow a constant unknown bias, as shown in Equation (3.8-13a).

The fourth use of angle-of-attack measurements is in the transformation of the C_N and C_A derivatives into the C_L derivatives used in the α state equation. The measurement does not have to be particularly accurate for this purpose, so rough approximations based on other data can often substitute adequately for an α measurement. For cruise conditions, zero is often an adequate approximation for the angle of attack; this approximation simplifies the equations significantly (see Section 3.7). The transformation is most important at high angles of attack and low speeds.

The fifth use of angle-of-attack measurements is in the gravity term in the α state equation. The reason you need an α measurement here is that the term is nonlinear. For a nonlinear term, you cannot blithely subtract the initial condition and work only with perturbation quantities. Section 3.7 describes methods for dealing with nonlinearities. In general, if the maneuver is large, you will need point-by-point measurements or reconstructions of α. For small maneuvers, you can get by with average values of angle of attack, perhaps approximated from other data.

The sixth use of angle-of-attack measurements is in the thrust term in the α state equation. The same comments apply to this nonlinear term as to the gravity term.

8.5.2 Vanes

There are several methods for measuring flow angles. No single method is clearly optimum in all situations. All the methods have some shortcomings; do not expect flow-angle measurements to be of the same high quality as surface position, angular rate, or linear acceleration measurements. It requires work, both in transducer installation and in analysis, to compensate for the inherent measurement problems and get good flow-angle data.

Most test programs at Ames-Dryden obtain flow-angle measurements from boom-mounted vanes. The boom is most commonly attached to the aircraft nose, but we have occasionally used booms attached to wingtips or other places. We consider boom-mounted vanes to be generally the best means of measuring flow angles, provided that you can put an adequate boom on the aircraft without influencing the aerodynamics.

Ideally, the boom should be long enough to position the vanes in the free stream, completely outside the area of measurable flow patterns around the aircraft. This ideal is unrealistic. A more practical goal is to position the vanes far enough from the aircraft that the flow influence of the aircraft is small enough and simple enough in form to correct for. Our empirical rule for nose booms is that boom length should be 2.5 to 3 times the fuselage diameter.

Boom-mounted vanes are invariably far from the aircraft center of gravity. Therefore, the observation terms involving the offset of the sensor from the center of gravity are important. You must know the vane locations to evaluate these terms. Either you can correct the measured data to the center of gravity or you can include the position effects in the observation equations; Section 3.5 discusses this choice.

Flow-angle vanes require considerable calibration to determine the effect of the flow field around the aircraft. You need to calculate this effect in order to subtract it and leave the desired free-stream flow angles (or in order to include it in the observation equation). The calibration data can come from theoretical computations, wind-tunnel tests, or flight tests. The effects get smaller as the vanes move farther from the aircraft, but it requires unreasonable boom lengths to make the effects negligible. For moderate flow angles, a reasonable model of the effects is

$$\alpha_\ell = K_\alpha \alpha_f + \alpha_b \tag{8.5-1}$$

where α_ℓ is the local angle of attack of the flow stream passing by the vane, α_f the free-stream angle of attack at the vane location, K_α the upwash factor, and α_b the bias. Wolowicz (1966) discusses calibration of flow-angle vanes. Section 3.6.3 discusses some aspects of estimating K_α from flight maneuvers. We do not go into detail here, but note only that accurate vane calibration can require considerable flight and analysis time.

Other possible complications of boom-mounted vanes include boom and vane dynamics. Although boom dynamics are often detectable as in Figure 7.3-9, they seldom constitute serious problems if the boom is designed with reasonable strength. The dynamic modes of a well-designed vane are imperceptible in flight. Static effects such as vane warping (with balsa vanes) or bending (with metal vanes) cause large problems. Careful handling can prevent bending of metal vanes. Humidity-induced warping of balsa vanes is more difficult to avoid. Boom bending can be a problem at high load factors. Finally, vanes are subject to a variety of mechanical difficulties such as friction and grit in the bearings. Some analysts apply data corrections for problems such as boom bending. We recommend a cautious, conservative approach to such corrections; carelessly applied corrections sometimes cause more problems than they solve, particularly when the original problems are negligible.

In spite of these problems, we generally consider boom-mounted vanes to be the best source of flow-angle measurements. We prefer this source unless boom installation is ruled out. The cost and complexity of installing an adequate boom is prohibitive in some test programs. There are also situations where booms are unacceptable, regardless of cost. For instance, a nose boom can have a large effect on the aerodynamics, particularly vortex flow at high angles of attack.

Body-mounted vanes avoid the necessity of installing booms to measure flow angles. Such vanes are typically mounted on both sides of the forward fuselage or nose. Although body-mounted vanes avoid the problems associated with booms, they introduce a serious problem of their own. The flow patterns near the aircraft surface, where the body-mounted vanes are located, are extremely complicated. For instance, nonlinear effects of yaw rate and roll rate are plausible. Averaging vane measurements on the left and right sides will not eliminate nonlinear effects. It is essentially impossible to accurately account for all such effects, so you should anticipate measurements from body-mounted vanes having errors correlated with the other states, regardless of how much effort you spend on calibrating the vanes. The purpose of a boom is to reduce the size and complexity of the flow effects to a manageable level. Body-mounted vanes can give adequate flow-angle measurements for some purposes, particularly in steady flight, where there are fewer influences to consider. You can use body-mounted vanes in analyzing dynamic stability and control maneuvers, but recognize that the signal quality will probably be marginal and watch for potential problems from the correlated errors.

8.5.3 Pressure Ports

Pressure ports are an alternative to vanes for measuring flow angles (Wuest, 1980). The pressure patterns over the vehicle are strongly affected by flow angles. Thus, if you know these functions, you can deduce the flow angles from measurements of the pressure at several locations. (Three locations is a minimum to deduce the flow angle in one axis.) Siemers et al. (1983) and Foster (1980) describe specifics of this technique for measuring flow angles. We comment here only on some of the general characteristics of the method.

You must know the pressure patterns corresponding to different flow angles in order to deduce the flow angles from pressure measurements. This is an impractical requirement if the pressure port locations are arbitrary. You need locations specifically selected to make the job easier.

The space shuttle provides an example of two typical arrangements. Figure 8.5-2 shows the shuttle side probes. There are two of these side probes, one on each side. Each probe has three pressure ports, one each on the center, top, and bottom of the front of the probe, plus two static pressure ports farther back. The side probes are retracted during the hypersonic part of the entry, when they would burn off if extended. They extend at about Mach 3.5, a higher Mach than would be practical for vanes. In other vehicles, such probes could be boom mounted.

Figure 8.5-3 shows the shuttle entry air data system (SEADS) ports, planned for installation on future shuttle flights. The ports in this system will be built into the nose of the vehicle itself, allowing data gathering throughout the entire entry.

Pressure port measurements of flow angles have two major problems. First is the difficulty of calibration. No matter how well positioned the ports are, the relationship between flow angles and pressures is complicated. Accurate calibration of pressure port measurements of flow angles requires extensive work, combining theoretical analysis, wind-tunnel tests, and flight tests.

The second problem is the lag inherent in pressure sensors. Section 8.7 discusses such pressure lags. The predominant frequencies of the flow-angle signals in dynamic maneuvers are higher than those of the air-data signals, so lags in the flow-angle measurements are more critical. The pressure lag problem is further complicated for flow-angle measurements because each of the several pressure ports involved can conceivably have a different lag.

It is possible to overcome the difficulties of pressure measurement of flow angles, but it requires a lot of work, often a prohibitive amount. Pressure ports are seldom a viable option on small projects; even on large projects, you will mostly consider pressure ports when they are already installed for other purposes or when the use of vanes is precluded.

8.5.4 Other Sources

Another source of flow-angle information is inertial integration. By this term we mean any process of estimating the aircraft velocity vector by integrating accelerometer measurements. This could involve a specialized inertial navigation unit; at the other extreme, it could be merely integrating linearized equations of motion using the stability and control accelerometers and gyros as inputs.

Inertial integration using the stability and control instrumentation does not add new information. It is really a reconstruction rather than an independent measurement. It is reasonable to treat flow angles from inertial navigation units as independent measurements because the inertial navigation unit uses sensors different from the stability and control sensors; furthermore, the sensors have dramatically different characteristics.

There are three problems with using inertial integration to obtain flow-angle time histories for stability and control analysis. The first problem is drift in the open-loop integration of the neutrally stable equations. The drift in integrating typical instrumentation-quality accelerometer measurements can be large, even during a single maneuver, unless you use special tricks to reduce the drift (like adjusting biases to match reasonable end conditions). The drift in any inertial navigation unit used for more than spare parts will be too small to measure during a single maneuver.

The second problem is one of response to dynamic maneuvers. Inertial navigation units are not designed to track the relatively high-frequency motion of a dynamic maneuver; the navigation problem concentrates in a much lower frequency range. Inertial navigation data tend to be heavily low-pass filtered, in some cases by the intrinsic characteristics of the transducer. Thus, although the steady-state data before and after a maneuver are extremely accurate, the inertial navigation unit may be of little value in tracking the dynamic response of a stability and control maneuver. It is possible to modify some inertial navigation units to obtain higher-frequency data, alleviating this problem. The problem does not apply to inertial integration using the stability and control sensors. Strap-down inertial systems typically have better dynamic response than do gimbaled inertial systems.

The third problem is intrinsic to all inertial integration methods: In the pure form, these methods ignore the effects of winds. Figure 8.5-4 shows sideslip angle from the space shuttle inertial naviga-

tion unit during the second entry. The large motions near the end of the entry are errors from neglecting winds; the true sideslip angle remained near zero.

There are several methods of accounting for winds in inertial navigation, but none really solve the problem. You can subtract known winds measured by balloons or other techniques, but such measurements can give only large-scale general patterns. No external sensor can accurately measure the winds encountered by the aircraft, particularly the small-scale phenomena.

Pressure ports, vanes, or other flow-angle and air-data measurements can be used to deduce winds by comparing the inertial integration flow angles and velocities with the measured flow angles and air data. This is a reasonable way to incorporate wind data into an inertial navigation system; the space shuttle system uses such an idea to adjust inertial velocity and angle of attack (but not angle of sideslip) after the side probes extend. This does not much help our quest for a means of inertially measuring flow angles, however; the inertial system here uses flow-angle measurements, it does not create them.

The final method of measuring flow angles is to deduce the flow angle required to get the measured accelerations. This method assumes knowledge of the force coefficient functions. After you have verified the force coefficient functions in flight test, this may be an excellent method of estimating flow angles on production vehicles, where you may want to remove the boom used in testing. For flight test use, where you are trying to determine force coefficients, it is circular to use flow-angle measurements based on knowledge of the force coefficients. You use such self-perpetrated data only at the severe risk of obtaining misleading results.

8.6 EULER ANGLES

We highly recommend that you obtain measurements of the Euler angles Θ and Φ unless your test program is completely restricted to small-perturbation maneuvers from steady wings-level flight. The Euler angle Ψ, on the other hand, is of little value in estimating stability and control derivatives.

The most important use of Euler angle measurements is in the gravity term in the α and β state equations. The Euler angle Ψ does not enter into these terms, which is why measurement of Ψ is less important than measurement of Θ and Φ. For small-perturbation maneuvers from wings-level flight, you can get by with small-angle approximations that do not require Euler angle measurements. For any other conditions, you will need measurements of Θ and Φ to get adequate results.

This is one place where equation-error methods can get by with less data than output-error methods. Since equation-error methods do not involve integration of the state equations, they do not directly need the Euler angles in this role. (They might, however, use Euler angles for data reconstruction and checking.)

Euler angles are also used as observations. Their value in this role is minimal. They add little to the angular rate data, which are generally of much higher quality. This role applies to all three Euler angles, but Ψ adds so little information that we seldom bother to use it, even if measurements are available.

Euler angles are usually measured by platform attitude gyros. Relative to other stability and control instrumentation, attitude gyros are bulky and use considerable power. They are also less reliable than accelerometers and rate gyros. The biggest problem with attitude gyros is drift. The gyros will typically drift significantly during flight. This means that the gyros must regularly be reset (often referred to as erecting the gyros). It also means that small bias errors, varying from maneuver to maneuver, are to be expected. Neither of these effects are disabling, but you should be aware of them so that you do not expect more from the attitude gyros than is realistic.

You should also be aware of the gyro erection procedure on your airplane and the possible effects of this procedure. Some aircraft have erection circuits that amount to high-pass filters. This means that the Euler angle data might be distorted in long, slow maneuvers if the erect circuit is active and the break frequency is high enough. You should know at least enough about the gyro erection in your airplane to determine whether this is a potential problem.

Another potential problem with most attitude gyros is gimbal lock. This is seldom a problem for stability and control maneuvers. Gimbal lock usually occurs only when the aircraft nose is pointed nearly straight up or down. These are rare (though not completely unheard of) conditions for stability and control maneuvers. We have encountered gimbal lock during tests of fighter airplanes.

The resolution of Euler angle measurements is often poor. This is not because of any limitations of the attitude gyros, but because of the large signal range and the small fraction of this range used in typical maneuvers. Section 7.5 gives some examples of Euler angle measurement resolutions. The poor resolution of the Euler angles is seldom a serious problem. None of the uses of the Euler angles require high resolution; the Euler angles are among the least resolution-critical signals. For instance, although we prefer better, we can usually accept Euler angles with the 1.4° resolution implied by scaling an 8-bit PCM channel over a full 360° range.

Strap-down inertial systems can obtain Euler angle data without using gimbaled platform gyros. Although we have had no experience in using strap-down inertial systems, they should be superior to gimbaled gyros.

8.7 AIR DATA

The term "air data" refers to the properties of the air in the vicinity of the airplane and the relative motion of the airplane and the air. The air-data parameters of most interest to us are angles of attack and sideslip, air-relative velocity, dynamic pressure, and Mach number. Related parameters such as Reynolds number are sometimes pertinent. The measurement of the flow angles presents somewhat different issues than the other air-data parameters; flow-angle measurement is separately discussed in Section 8.5.

The air-relative velocity, dynamic pressure, and Mach number are closely related by the equation

$$\bar{q} = \frac{1}{2}\rho V^2 = 0.7P_sM^2 \tag{8.7-1}$$

This section concentrates on these three parameters.

Dynamic pressure and velocity are used in the nondimensionalization of the stability and control derivatives. Therefore some form of dynamic pressure and velocity data is mandatory for most applications. Section 5.2.2 discusses the importance of nondimensionalization. Errors in dynamic pressure and velocity measurement cause proportional errors in the nondimensional derivative estimates. If your application requires only dimensional derivative estimates and if the dynamic pressure and velocity are constant during each maneuver, then you can dispense with dynamic pressure and velocity measurements, but this situation is rare.

The dynamic pressure measurement is particularly critical for normal force derivatives. This arises from the relatively large normal force in level cruise flight (equal to the airplane weight). Suppose that dynamic pressure changed by 10 percent during a maneuver where the normal acceleration varied between 0.9 and 1.1 g. The dynamic pressure change would directly cause a 0.1-g change in the normal acceleration. Therefore, if dynamic pressure measurements were unavailable, you would have modeling errors on the order of 50 percent of the perturbation in the normal force. Dynamic pressure effects on normal acceleration are amplified because the dynamic pressure effects are proportional to the total normal force, which is much larger than the perturbation.

Figure 8.7-1 shows a typical example of the problems that arise from neglecting the changes of dynamic pressure during a maneuver. Although dynamic pressure measurements were available, this analysis used only the constant average value (use of a constant value often divides computer time by a factor of 2 or 3). Note the fit error in the low-frequency range in normal acceleration and the correlation of this error with dynamic pressure.

A secondary use for velocity measurements is in the gravity term of the α and β state equations. Precise measured values are of less importance in these terms. On occasion we have even estimated these terms as unknowns, usually as a diagnostic tool to check the velocity measurement.

Mach number measurements are necessary to define the flight conditions of transonic and supersonic airplanes. The Mach number is needed for looking up starting values and for comparing with other estimates. Constant Mach number is a requirement for usable transonic data; a Mach number measurement is needed to evaluate how well each maneuver adheres to this model. The Mach number is not used in the estimation process itself.

The air-data parameters are usually derived from three basic measurements: static pressure, total pressure, and temperature. A direct measurement of differential pressure (total minus static) often substitutes for the total pressure measurement. Gracey (1980) and Wuest (1980) give comprehensive treatments of these air-data measurements.

The total pressure (or differential pressure) measurement is the most important; you should make every effort to record a total or differential pressure measurement on a data tape in order to allow modeling of the effects of dynamic pressure changes during a maneuver. The pilot's airspeed gauge really measures differential pressure, so lap notes of indicated airspeed are a potential alternative form of these data. Although we have used such lap notes on occasion, we recommend them only as a last resort. The use of lap note airspeed data forces you to assume constant dynamic pressure during a maneuver and provides no good way to check the validity of the assumption. The maneuver requirements are therefore considerably more stringent. Lap notes are also less accurate and more susceptible to error than automatic data systems.

There are several acceptable alternatives for static pressure and temperature data. The preferred approach is to record static pressure and air temperature probe measurements on the data tape. Static pressure and temperature are essentially constant during most (but not all) reasonable stability and

control maneuvers. Therefore, lap note recording of pressure altitude and outside air temperature gauges is often acceptable. For stability and control analysis, you can probably even get by with approximations, like using ground-level temperature minus a standard lapse rate to extrapolate to temperature at altitude.

Calibration of air-data sensors is too broad a subject to reasonably cover here. This calibration can be a major consideration in a flight test program. Gracey (1980), Herrington et al. (1966), and Wolowicz (1966) discuss air-data sensor calibration in more detail.

Pressure lags in air-data sensors are often large. Lags are commonly on the order of several tenths of a second, and are sometimes longer than a second. For most stability and control maneuvers, the air-data lags present no problems, because the air-data parameters are nearly constant during the maneuvers. Pressure lags are important, however, in analyzing maneuvers such as that shown in Figure 6.2-2 involving large, rapid changes in static and dynamic pressure.

Figure 8.7-2 shows data from a step response test of the air-data sensors of the airplane that flew the maneuver shown in Figure 6.2-2. The test began with 50 lb/ft^2 of suction applied to the static port. The total pressure port was at ambient conditions in the hangar, giving a 50-lb/ft^2 differential pressure. We then suddenly removed the suction from the static port and recorded the instrument response to this step change in pressure. The static pressure response is plain, showing an approximately exponential rise from the reduced pressure to ambient conditions. The time constant of this rise is about 0.4 sec. The digital resolution is apparent in the stairstep appearance.

The response of the differential pressure sensor requires a little more explanation. The differential pressure transducer has a multiturn potentiometer; any particular digital counts value can correspond to several different engineering units values, depending on which turn the potentiometer is on (you determine the correct turn by examining a coarse measurement). The data in Figure 8.7-2 are in raw, uncalibrated counts. The 50-lb/ft^2 pressure change crosses several complete turns, which is the cause of the sawtooth appearance. The sudden jumps at about 0.5, 1.0, and 1.8 sec are the turn changes. The engineering units decrease monotonically from 50 to 0 lb/ft^2.

Discounting the turn jumps, the response appears to consist of two segments of constant slope, changing from the higher slope to its lower slope at about 1.2 sec. The constant slope before 1.2 sec is from the sensor rate limit. The apparent constant slope after 1.2 sec is an artifact of the highly nonlinear sensor calibration; the engineering units here have a more nearly exponential shape, as we should expect. Finally, the sensor reaches the lower calibration limit of the lowest turn at 1.8 sec and then stays limited there.

Gracey (1980) and Herrington et al. (1966) give data on pressure lags. It is important to note that pressure lags are strongly affected by density; the lags are much larger at altitude than at ground level, up to four times larger for reasonable cruise altitudes. Therefore, be sure to ask what altitude the data apply to when you look at pressure lag figures. A system with acceptably small lags in the hangar may be unusable for the maneuvers you plan at cruise altitudes.

Inertial navigation systems or radar tracking are possible alternative sources of velocity measurements; assuming some source of density and static pressure data, Equation (8.7-1) gives dynamic pressure and Mach number. Wind effects are a major obstacle to these methods. If you know the winds, you can easily account for them, but accurate wind data are difficult to obtain. Accurate static pressure and density data are also a problem at extreme altitudes.

8.8 ANGULAR ACCELERATIONS

Angular acceleration measurements are useful but are not required. The primary use of angular accelerometer measurements is as an observation. You can also use angular accelerometer measurements to correct linear accelerometers to the center of gravity if you choose that approach to dealing with instrument position (see Section 3.5).

In our experience, the quality of angular accelerometers usually varies from marginal to unusable for stability and control analysis. We have seen a few good angular accelerometers, but they are more the exception than the rule. Many angular accelerometers are extremely noisy, to the extent that the response to a stability and control maneuver is lost in the noise. Figure 5.2-4 shows such data from the space shuttle. The noise sensitivity is the worst in buffet and turbulence. Those angular accelerometers that have reasonably low noise levels often have large lags and severe distortion at frequencies of interest for stability and control analysis.

It is fairly common practice, when angular accelerometer data are unavailable or of poor quality, to substitute differentiated angular rates, smoothed as necessary to alleviate noise problems in the differentiation. The merits of this practice depend on the intended use of the data. Regression methods require angular acceleration data, justifying the differentiation of the angular rates if that is the best (or only) source of such data. Differentiated angular rates are acceptable for use in correcting linear accelerometers to the center of gravity, provided that the differentiated signal has a low noise level. However, modeling of linear accelerometer positions in the observation equation is less susceptible to noise and is thus generally preferable to correcting the data using differentiated signals.

We usually advise against using differentiated angular rates as observations for output-error or filter-error methods. Such use seldom causes problems (unless the differentiated data are excessively noisy or are distorted by poorly chosen filtering) but seldom helps. The main effect is to spend extra effort and computer time without measurably changing the resulting derivative estimates. This waste arises from a mistaken impression that just because output-error methods can use angular acceleration observations, such observations are required. Angular acceleration data are required for regression methods but not for output-error methods. We have done more stability and control analysis without angular acceleration data than with them.

In output-error methods, angular rates and angular accelerations are largely mutually redundant. If you have truly independent angular rate and angular acceleration data, they are not redundant and you have more information than from either signal alone. If the angular acceleration signal is derived by differentiating the angular rate signal, it adds no new information; it only repeats the same information in a different form.

This change in form can be useful. The differentiated signal emphasizes the high-frequency content of the data, possibly limited by any filtering done with the differentiation. Whether this emphasis is helpful or detrimental depends on specific circumstances. In many cases, the useful data are all at relatively low frequencies, and the high-frequency content is little more than noise. In other circumstances, the high-frequency part of the data gives good control derivative information. This occurs when there are sharp control inputs and the random noise level and quantization level of the rate gyros are low. If, in addition, the low-frequency response is distorted by drift and small modeling errors (large modeling errors would make the analysis unacceptable), then the high-frequency part of the data can contain the best control derivative information. In such situations, the emphasis change caused by using differentiated rate gyro signals improves the results.

In summary, the use of differentiated rate gyro measurements is sometimes justifiable, but we recommend against indiscriminant use without specific benefits in mind. If you have no useful angular accelerometer measurements and no specific reasons for using differentiated rate gyro measurements, then simply remove the angular accelerations from the observation vector. Avoid differentiating angular rates just because they are there.

8.9 ENGINE PARAMETERS

Engine thrust and engine speed measurements are useful in some applications, but they can be ignored in most stability and control work.

The primary need for engine thrust data is in determining drag coefficients. Without some means of subtracting the thrust component, you can only estimate the sum of the thrust and drag forces. The thrust also appears in the α state equation, but that term is usually of minor importance. The importance of obtaining thrust measurements is directly related to the importance of the thrust to the vehicle stability and control. Thrust data are critical, for instance, for vectored-thrust configurations.

Inclusion of engine gyroscopic effects requires measurement of engine speed. Such effects are often small enough to ignore in small-perturbation stability and control maneuvers; they are most likely to be important in small high-powered airplanes. Large gyroscopic effects cause coupling between the longitudinal and lateral-directional motions, which is detrimental to the handling qualities.

Methods for measuring engine thrust are outside the scope of this document. Signals required to compute engine thrust are likely to include pressures and temperatures at various places in the propulsion system, rotational speeds, fuel flows, and engine control settings. Calibration of engine thrust measurement is a major portion of a flight test program.

8.10 CONFIGURATION PARAMETERS

The vehicle configuration parameters are the positions of all flaps, canards, landing gear, wings (sweep or skew), engine controls, external stores, and other items that affect the aerodynamic or control characteristics of the vehicle and are constant during a maneuver. Changeable control system characteristics, such as control system modes and gains, are also part of the configuration data.

It is important to know the configuration to which the flight data apply. For small, simple programs, lap notes or other manual records are adequate for configuration data. For large programs with many configurations and maneuvers, digital recording of configuration variables is advisable to lessen the chances of confusion and error. Digitally recording configuration data also aids automated handling of large quantities of data.

Such digital recording can be a great convenience, but it is not a requirement. Hand recording is adequate for parameters that remain constant during a maneuver. Any parameter that changes value during a maneuver is more properly classified as a control than as a configuration parameter. There is obviously some overlap in these classifications. Flap position can be a configuration parameter for some maneuvers and a control in others. We can even have, for instance, the average elevator position during a maneuver as a configuration parameter, with perturbation in elevator position as a control in the same maneuver. In such cases of overlap, the requirements for measuring control positions apply.

8.11 LOADING DATA

The loading data are those measurements needed to define the mass characteristics of the vehicle during each flight maneuver. These data include fuel weights in various tanks (or fuel flow rates, from which you deduce remaining fuel weights by integration) and payload weights and positions. Geometric configuration parameters are also necessary to define the mass distribution. Section 5.2 discusses the importance of knowing the vehicle mass characteristics.

Digital recording of loading data is a great convenience to automated data handling but is seldom a requirement. Hand-recorded data are adequate in most cases. Automated recording of fuel flow is necessary if your method of measuring fuel weights involves integrating fuel flows.

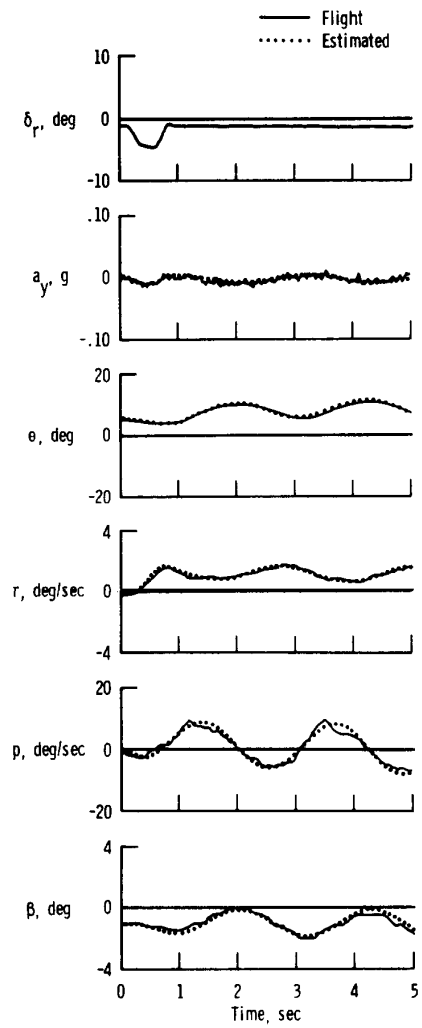

Figure 8.3-1. Apparent roll rate gyro failure.

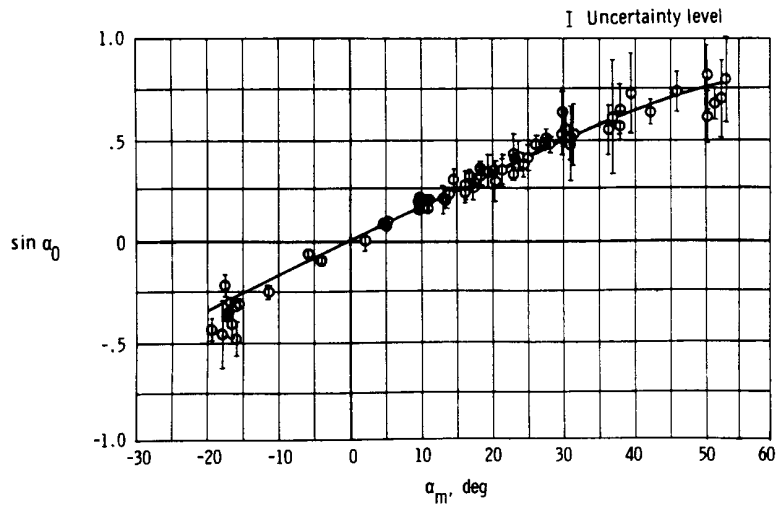

Figure 8.5-1. Estimates of sin α based on β state equation.

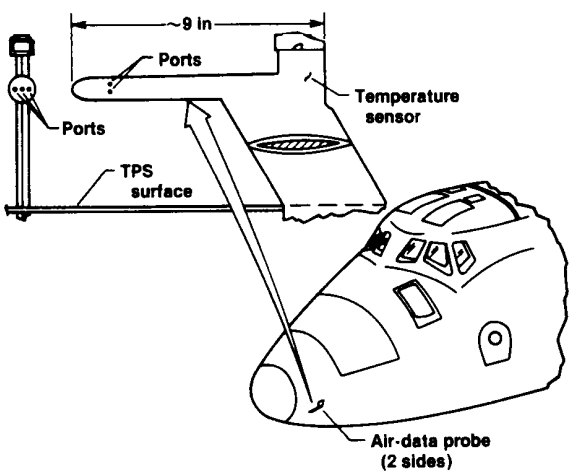

Figure 8.5-2. Shuttle side probes.

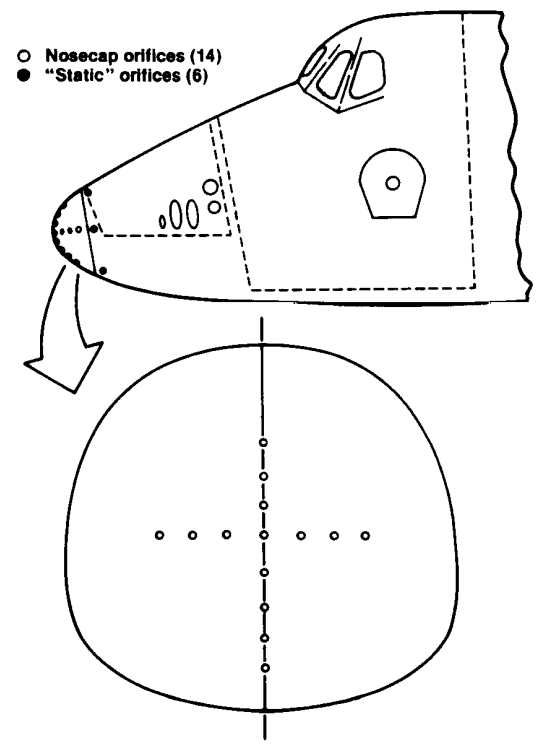

Figure 8.5-3. Shuttle entry air data system (SEADS) ports.

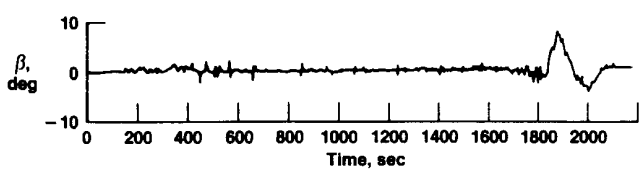

Figure 8.5-4. Inertial sideslip angle during second shuttle entry.

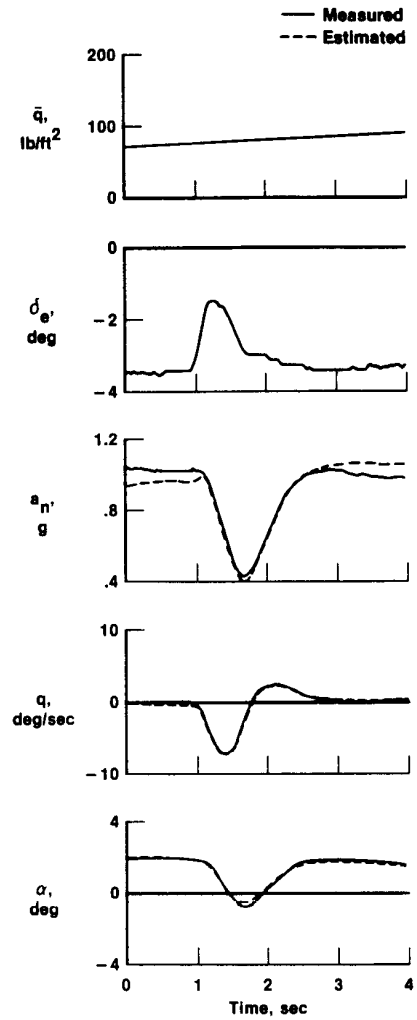

Figure 8.7-1. Fit neglecting dynamic pressure change.

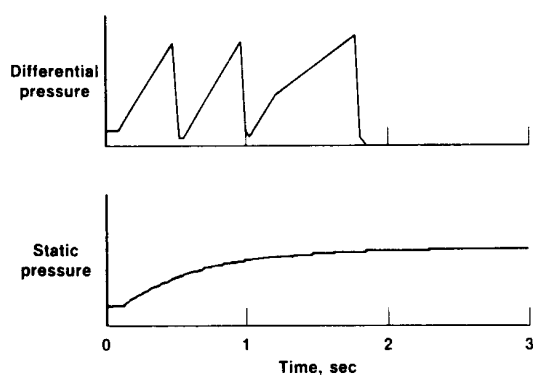

Figure 8.7-2. Step response test of pressure sensor.

9.0 EVALUATION OF RESULTS

This chapter treats the evaluation of the results of the parameter estimation process. There is some discussion of evaluation throughout this document because of the closed-loop nature of the analysis process. Evaluating the results is an integral part of the analysis process, not something that waits until the analysis is done. You must investigate the reasons for poor or questionable results and endeavor to acquire new data or reanalyze the existing data to improve the results.

Maine and Iliff (1984, chapter 11) discusses many of the general principles of evaluating parameter estimation results. The evaluation tools presented include the Cramér-Rao bound, bias, scatter, and engineering judgment. We do not repeat those general discussions here.

In this chapter we address specific issues of computing the Cramér-Rao bound, which is the primary analytical tool for evaluating the accuracy of the estimates. We use the Cramér-Rao bound in other chapters, but this chapter considers the practical aspects of its computation in detail. This chapter also shows several examples of evaluating data using the Cramér-Rao bound in conjunction with other tools. Finally, we briefly discuss the importance of considering modeling issues in evaluation.

9.1 COMPUTATION OF THE CRAMÉR-RAO BOUND

The Cramér-Rao bound is one of the primary analytical tools for evaluating parameter identification results; it gives an approximation to the variance of the estimates. Maine and Iliff (1984, chapter 11) shows the derivation of the Cramér-Rao bound.

Because of the numerous qualifications involving modeling error and noise statistics, it is necessary to validate the theoretical properties of the computed Cramér-Rao bounds before much confidence can be placed in them. Comparing the Cramér-Rao bounds and the sample standard deviations 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 section examines these discrepancies and explains their cause. We also suggest and evaluate an approximate correction for these effects.

9.1.1 Discrepancy in the Cramér-Rao Bound

We choose data from a Piper PA-30 aircraft (Figure 9.1-1) to evaluate the Cramér-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 using an output-error maximum likelihood method. Figure 9.1-2 shows a typical time history match.

Figure 9.1-3 presents the estimates of C_{n_β} and C_{ℓ_β} and the Cramér-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; you can regard the 18 maneuvers as being at essentially the same flight condition. The Cramér-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 Cramér-Rao bounds. Quantitatively, the sample standard deviation is about nine times the average Cramér-Rao bound.

This discrepancy between the data scatter and the Cramér-Rao bound has long been known. It has become a common practice to multiply the Cramér-Rao bounds by an empirically derived "fudge factor" of 5 to 10 (Iliff and Maine, 1975). The resulting values have proven useful for evaluating the accuracy of estimates, but the necessity of the unexplained fudge factor has detracted from the confidence. Because of this problem with the Cramér-Rao bounds, several reports have used the estimated correlations as primary indicators of accuracy (Wells and Ramachandran, 1976; Suit, 1972) in spite of the previously mentioned problems with the correlations. The estimated correlations are based on the same theoretical foundation as the Cramér-Rao bounds and thus should be equally suspect if errors are known to exist. Other reports have ignored the fudge factor and quoted overly optimistic values of the accuracy (Williams and Suit, 1974); 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 provides a more controlled environment, which may aid preliminary work. We therefore repeated the preceding experiment with simulated data.

To mimic the flight data experiment as closely as feasable, we used the control inputs measured from the flight data to create the simulated data. We created 18 simulated maneuvers using the same flight conditions as those of the 18 actual maneuvers. 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. We used a pseudorandom noise generator 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. We adjusted the proportionality constant to obtain the same magnitude of scatter in the estimates from the simulated data as was observed in the estimates

from the flight data. The derivatives and Cramér-Rao bounds were estimated from these simulated data using the same program as for the flight data. Figure 9.1.4 shows a typical time history match.

The estimates of C_{n_β} and C_{ℓ_β} from the simulated data are shown in Figure 9.1-5 plotted to the same scale as in Figure 9.1-4. The true values are $C_{n_\beta} = 0.001$ and $C_{\ell_\beta} = -0.0007$.

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 Cramér-Rao bounds differ drastically on the flight and simulated data. The Cramér-Rao bounds on the simulated data are about 10 times those 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 deviations to the Cramér-Rao bounds range from 0.68 to 1.32. This is excellent agreement for a sample of only 18 maneuvers.

The Cramér-Rao bound thus agrees well with the scatter on simulated data, but disagrees drastically on flight data. Sections 9.1.2 and 9.1.3 examine the reasons for this discrepancy.

9.1.2 Previous Attempts at Explanation

Researchers have advanced several ideas attempting to resolve the discrepancy in the Cramér-Rao bound. Few of these ideas have been discussed extensively in the literature because of their speculative nature. Maine and Iliff (1981a) examined these proposals and showed that none of them provided a satisfactory explanation for discrepancies of the magnitude observed in practice.

The excellent performance of the Cramér-Rao bound on simulated data refutes many conceivable explanations of the discrepancy: One 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. A second refuted explanation is the fact that the Cramér-Rao bound is only a lower bound. The maximum likelihood estimator is proven to be asymptotically efficient, that is, the Cramér-Rao inequality approaches an equality as time approaches infinity. For finite time, the equality does not hold. However, 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 scatter of the simulated data agrees so well with the Cramér-Rao bound verifies that the time was long enough to make the equality a very good approximation.

These attempts to explain the discrepancy between the Cramér-Rao bounds and the scatter implicitly assume that the scatter is a reasonable measure of the accuracy. For the simulated data, where the true values are known to be constant, the scatter almost defines the accuracy for an unbiased estimator. For the flight data, other possibilities must be considered. Recall that the simulated data noise level was chosen to make the scatter match that of the flight data, and the resulting Cramér-Rao bounds of the simulated data were larger than those of the flight data. As is evident by comparing Figures 9.1-2 and 9.1-4, the noise power in this simulated data is much larger than that of the flight data residuals. The difference in the Cramér-Rao bounds arises directly from the difference in the noise power. If the simulated data noise power were lowered to the same level as that of the flight data residuals, the simulated and flight Cramér-Rao bounds would be the same. Of course, the scatter of the simulated data would be much less than that of the flight data.

This suggests the possibility that instead of the Cramér-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 Cramér-Rao bounds and that the scatter reflects actual changes in the coefficients. In principle, 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 like the PA-30 shown here, airliners (Tanner and Montgomery, 1979), military fighters (Iliff et at., 1978; Foster, 1977), large supersonic aircraft (Powers, 1977), and unconventional vehicles (Sim, 1976; Maine, 1978). 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 Cramér-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 Cramér-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 in data from many different aircraft and data systems argues strongly against this as a cause. We believe that most of the data systems used are accurate enough to eliminate problems of errors in the flight condition.

None of these possibilities has proven to be 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 that invalidated the Cramér-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 errors 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 Cramér-Rao bound as invalid.

This situation has more far-reaching implications than the invalidation of the Cramér-Rao bound for use with flight data. The theoretical derivation of the Cramér-Rao bound rests on the likelihood functional. If this theory is found inadequate, the theoretical justification for the use of maximum likelihood estimators must also be questioned since the estimator is based on the same likelihood functional. Thus, an invalidation of the Cramér-Rao bound might imply an invalidation of the estimates, leaving us with nothing theoretically worthwhile.

9.1.3 Explanation of the Discrepancy

Iliff and Maine (1977) advanced the first satisfactory explanation of the discrepancy in the Cramér-Rao bound. The discrepancy was traced to the theoretical assumptions about independence of the noise samples.

The long period of time that elapsed before this explanation was advanced points out the necessity for researchers who are well grounded in both theory and practice, because the problem lay essentially in a lack of communication between the theoreticians and the practitioners. The theory is most naturally developed assuming independence of the noise samples, and it had been passed to the practitioners in this form. Experienced practitioners were well aware that residual spectra are seldom even close to white, but they accepted the assumptions as being necessary to the theory. In fact, the theory can easily give much information about the effects of colored noise. The theoreticians were unaware that the results were needed, and the practitioners were unaware that helpful theory was available.

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. Estimating 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 issue. For a first approximation, assume that the noise is band-limited white with band limit B. Rather than derive a maximum likelihood estimator for this system, we analyze the performance of the white-noise-based estimator when the actual noise is colored as described. The results of this analysis 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 for actual system response and some percentage of the actual system response for noise. 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 easy to generalize. It is obvious that the exact shape of the noise roll-off was of little consequence in the preceding analysis. We could repeat the analysis with different types of roll-off characteristics and would obtain the same results. 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 functions of the estimates are equally unaffected. Thus, in particular, the same expression for the Cramér-Rao bound should still be valid. This would seem to refute the thesis that the discrepancy in the Cramér-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 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 Cramér-Rao bounds should be correct. The noise power spectral density of this system is $2R\Delta t$ because the noise spectrum is flat up 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 9.1-6 shows the noise spectra of these two systems.

The Cramér-Rao inequality for both systems is

$$\mathrm{var}(\hat{\xi}) \geq \left\{ \sum_{i=1}^{N} (\nabla_{\xi}\tilde{z}^{*}) R^{-1} (\nabla_{\xi}\tilde{z}^{*})^{*} \right\}^{-1} \tag{9.1-1}$$

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

$$\mathrm{var}(\hat{\xi}) \geq \left\{ \sum_{i=1}^{N} (\nabla_{\xi}\tilde{z}^{*}) \left(\frac{1}{2\Delta t} GG^{*} \right)^{-1} (\nabla_{\xi}\tilde{z}^{*})^{*} \right\}^{-1} \tag{9.1-2}$$

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

$$\mathrm{var}(\hat{\xi}) \geq \frac{1}{2B\Delta t} \left\{ \sum_{i=1}^{N} (\nabla_{\xi}\tilde{z}^{*}) R^{-1} (\nabla_{\xi}\tilde{z}^{*})^{*} \right\}^{-1} \tag{9.1-3}$$

The Cramér-Rao bound computation based on the total power was therefore too small by a factor of $1/(2B\Delta t)$ in the variance (or $\sqrt{1/(2B\Delta t)}$ in the standard deviation). The Cramér-Rao bound computation based on the power assumes that the power is evenly spread over the Nyquist range as in the first case. Thus, in Equation (9.1-3), B is assumed to be $1/(2\Delta t)$. If this assumption is incorrect, the resulting variance computation will be proportionally incorrect.

These results excellently explain the discrepancies previously observed. The noise samples on the simulated data were independent, and thus the corresponding Cramér-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 addressed in Maine and Iliff (1984). Since our theory contains only the linear system model and measurement noise, any effects not included in the system model must contribute to the measurement noise (otherwise our theory is denying their existence; such a solipsistic approach seems unwise). The measurement noise in real flight test data therefore tends to be quite colored. Figure 9.1-7 shows power spectral density plots of the residuals from the flight test data used in Figure 9.1-3. A precise break point is not obvious, but a value near 0.3 Hz seems reasonable. The Nyquist frequency for these 50-sample/sec data is 25 Hz, so the computed Cramér-Rao bounds should be increased by a factor of about $\sqrt{25/0.3} \cong 9$. The sample standard deviations were about nine times the Cramér-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 these data was created by passing the pseudorandom independent noise through a fifth-order Chebyshev filter with a break frequency of 1 Hz; this filter has a sharp break at 1 Hz as shown by Figure 9.1-8, which is a power spectral density of one of the resulting measurement noise signals.

Figure 9.1-9 shows a typical time history match to these data. Note that this match exhibits deterministic-appearing characteristics, such as phase shift and flattened peaks. The matches of β, p, r, and Φ are more typical of flight data than are the simulations shown in Figure 9.1-4. Accelerometers, by their nature, have more high-frequency noise than such instruments as rate gyros. Therefore, the flight accelerometer matches tend to be of a character intermediate between that of Figures 9.1-4 and 9.1-9. Figure 9.1-10 shows the Cramér-Rao bounds and estimates of C_{n_β} and C_{ℓ_β} from the simulated data with filtered noise. 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 Cramér-Rao bounds are corrected for the colored noise, they agree excellently with the scatter.

These results support the conclusion that the discrepancy in the Cramér-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 Cramér-Rao bound must represent the entire residual error including modeling error contributions (which may be much larger than the actual instrumentation error). Therefore, studies made solely on the basis of instrumentation characteristics (Hodge and Bryant, 1975; Sorensen, 1972; Gupta and Hall, 1978) are likely to be extremely overoptimistic. Colored noise would also be expected to affect the insensitivities, but it should not directly affect the correlations. The correlations might, however, be affected by the fact that the noise contributions from modeling error on the various signals tend to be correlated.

9.1.4 Suggested Implementations

Section 9.1.3 explains the reasons for the discrepancies observed in the Cramér-Rao bounds. It remains to discuss practical implementation of a corrected computation of the bound. We consider three approaches.

The first approach is to continue the use of fudge factors. The Cramér-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 that is based on past experience rather than analytically computed for each maneuver. It could be regarded 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, further justifying this approach. The advantages to this approach 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 that 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 criticism on the basis of arbitrariness, but when considered as a tool to aid the engineer's evaluation, instead of as an absolute value of accuracy, it can continue to be useful, as it has been in the past.

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 used directly. 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 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 required. It is always a good practice to examine at least a few sample residual power spectral density plots anyway, but it is simplest 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 9.1-7 shows 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 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 with very little work it does pull out a reasonable estimate of the single value we are after. The bulk of the implementation is in 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 obvious approximations include the assumption that the noise spectrum is flat at least to the filter break frequency, the choice of filter break frequency, and the use of a first-order filter. In practice, the noise tends to be closely concentrated around the system natural frequencies. This violates the requirement that the noise bandwidth be well above the system frequencies. Because of these approximations, the method should not be regarded as giving a precise measure of the expected scatter; our only claim is that the method is an improvement over previous computations. Furthermore, some of the proposals of Section 9.1.3, although unable to explain discrepancies of a factor of 10, might contribute measurably to smaller discrepancies. The approximations in the proposed approach are such that a factor of about 2 typically remains to be multiplied, based on empirical observation.

The method might therefore seem to provide 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 approximately account for large changes in the spectral characteristics. If the spectral characteristics of all the maneuvers are similar, the results from this and the first approach are equally applicable.

Figure 9.1-11 shows the flight data from Figure 9.1-3 with Cramér-Rao bounds computed from the filtered residuals. A break frequency of 0.5 Hz was used for the filters. The correction factors computed from the filtered residual power ranged from 3 to 6, and a remaining empirical factor of 2 was used. The magnitudes of the Cramér-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 Cramér-Rao bounds than the outliers. This approach results in a useful estimate of the accuracy.

9.2 EXAMPLES OF APPLICATION

This section presents several examples of application of the Cramér-Rao bound to actual flight data. The examples illustrate the kinds of information that can be deduced with the aid of the bounds.

9.2.1 Example 1

The first example uses data from a remotely piloted oblique-wing vehicle (Maine, 1978) shown in Figure 9.2-1. Cross-coupling derivatives between the longitudinal and lateral-directional modes were of significant interest for this nonsymmetric vehicle. Figure 9.2-2 shows the estimates of two such derivatives, C_{Y_α} and C_{m_r}. An empirically determined spectral adjustment factor of 5 was used in the

Cramér-Rao bounds shown. The C_{Y_α} data show relatively small scatter and correspondingly good Cramér-Rao bounds. The effect of this derivative is readily identified from the flight data, and it agrees well with the predictions.

For C_{m_r}, on the other hand, the scatter and Cramér-Rao bounds are about the same magnitude as the estimates. There is very little information about the value of C_{m_r} in the flight data. With the wing skewed 15° to 45°, the data show that C_{m_r} is positive and probably somewhere in the range of 0.1 to 0.3. It cannot be estimated better than this from the data available.

This example illustrates the use of the scatter and the Cramér-Rao bounds to determine which derivatives can be accurately estimated from the data. Confidence in the determination is improved by using both the scatter and the Cramér-Rao bounds to support the conclusion.

9.2.2 Example 2

The second example uses data from a remotely piloted 3/8-scale F-15 model (Iliff et al., 1976) shown in Figure 9.2-3. A spectral adjustment factor of 5 was used for these data. Figure 9.2-4 shows the estimates of C_{m_q}. The behavior around 25° angle of attack is of particular interest. It is difficult to justify the fairing shown based on the data scatter alone. The fairing was based mostly on the data points with small Cramér-Rao bounds. In particular, the cluster of points near 25° angle of attack showing C_{m_q} near 0 have small Cramér-Rao bounds, while the points at the same angle of attack showing C_{m_q} values of -3 and -4 have larger bounds. The fairing therefore goes to 0 at 25° angle of attack. The estimates of C_{m_α} in Figure 9.2-5 unambiguously show a similar shape near 25° angle of attack. Subsequent results have tended to verify the validity of the C_{m_q} fairing.

In this example, the Cramér-Rao bounds provide more information than the scatter alone. The bounds indicate that some points have good accuracy, even though the overall scatter is large because of a few poor estimates. The Cramér-Rao bounds indicate that some of the estimates above 35° angle of attack and also some near -10° are quite unreliable even though the scatter is not large (see example 5, Section 9.2.5, for more on this point).

9.2.3 Example 3

Figure 9.2-6 shows estimates of C_{n_β} taken from the same 3/8-scale F-15 data as in example 2, Section 9.2.2, using the same spectral adjustment factor of 5. The scatter above 35° angle of attack is quite large, as are most of the Cramér-Rao bounds. A few points with small Cramér-Rao bounds form the basis for the fairing used. It is interesting that the three points (shaded) farthest from the fairing were obtained immediately before aircraft upsets. These more negative values of C_{n_β} may be an early manifestation of the aerodynamic phenomenon causing the upsets.

9.2.4 Example 4

The fourth example uses data from an F-111A airplane shown in Figure 9.2-7. Figure 9.2-8 shows estimates of C_{m_α} for the F-111A at 26° of wing sweep (Iliff et al., 1978). The Cramér-Rao bounds were multiplied by an empirically determined factor of 10 to correct for the effects of colored noise. The factor was chosen to make the Cramér-Rao bounds of about the same magnitude as the scatter for most of the derivatives. The scatter on C_{m_α} is much larger than the Cramér-Rao bounds, even after using the factor. The M = 0.9 points are particularly puzzling. Since using the same factor resulted in good agreement of the Cramér-Rao bounds and the scatter for the other derivatives, the C_{m_α} estimates deserve careful study.

To study these results, we replotted the data as a function of Mach number (Figure 9.2-9). The reason for the apparent discrepancies was then evident. As shown by the fairing in Figure 9.2-9 the data show a significantly less stable C_{m_α} in a region around Mach 0.85. In Figure 9.2-8 the data were rounded to the nearest M = 0.7, 0.8, or 0.9. The points near Mach 0.85 were all automatically rounded to Mach 0.9 (they were slightly above 0.85). This merging of the Mach 0.85 and 0.9 data made the Mach effect look like an inconsistency. Figure 9.2-10 shows the same data as Figure 9.2-8 with the points near Mach 0.85 shaded and fairings of the flight data added. The scatter about these fairings is reasonable.

In this example, the Cramér-Rao bound did not directly indicate the nature of the problem, but it did draw attention to an area that needed more careful study. The ensuing study disclosed the unanticipated Mach effect.

9.2.5 Example 5

Figure 9.2-11 shows estimates of C_{ℓ_r} for the F-111A airplane at 35° of wing sweep. A spectral adjustment factor of 10 was used, as described in example 4, Section 9.2.4. The data in this figure were gathered in an effort to investigate derivative changes as a function of normal acceleration; such changes could be caused by structural deformation. The solid line is a fairing of previous 1-g flight data.

This figure has characteristics the opposite of those in example 4; the adjusted Cramér-Rao bounds are much larger than the data scatter. The estimates are all very close to the fairing of the previous data. This behavior is typical of a problem that can occur when you use an *a priori* weighting (that is, when you use an MAP estimator).

Convergence difficulties were encountered in analyzing several of the elevated-g maneuvers, which were often not well-stabilized maneuvers. Therefore, an *a priori* weighting was used to improve the convergence. The *a priori* values are indicated by the solid line in Figure 9.2-11.

The effect of the *a priori* weighting is to hold estimates near the *a priori* values unless there is significant evidence that the *a priori* values are incorrect; the level of significance required can be adjusted. The *a priori* weighting thus tends to eliminate large estimate changes that are based on minimal information. This can improve convergence if the information content of the manuever is poor.

The Cramér-Rao bounds in Figure 9.2-11 indicate that the maneuvers contained very little information about the derivative C_{ℓ_r}. The small scatter is probably due almost entirely to the *a priori* weighting, rather than to information from the maneuvers. Thus, the data in Figure 9.2-11 do not verify the previous flight data, a conclusion that might be mistakenly drawn if the Cramér-Rao bounds were not considered. Instead, Figure 9.2-11 indicates that the new data do not have sufficient information to contradict or verify the *a priori* values. In such a case, the estimator holds near the *a priori* values when an *a priori* weighting is used.

This example also illustrates an important fact about the computation of the Cramér-Rao bound when an *a priori* weighting is used. The Cramér-Rao bounds are computed from an approximation to the second gradient of the cost function. When an *a priori* weighting is used, a penalty function is added to the cost function. To estimate the information content of the flight maneuver, the matrix used for the Cramér-Rao bound must be computed based only on the original cost function; it should not have a term added for the second gradient of the penalty function. If the second gradient of the penalty function were added into the computation of the bound, the bound would reflect the sum of the information from the maneuver and from the *a priori* weighting. The phenomenon in Figure 9.2-11 would not be observed (the Cramér-Rao bounds in the figure would be much smaller).

There are applications where you want to use the information from both sources and to have a corresponding accuracy estimate. The Cramér-Rao bound with the second gradient of the penalty function included is a possible source for such an accuracy estimate. However, it gives no hint as to how much of the information was obtained from the maneuver and how much was obtained from the prior information. Knowing how much of the information came from each source helps avoid misinterpretations.

9.2.6 Example 6

Figure 9.2-12 shows the estimates of $C_{n_{\delta_r}}$ obtained from the same flight data used in example 5, Section 9.2.5. As in example 5, the data scatter is small and agrees well with the fairing of the previous flight data. *A priori* weighting was used for both examples. In this example, however, the Cramér-Rao bounds are small and are consistent with the scatter. It can be concluded that the maneuvers contain significant information about the derivative $C_{n_{\delta_r}}$. Therefore, the data shown give positive verification of the *a priori* values. This contrasts with example 5, where we could only conclude that the new data did not contradict the *a priori* values.

9.3 MODELING CONSIDERATIONS

In this document and in Maine and Iliff (1984), we have emphasized issues of modeling. All the analysis and evaluation tools depend to some degree on assumptions of model validity. If the assumed model is inappropriate, then the parameter estimates are of little value. Evaluation tools such as the Cramér-Rao bound will sometimes detect modeling errors, but they cannot be relied on for such purposes. Critical consideration of model adequacy must pervade the entire estimation process.

In evaluating and using results, it is not sufficient to know that the estimates came from an adequate model. For any given problem, there are numerous adequate model forms. It is important to know precisely which model was used so that you can evaluate the parameter estimates in the context of that model. Estimates obtained with one model structure may not be directly comparable with values derived using different model assumptions. This need not mean that either set of values is wrong. It does mean that knowledge of the parameter values alone is insufficient; you also need to know what model the parameters apply to.

A classic example concerns estimates of q and $\dot{\alpha}$ derivatives. Consider the following simplified longitudinal model, obtained from Equation (3.4-3):

$$\dot{\alpha} = -\frac{\bar{q}s}{mV} C_{L_\alpha}\alpha + q + \left(\frac{g}{V} - \frac{\bar{q}s}{mV} C_{L_0}\right) \tag{9.3-1a}$$

$$\frac{I_y}{\bar{q}sc}\dot{q} = C_{m_\alpha}\alpha + C_{m_q}\frac{qc}{2V} + C_{m_0} \tag{9.3-1b}$$

Suppose, now, that we want to expand this model by the addition of a $C_{m_{\dot{\alpha}}}$ term. Equation (9.3-1a) is unaffected, and Equation (9.1-1b) is replaced by

$$\frac{I_y}{\bar{q}sc}\dot{q} = C_{m_\alpha}\alpha + C_{m_q}\frac{qc}{2V} + C_{m_0} + C_{m_{\dot{\alpha}}}\frac{\dot{\alpha}c}{2V} \tag{9.3-2}$$

Substituting Equation (9.3-1a) into Equation (9.3-2) and rearranging gives

$$\dot{\alpha} = -\frac{\bar{q}s}{mV}C_{L_\alpha}\alpha + q + \frac{g}{V} - \frac{\bar{q}s}{mV}C_{L_0} \tag{9.3-3a}$$

$$\frac{I_y}{\bar{q}sc}\dot{q} = \left(C_{m_\alpha} - \frac{\bar{q}sc}{2mV^2}C_{m_{\dot{\alpha}}}C_{L_\alpha}\right)\alpha + \left(C_{m_q} + C_{m_{\dot{\alpha}}}\right)\frac{qc}{2V} + C_{m_0} + C_{m_{\dot{\alpha}}}\frac{c}{2V}\left(\frac{g}{V} - \frac{\bar{q}s}{mV}C_{L_0}\right) \tag{9.3-3b}$$

For current purposes, we will assume that the value of $C_{m_{\dot{\alpha}}}$ is known. This avoids an identifiability problem, which, although important, is not relevant to the current discussion.

The sum $C_{m_q} + C_{m_{\dot{\alpha}}}$ in Equation (9.3-3) plays the same role as the term C_{m_q} in Equation (9.3-1). Parameter estimates based on these two models will reflect this equivalence. The estimates of C_{m_q} obtained using the model of Equation (9.3-1) will equal the sum of the estimates $C_{m_q} + C_{m_{\dot{\alpha}}}$ obtained using the model of Equation (9.3-3). Unless $C_{m_{\dot{\alpha}}}$ is zero, this implies that the C_{m_q} estimate from Equation (9.3-1) is not equal to that from Equation (9.3-3); analysis based on such an equivalence will give unreasonable results. Similarly, the C_{m_α} from Equation (9.3-1) is not directly comparable with that from Equation (9.3-3); however, the numerical difference between these two C_{m_α} values is small for most conventional aircraft configurations, and it is usually neglected.

The issue of model comparability often arises in comparing wind-tunnel and flight estimates. Models like Equation (9.3-1) are often used in flight data analysis, whereas Equation (9.3-3) is more representative of most wind-tunnel test methods. The result is that wind-tunnel C_{m_q} estimates are significantly different from flight C_{m_q} estimates obtained using Equation (9.3-1). This difference does not mean that either of the models or their corresponding derivative values are wrong. Both can be useful models of aircraft behavior. The important thing is that C_{m_q} estimates from Equation 9.3-1 apply only to that model; it is incorrect to use them in Equation 9.3-3. This is what we mean by the statement that you must evaluate parameter estimates in the context of the assumed model structure.

If you apply estimated parameter values to a different model structure, you must consider the effects of the changed structure. In some cases there is no effect, and the same parameter values apply. In other cases there are simple equivalence relationships, such as the one derived here between C_{m_q} and $C_{m_q} + C_{m_{\dot{\alpha}}}$. Finally, there are cases where the model structures are too different for the parameter to be meaningfully translated.

The first, and most crucial, step in considering the effects of model differences is to recognize when a model difference exists. This requires you to be conscious of what model was used in the estimation and what model you are using in application or evaluation. Several misapplications of parameter estimates have arisen from failure to consider model structure as an issue.

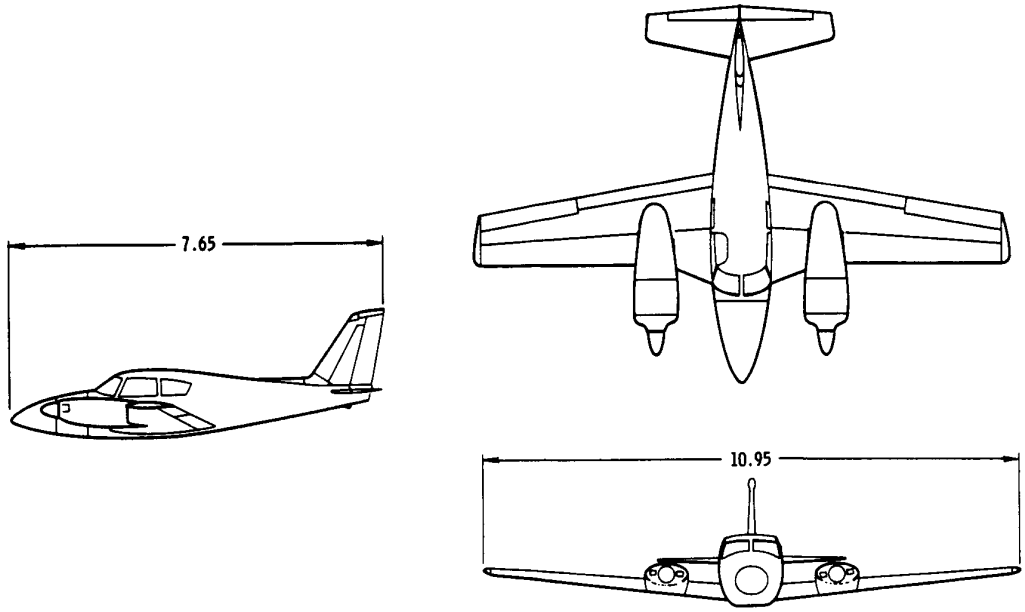

Figure 9.1-1. The PA-30 airplane (dimensions in meters).

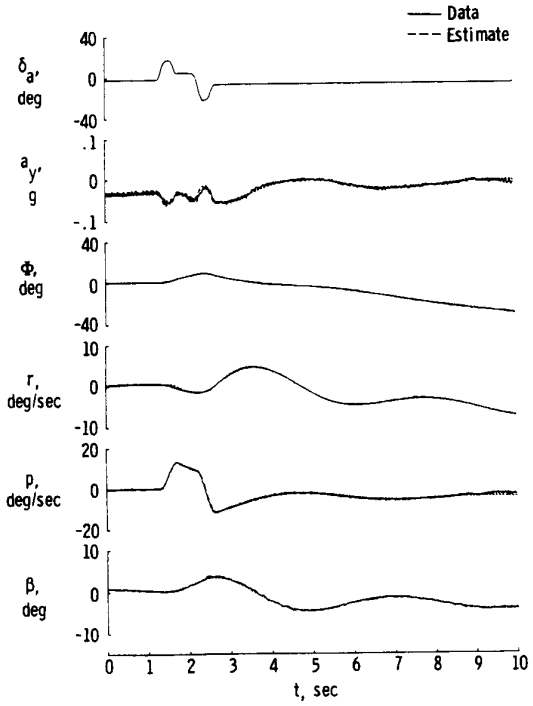

Figure 9.1-2. Time history of PA-30 flight maneuver.

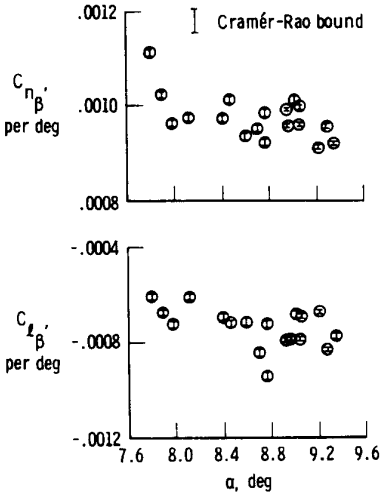

Figure 9.1-3. Estimates of C_{n_β} and C_{ℓ_β} from flight data.

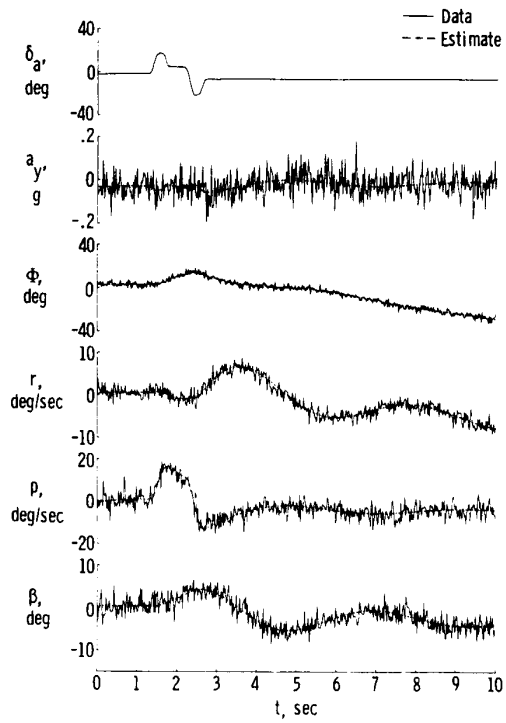

Figure 9.1-4. Time history of simulated data.

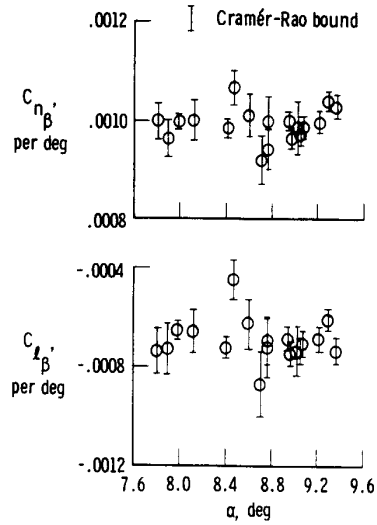

Figure 9.1-5. Estimates from simulated data with white noise.

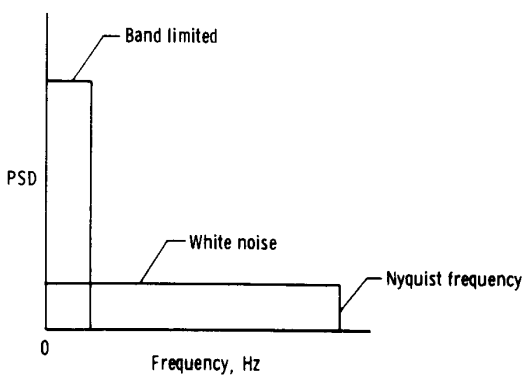

Figure 9.1-6. Idealized noise spectra.

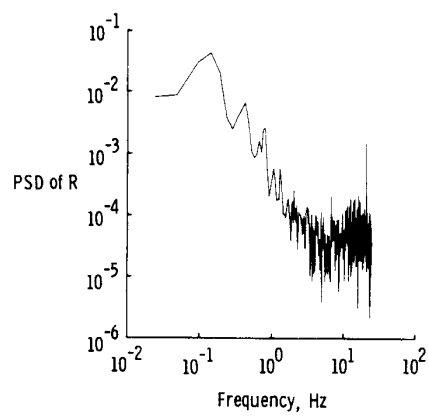

Figure 9.1-7. Noise spectra of flight data residuals.

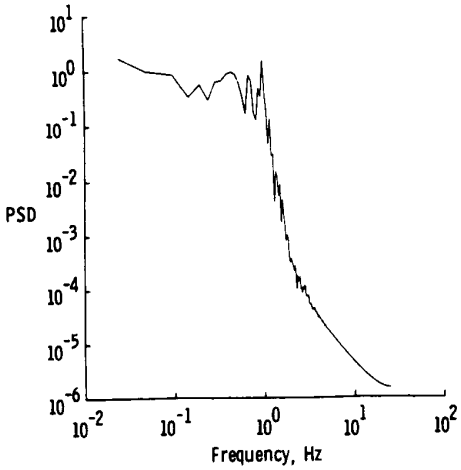

Figure 9.1-8. Spectrum of simulated colored noise.

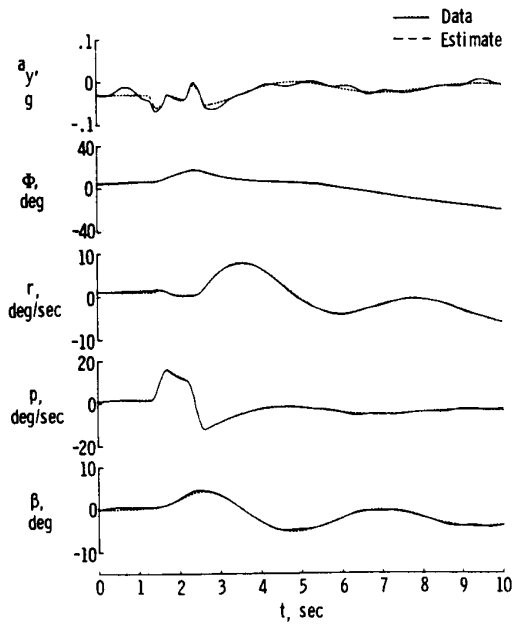

Figure 9.1-9. Time history of simulated data with colored noise.

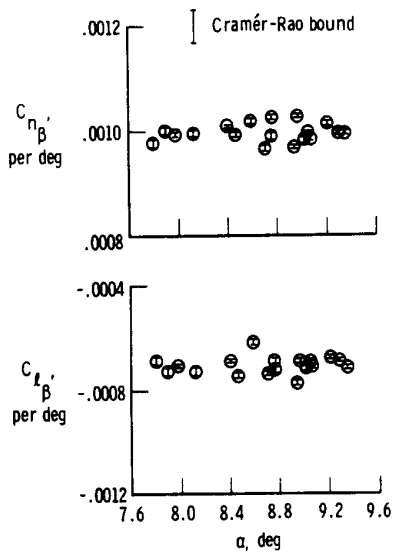

Figure 9.1-10. Estimates from simulated data with colored noise.

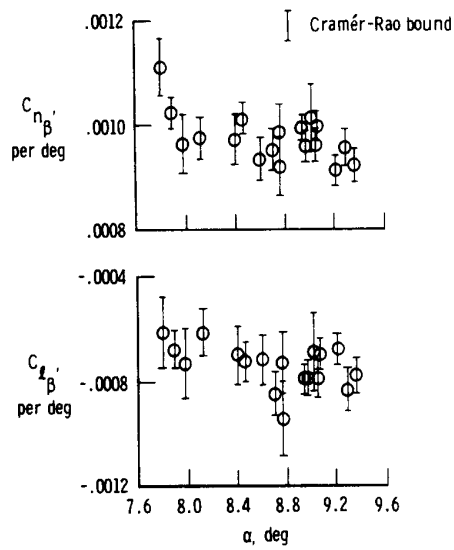

Figure 9.1-11. Estimates from flight data corrected using filter and empirical factor.

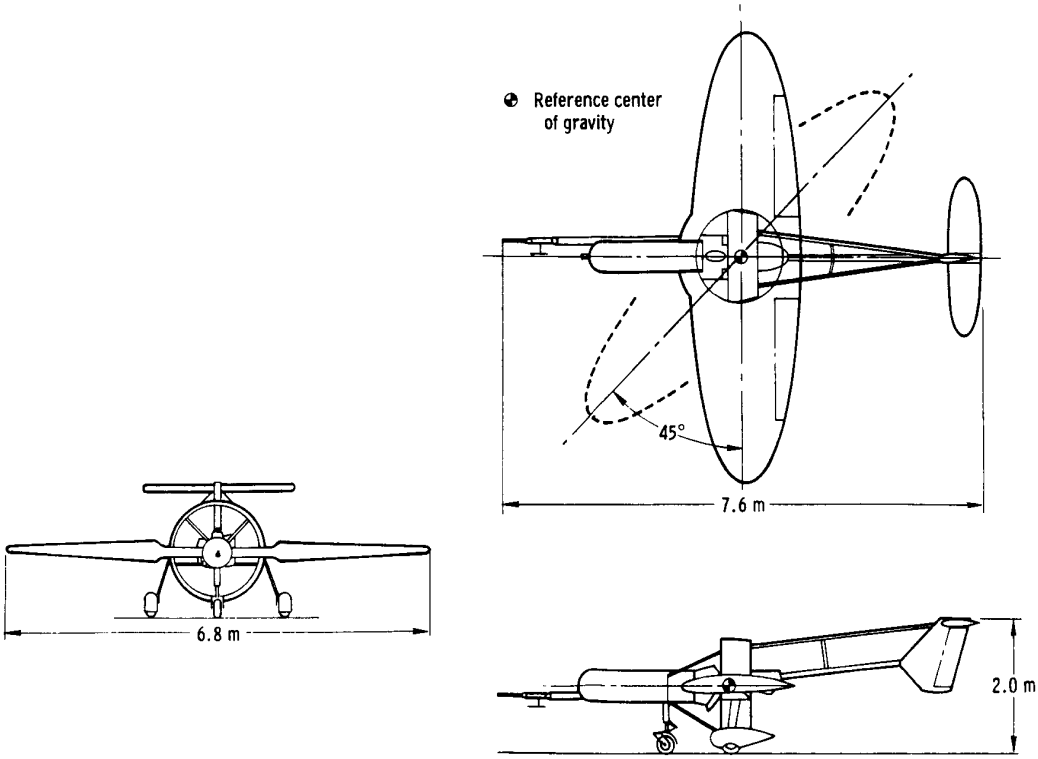

Figure 9.2-1. Oblique-wing aircraft.

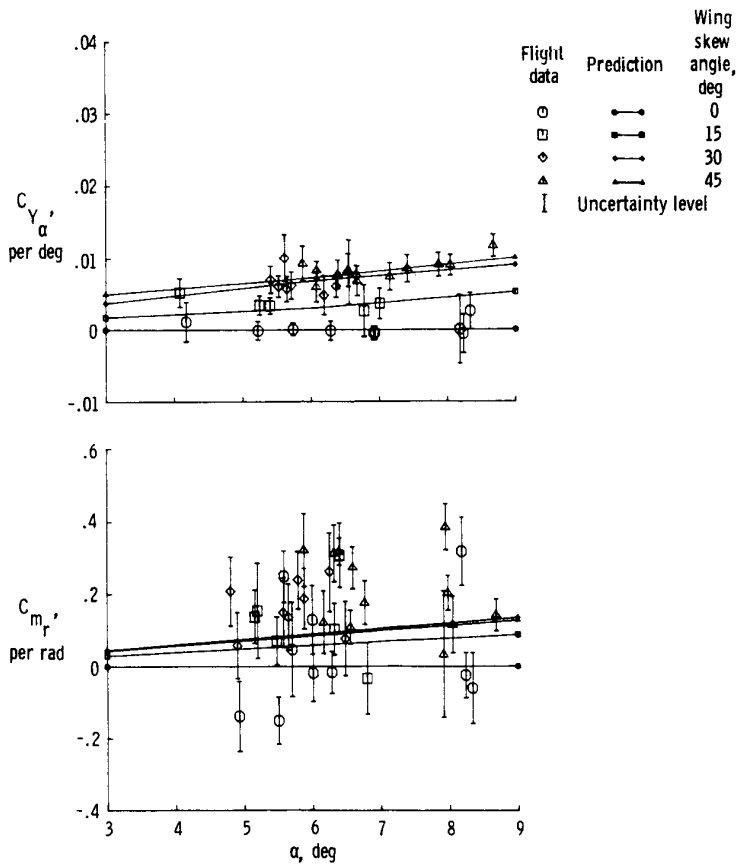

Figure 9.2-2. Estimates of C_{Y_α} and C_{m_r} for obilique-wing aircraft.

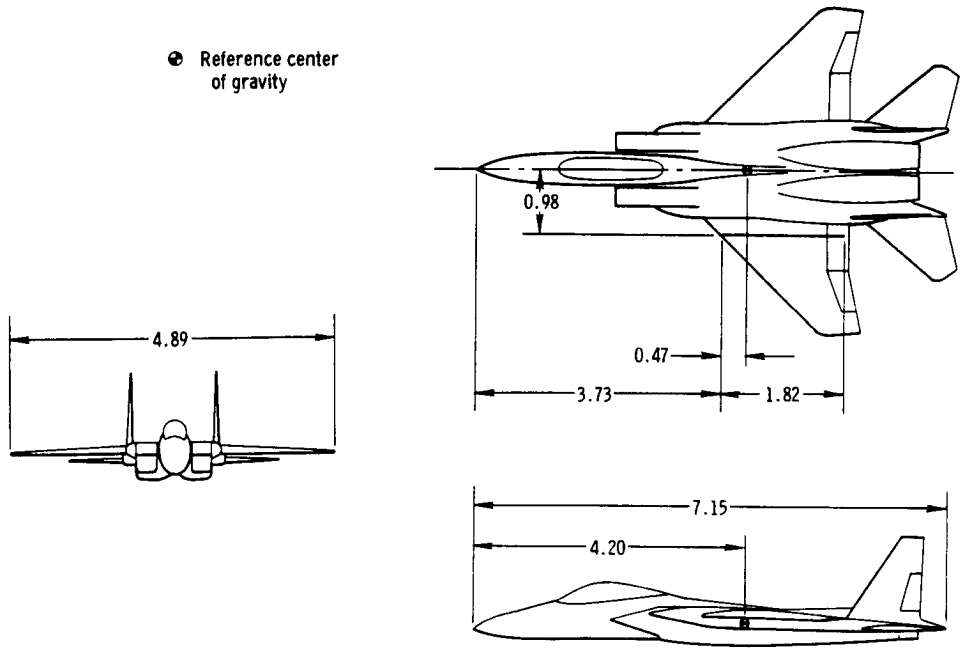

Figure 9.2-3. The 3/8-scale F-15 model (dimensions in meters).

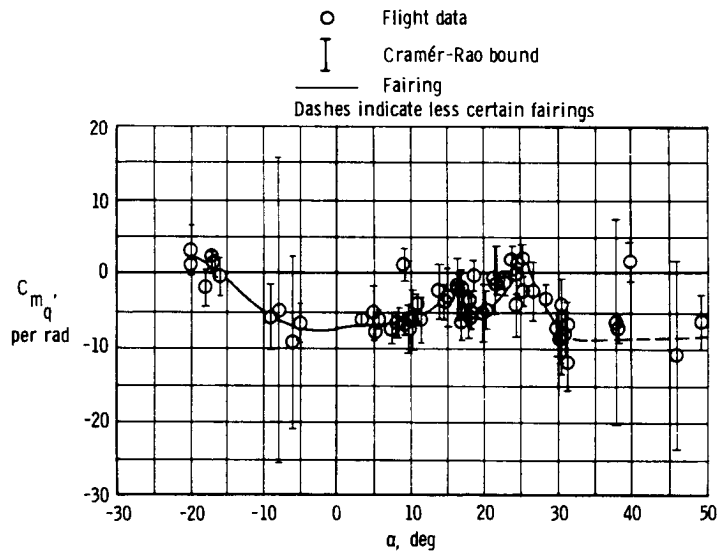

Figure 9.2-4. Estimates of C_{m_q} for 3/8-scale F-15 model.

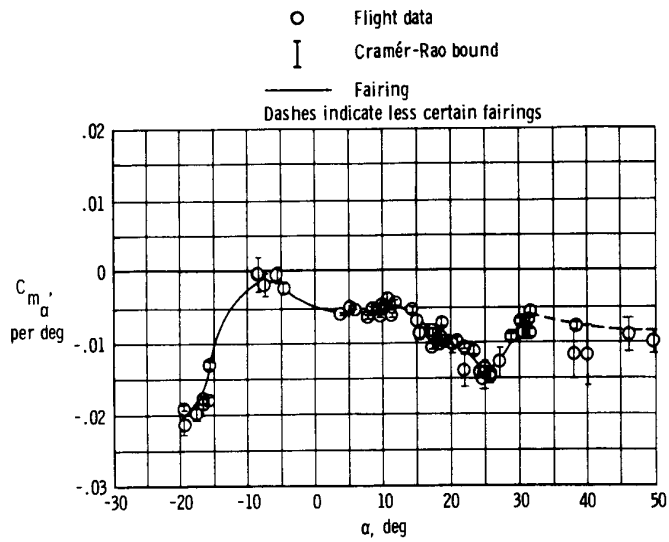

Figure 9.2-5. Estimates of C_{m_α} for 3/8-scale F-15 model.

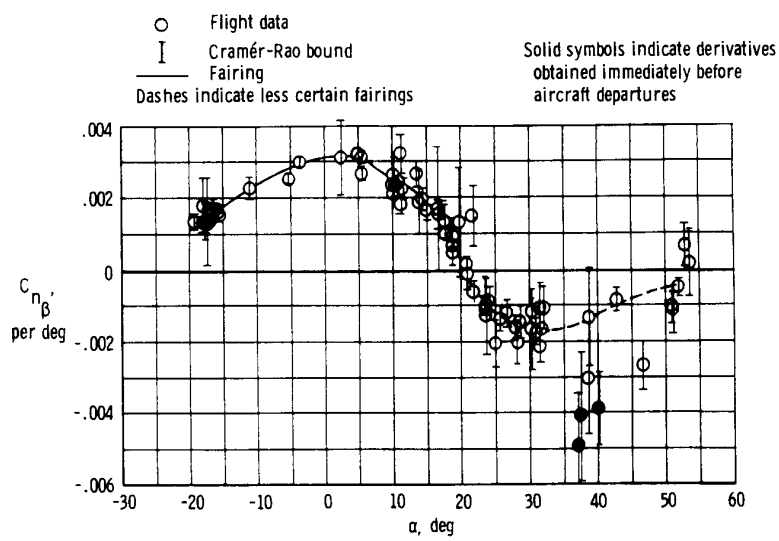

Figure 9.2-6. Estimates of C_{n_β} for 3/8-scale F-15 model.

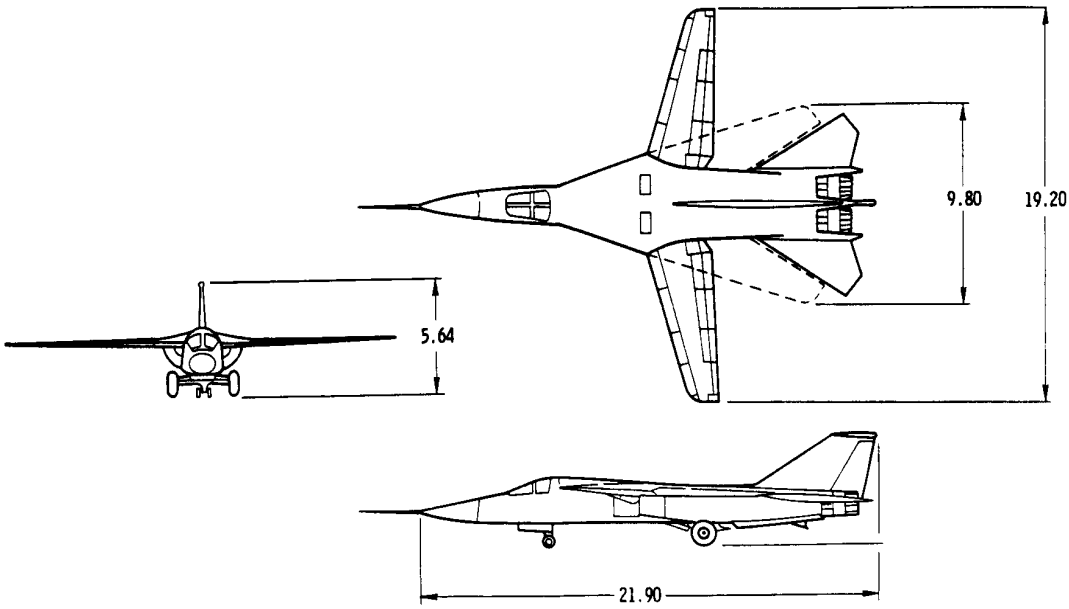

Figure 9.2-7. F-111A airplane (dimensions in meters).

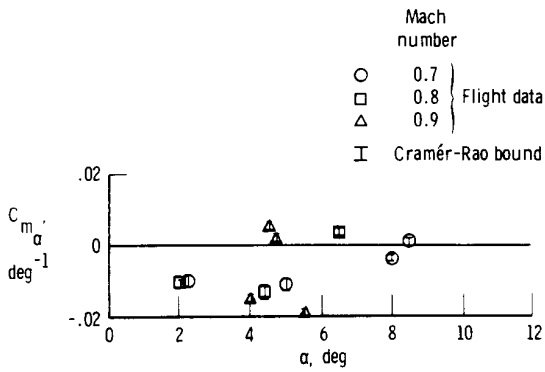

Figure 9.2-8. Estimates of C_{m_α} for F-111A airplane.

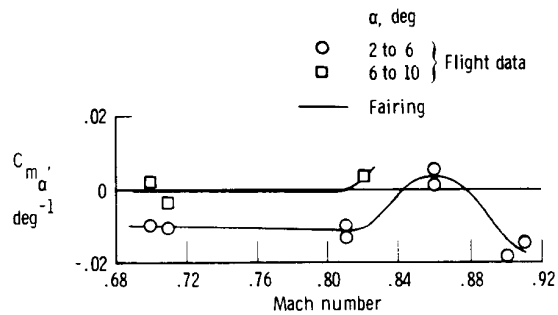

Figure 9.2-9. Estimates of C_{m_α} as a function of Mach number for F-111A airplane.

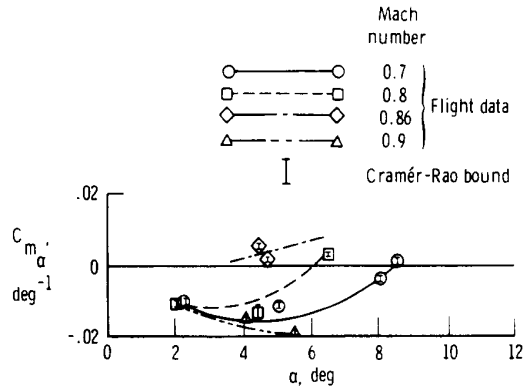

Figure 9.2-10. Estimates of C_{m_α} for F-111A airplane.

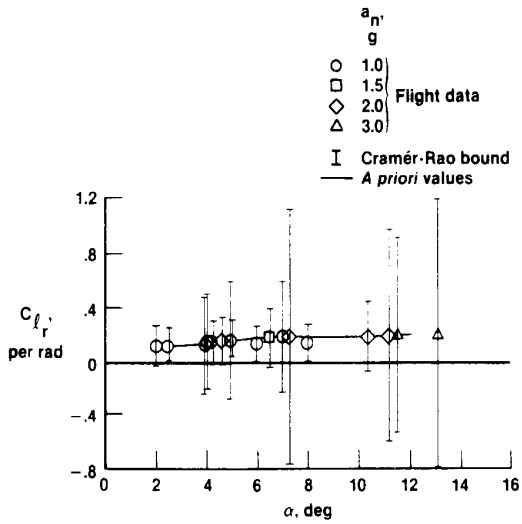

Figure 9.2-11. Estimates of C_{ℓ_r} for F-111A airplane.

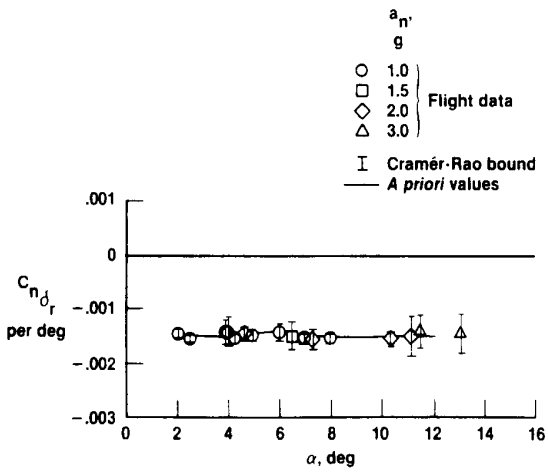

Figure 9.2-12. Estimates of $C_{n_{\delta_r}}$ for F-111A airplane.

10.0 CONCLUDING REMARKS

This document has examined the practical application of parameter estimation methodology to the problem of estimating aircraft stability and control derivatives from flight test data. We have concentrated on the output-error method in order to provide a focus for detailed examination. There are, however, other viable approaches not covered in this document. It is not realistically possible for one document both to survey the wide variety of feasible approaches and to treat each in the depth that we have attempted here.

Parameter estimation plays a major role in current flight test programs. Parameter estimation results are considered a major test objective required for successful completion of flight test programs. Although parameter estimation research continues in several areas, it is no longer purely a research topic.

In this document we first derived the aircraft equations of motion as used for stability and control derivative estimation. Successful application of parameter estimation techniques requires careful consideration of what effects to include in the equations. Such complications as pilot and control system models are peripheral to the purpose of estimating aerodynamic characteristics; the most successful models avoid such peripheral issues where possible.

We next discussed simplifying the general equations of motion and adapting them to a specific computer program. Although the precise details may vary for different programs, the process we examined is always important. Numerous attempted applications have floundered because of such mundane issues as the treatment of biases. We presented an idealized example using computed data, simplified in order to help gain an intuitive grasp of the estimation process.

We discussed the role of aircraft mass distribution data and predicted derivative data in the parameter estimation process. We particularly emphasized the effects of errors in these data and the various requirements for different objectives.

The design of flight test maneuvers involves trade-offs between many conflicting criteria. We discussed maneuver size, flight safety considerations, and pilot involvement. We examined identifiability as a practical issue involving independent inputs, modal excitation, and frequency content, more than as an abstract notion relating to singularity of the information matrix. We also considered flight scheduling issues and allowances for instrumentation failures and other contingencies.

We discussed the data acquisition system in detail. This discussion included consideration of time tagging, frequency aliasing, filtering, sample rates, and resolution. We also discussed the various sensors, their characteristics, and their relative merits.

Finally, we considered evaluation of the parameter estimation results. The tools used for evaluation include the Cramér-Rao bounds, scatter, correlations, bias, and engineering judgment. We discussed the roles of these various tools, their relative merits, and their interplay. We specifically emphasized the use of the Cramér-Rao bound as the primary analytic measure of accuracy, while recognizing that no analytic tool can substitute for the application of sound engineering judgment.

We also emphasized the iterative nature of the evaluation process. Evaluation of the results must be accompanied by a consideration of how they might be improved — by better instrumentation, modeling, maneuvers, analysis methods, or whatever. To obtain good results, it is mandatory to thoroughly review the entire estimation process for errors of equipment, judgment, or fact. Without such a conscious review, it is almost certain that some such errors will exist, corrupting the validity of the results.

BIBLIOGRAPHY

Anon.: Total In-Flight Simulator (TIFS) — Preliminary Design Report. Technical Report AFFDL-TR-71-119, pp. 25-26, Aug. 1971.

Anon.: Parameter Estimation Techniques and Applications in Aircraft Flight Testing. NASA TN D-7647, 1974.

Abbot, W.Y.; Boirun, B.H.; Hill, G.E.; and Tavares, E.J.: Flight Evaluation Pacer Systems Low-Range Airspeed System LORAS 1000, Final Report. US AAEFA Project No. 75-17-1, Edwards Air Force Base, California, May 1977.

Acton, Forman S.: Numerical Methods that Work. Harper & Row, New York, 1970.

AGARD: AGARD Flight Test Instrumentation Series, vols. 1-15. AGARD-AG-160, 1972-1983.

AGARD: Methods for Aircraft State and Parameter Identification. AGARD-CP-172, May 1975.

AGARD: The Effects of Buffeting and Other Transonic Phenomena on Maneuvering Combat Aircraft. AGARD-AR-82, July 1975.

AGARD: Stall/Spin Problems of Military Aircraft. AGARD-CP-199, June 1976.

AGARD: Flight/Ground Testing Facilities Correlation. AGARD-CP-187, 1976.

AGARD: Flight Test Techniques — of Aircraft and Weapon Systems Control. AGARD-CP-223, 1977.

AGARD: Dynamic Stability Parameters. AGARD-CP-235, 1978.

AGARD: High Angle of Attack Aerodynamics. AGARD-CP-247, Jan. 1979 (a).

AGARD: Parameter Identification. AGARD-LS-104, 1979 (b).

AGARD: Aerodynamic Characteristics of Controls. AGARD-CP-262, May 1979 (c).

AGARD: Dynamic Stability Parameters. AGARD-LS-114, 1981.

AGARD: Ground/Flight Test Techniques and Correlation. AGARD-CP-339, 1983.

A'Harrah, R.C.; Lamanna, W.J.; and Hodgkinson, J.: Are Today's Specifications Appropriate for Tomorrow's Aeroplanes? Stability and Control, AGARD-CP-260, Paper 23, May 1979.

Akaike, Hirotugu: A New Look at Statistical Model Identification. IEEE Trans. Automat. Control, vol. AC-19, no. 6, pp. 716-723, Dec. 1974.

Anglin, E.: Aerodynamic Characteristics of Fighter Configurations During Spin Entries and Developed Spins. J. Aircraft, vol. 15, no. 11, pp. 769-776, Nov. 1978.

Aoki, Masanao: Optimization of Stochastic Systems; Topics in Discrete-Time Systems. Academic Press, New York, 1967.

Aoki, Masanao; and Staley, R.M.: On Input Signal Synthesis in Parameter Identification. Preprints Fourth International Federation of Automatic Control Congress, Warsaw, Poland, 1969.

Apostol, Tom M.: Calculus: Volume II. Second ed. Xerox College Publishing, Waltham, Massachussetts, 1969.

Ash, Robert B.: Basic Probability Theory. John Wiley & Sons, New York, 1970.

Åstrom, Karl J.: Introduction to Stochastic Control Theory. Academic Press, New York, 1970.

Åstrom, K.J.; and Eykhoff, P.: System Identification — a Survey. Automatica, vol. 7, no. 2, pp. 123-162, Mar. 1970.

Attlfellner, S.; and Rade, M.: BO 105 In-Flight Simulator for Flight Control and Guidance Systems. First European Rotorcraft and Powered Lift Aircraft Forum, Southhampton, Great Britain, Sept. 1975.

Bach, R.E.; and Wingrove, R.C.: Applications of State Estimation in Aircraft Flight Data Analysis. AIAA Paper 83-2087, Aug. 1983.

Balakrishnan, A.V.: Communication Theory. McGraw-Hill Book Co., New York, 1968.

Balakrishnan, A.V.: Introduction to Optimization Theory in a Hilbert Space. Springer-Verlag, Berlin, 1971.

Balakrishnan, A.V.: Stochastic Differential Systems I. Filtering and Control; A Function Space Approach. M. Beckmann, G. Goos, and H.P. Kunzi, eds., Springer-Verlag, Berlin, 1973. (Lecture Notes in Economics and Mathematical Systems, vol. 84.)

Balakrishnan, A.V.: Stochastic Filtering and Control. Optimization Software, Inc., Los Angeles, 1981.

Balakrishnan, A.V.: Kalman Filtering Theory. Optimization Software, Inc., New York, 1984.

Balakrishnan, A.V.; and Peterka, V.: Identification in Automatic Control System, Automatica, vol. 5, pp. 817-829, Nov. 1969.

Banerjee, D.; and Hohenemser, K.H.: Optimum Data Utilization for Parameter Identification with Application to Lifting Rotors. J. Aircraft, vol. 13, pp. 1014-1016, Dec. 1976.

Bard, Yonathon: Comparison of Gradient Methods for the Solution of Nonlinear Parameter Estimation Problems. SIAM J. Numer. Anal., vol. 7, no. 1, pp. 157-186, Mar. 1970.

Barnard, G.A.: Thomas Bayes – A Biographical Note. Biometrika, vol. 45, pp. 293-295, 1958.

Batterson, James G.: Analysis of Oscillatory Motion of a Light Airplane at High Values of Lift Coefficient. NASA TM-84563, 1983.

Bayes, Thomas: An Introduction to the Doctrine of Fluxions, and a Defence of the Mathematicians Against the Objections of the Author of "The Analyst." John Noon, London, 1736. (For additional information see Barnard, 1958.)

Bendix Corp.: ACIP Error Correction Models Final Report. BSR 4426, Contract NAS9-15588, Bendix Corporation, Ann Arbor, Michigan, Oct. 1980.

Bennett, G.E.; Mace, W.D.; and Pool, A.: Magnetic Recording of Flight Test Data. Flight Test Instrumentation Series, AGARD-AG-160, vol. 5, Feb. 1974.

Bennett, Robert M.; Farmer, Moses G.; Mohr, Richard L.; and Hall, W. Earl, Jr.: A Wind-Tunnel Technique for Determining Stability Derivatives from Cable-Mounted Models. J. Aircraft, vol. 15, no. 5, pp. 304-310, May 1978.

Best, M.R.: Three-Dimensional Flight-Path Reconstruction by Means of Spline Approximation. National Aerospace Laboratory, NLR TR 83091 U, The Netherlands, 1983.

Bierman, Gerald J.: Factorization Methods for Discrete Sequential Estimation. Academic Press, New York, 1977. (Mathematics in Science and Engineering, vol. 128.)

Box, George E.P.; and MacGregor, John F.: The Analysis of Closed-Loop Dynamic-Stochastic Systems. Technometrics, vol. 16, no. 3, pp. 391-398, Aug. 1974.

Box, G.E.P.; and MacGregor, J.F.: Parameter Estimation with Closed-Looped Operating Data. Proc. Sixth International Federation of Automatic Control, Triennial World Congress, Boston and Cambridge, Massachussetts, 1975, Paper 38-5, 1975.

Bramwell, A. R. S.: Helicopter Dynamics. Edward Arnold Publishers, London, 1976.

Brauer, Fred; and Nohel, John A.: Qualitative Theory of Ordinary Differential Equations; an Introduction. W.A. Benjamin, New York, 1969.

Breeman; J.H.: Parameter Identification Results of Tests in Non-Steady Symmetric Flight with the Hawker Hunter Mk 7. National Aerospace Laboratory NLR TR 83042, The Netherlands, 1983.

Breeman, J.H.; and Simons, J.L.: Evaluation of a Method to Extract Performance Data from Dynamic Maneuvers for a Jet Transport Aircraft. Proc. Eleventh Congress of the International Council of the Aeronautical Sciences, Lisbon, Sept. 1978, vol. 1, pp. 475-488, International Council of the Aeronautical Sciences Secretariat (DGLR), Cologne, 1978.

Breeman, J.H; Erkelens, L.J.J.; and Nieuwpoort, A.M.H.: Determination of Performance and Stability Characteristics from Dynamic Manoeuvers with a Transport Aircraft Using Parameter Identification Techniques. National Aerospace Laboratory NLR MP 84024, The Netherlands, 1984.

Breeman, J.H.: Van Woerkom, K.; Jonkers, H.L.; and Mulder, J.A.: Aspects of Flight Test Instrumentation. Parameter Identification, AGARD-LS-104, Paper 4, Oct. 1979.

Brenner, Martin J.; Iliff, Kenneth W.; and Whitman, Robert K.: Effect of Sampling Rate and Record Length on the Determination of Stability and Control Derivatives. NASA TM-72858, 1978.

Briczinski, S.J., and Cooper, D.E.: Flight Investigation of Rotor/Vehicle State Feedback. NASA CR-132546, 1975.

Brockett, Roger W.: Finite Dimensional Linear Systems. John Wiley & Sons, New York, 1970.

Bryant, W.H.; and Hodge, W.F.: A Monte Carlo Analysis of the Effects of Instrumentation Errors on Aircraft Parameter Identification. Methods for Aircraft State and Parameter Identification, AGARD-CP-172, Paper 5, May 1975.

Bryson, Arthur E., Jr.; and Ho, Y. Ch.: Applied Optimal Control, Optimization, Estimation and Control. Ginn and Company, Lexington, Massachusetts, 1969.

Bucy, Richard S.; and Joseph, Peter D.: Filtering for Stochastic Processes with Applications to Guidance. Interscience Publishers, New York, 1968.

Burton, R.A.: Advancement in Parameter Identification and Aircraft Flight Testing. Methods for Aircraft State and Parameter Identification, AGARD-CP-172, Paper 15, May 1975.

Caines, Peter E.; and Chan, C.W.: Feedback Between Stationary Stochastic Processes. IEEE Trans. Automat. Control, vol. AC-20, pp. 498-508, Aug. 1975.

Caines, P.E.; and Sinha, S.: An Application of the Statistical Theory of Feedback to Power System Identification. Proc. 1975 IEEE Conference on Decision and Control, Including the 14th Symposium on Adaptive Processes, Houston, pp. 584-589, Dec. 1975.

Cannaday, R.L.; and Suit: W.T.: Effects of Control Inputs on the Estimation of Stability and Control Parameters of a Light Airplane. NASA TP-1043, 1977.

Carter, J.: The Measurement of Helicopter Air Data Using a Swivelling Pitot-Static Pressure Probe. Air Data Symposium, Naval Postgraduate School, Monterey, June 1976; Proceedings, Session V: Helicoptor/VSTOL Sensors, Monterey, California, 1976.

Chalk, C.R.: Technical Evaluation Report on the Flight Mechanics Panel Symposium on Stability and Control. AGARD-AR-134, Jan. 1979.

Charon, W.; and Verbrugger, R.A.: Nouvelle Technique d'essais sur Maquettes Libres Pour la Determination de Caractéristiques Aérodynamiques. Dynamic Stability Parameters, AGARD-CP-235, Paper 2, 1978.

Chen, R.T.N.: Input Design for Aircraft Parameter Identification: Using Time-Optimal Control Formulation. Methods for Aircraft State and Parameter Identification, AGARD-CP-172, Paper 13, May 1975.

Chen, R.T.N.; Eulrich, B.J.; and Lebecqz, J.V.: Development of Advanced Techniques for the Identification of V/STOL Aircraft Stability and Control Parameters. Cornell Aeronautical Lab., Inc., Rept. No. CAL-BM-2820-F-1, AD-730121, 1971.

Clark, Daniel C.; and Droll, John: General Purpose Airborne Simulator-Conceptual Design Report. NASA CR-544, 1966.

Corbin, M.: A Preliminary Study of an Adaptive Flight Control System Using In-Flight Identification of Three Aircraft Coefficients. RAE TR 75012, Mar. 1975.

Cox, A.B.; and Bryson, A.E., Jr.: Identification by a Combined Smoothing Nonlinear Programming Algorithm. Automatica, vol. 16, no. 6, pp. 689-694, Nov. 1980.

Cramér, Harald. Mathematical Methods of Statistics. Princeton University Press, Princeton, New Jersey, 1946.

Davidon, W.C.: Variable Metric Methods for Minimization. A.E.C. Research and Development Report ANL-5990, Rev. Argonne National Lab., Argonne, Illinois, 1959.

Defalque, B.; Gevers, M.; and Installé, M.: Combined Identification of the Input-Output and Noise Dynamics of a Closed-Loop Controlled Linear System. Int. J. Control, vol. 24, no. 3, pp. 345-360, Sept. 1976.

Dennis, J.E., Jr.: Non-Linear Least-Squares and Equations. Conference on the State of the Art in Numerical Analysis, York, 1976, D. Jacobs, ed., The State of the Art of Numerical Analysis, pp. 269-312. Academic Press, New York, 1977.

De Leo, R.V.; and Jensen, D.P.: Low Range Orthogonal Airspeed System. Air Data Symposium, Naval Postgraduate School, Monterey, June 1976; Proceedings, Session V: Helicopter/VSTOL Sensors, Monterey, California, 1976.

Derbyshire, Thomas; and Sidwell, Kenneth W.: PAN Air Summary Document (Version 1.0). NASA CR-3250, 1982.

Determann, O.: Ermittlung von dynamischen Derivativen der Längs-und Seitenbewegung mit der mobilen oszillierenden Derivativwaage und systematische Untersuchungen zum Einfluss eininger Parameter auf die Ergebinisse. DGLR-78-115, Jahrestagung der DGLR, Darmstadt, 19-23 Sept. 1978.

Deutsch, R.: Estimation Theory. Prentice Hall, Englewoods Cliffs, New Jersey, 1965.

Dixon, L.C.W.: Nonlinear Optimization. Crane, Russak & Co., New York, 1972.

Doetsch, Karl H.: The Time Vector Method for Stability Investigations. A.R.C. R. & M. 2945, Aug. 1953.

Dongarra, J.J.; Moler, C.B.; Bunch, J.R.; and Stewart, G.W.: LINPACK Users' Guide. Society for Industrial and Applied Mathematics, Philadelphia, 1979.

DuBro, G.A.; Kim, D.G.; and Rudd, M.J.: An Electro-Optic Airspeed Sensor. Air Data Symposium, Naval Postgraduate School, Monterey, June 1976; Proceedings, Session IV: Helicopter/VSTOL Sensors, Monterey, California, 1976. (Also published in Electro-Optical Systems Design Conference and International Laser Exposition, New York, Sept. 14-16, 1976; Proc. Technical Program, pp. 384-392, Chicago, 1976.)

Edwards, John W.; Deets, Dwain A.: Development of a Remote Digital Augmentation System and Application to a Remotely Piloted Research Vehicle. NASA TN D-7941, 1975.

Ericsson, L.E.: Technical Evaluation Report on the Fluid Dynamics Panel Symposium on Dynamic Stability Parameters. AGARD-AR-137, Apr. 1979.

Ericsson, L.E.: Support Interference. Dynamic Stability Parameters, AGARD-LS-114, Paper 8, May 1981.

Etkin, Bernard: Dynamics of Flight. John Wiley & Sons, New York, 1959.

Eulrich, B.J.; and Rynaski, E.G.: Identification of Nonlinear Aerodynamic Stability and Control Parameters at High Angle of Attack. Methods for Aircraft State and Parameter Identification, AGARD-CP-172, pp. 2-1 to 2-15, 1975.

Eulrich, B.J.; Govindaraj, K.S.; and Harrington, W.W.: Estimation of the Aerodynamic Stability and Control Parameters for the F-106A Aircraft from Flight Data: Maneuver Design and Flight Data Analysis. AIAA Paper 78-1326, 1978.

Eykhoff, Pieter: System Identification, Parameter and State Estimation. John Wiley & Sons, London, 1974.

Fagin, S.L.: Recursive Linear Regression Theory, Optimal Filter Theory and Error Analysis of Optimal Systems. IEEE International Convention Record, Part I, Session 30, pp. 216-240, 1964.

Feik, R.A.: On the Application of Compatibility Checking Techniques to Dynamic Flight Test Data. Department of Defence Aerodynamics Report 161, Defence Science and Technology Organisation, Aeronautical Research Laboritories, Melbourne, Victoria, Australia, 1984.

Ferguson, Thomas S.: Mathematical Statistics: A Decision Theoretic Approach. Academic Press, New York, 1967.

Findlay, J.T.: ACIP Performance Assessment During STS-1 Re-Entry — Comparisons with IMU Measurements and Trajectory Prediction Considerations. AMA Rep. No. 81-26, Contract NAS1-16087, Sept. 1981.

Fink, Marvin P.; and Freeman, Delma C., Jr.: Full-Scale Wind-Tunnel Investigation of Static Longitudinal and Lateral Characteristics of a Light Twin-Engine Airplane. NASA TN D-4983, 1969.

Fisher, R.A.: On the Mathematical Foundations of Theoretical Statistics. Philos. Trans. Roy. Soc. London, vol. 222, pp. 309-368, 1922.

Fisher, Ronald A.: Contributions to Mathematical Statistics. John Wiley & Sons, New York, 1950.

Fiske, P.H.; and Price, C.F.: A New Approach to Model Structure Identification. AIAA Paper 77-1171, 1977.

Flack, Nelson D.: AFFTC Stability and Control Technique. AFFTC-TN-59-21, Edwards, California, 1959.

Fletcher, R.; and Powell, M.J.D.: A Rapidly Convergent Descent Method for Minimization. Comput. J., vol. 6, no. 2, pp. 163-168, July 1963.

Foster, G.W.: Identification of Hunter Mk 12 Longitudinal and Lateral Aerodynamic Stability and Control Derivatives. RAE TR 77009, 1977.

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

Foster, G.W.: Wind Tunnel Calibration of An Aircraft Pitot-Static and Flow-Direction Measuring Nose Probe. RAE TR 80077, 1980.

Foster, G.W.: The Identification of Aircraft Stability and Control Parameters in Turbulence. RAE TR 83025, 1982. (Also available as Ph.D. Thesis, California Insitute of Technology, Pasadena, California, 1982.)

Froisy, J.B.: Smith, C.L.; Corripio, A.B.; and Murrill, P.W.: Closed-Loop Identification of System Dynamics in the Presence of Noise and Unmeasured Disturbances, Proc. 15th Joint Automatic Control Conference, Austin, Texas, pp. 293-302, 1974.

Fu, K.H.; and Marchand, M.: Helicopter System Identification in the Frequency Domain. Ninth European Rotorcraft Forum, Paper 96, Stresa, Italy, September 13-15, 1983.

Gainer, Thomas G.; and Hoffman, Sherwood: Summary of Transformation Equations and Equations of Motion Used in Free-Flight and Wind Tunnel Data Reduction and Analysis. NASA SP-3070, 1972.

Galbraith, T.J.; and Petersen, T.J.: Nonlinear Parameter Identification and Its Application to Transport Aircraft. Dynamic Stability Parameters, AGARD-CP-235, Paper 18, 1978.

Garbow, B.S.; Boyle, J.M.; Dongarra, J.J.; and Moler, C.B.: Matrix Eigensystem Routines — EISPACK Guide Extension. Springer-Verlag, Berlin, 1977. (Lecture notes in Computer Science, 51.)

Garretson, H.C., III: Beaver Aircraft Parameter Identification Technical Preparations and Technical Preparations and Preliminary Results. DFVLR-Mitt. 78-01, DEVLR Institut Für Flugmechanik, Braunschweig, Federal Republic of Germany, July 1978.

Garretson, H.C., III: Beaver Aircraft Parameter Identification-Technical Preparations and Preliminary Results. DFVLR-Mitt. 78-01, DELVR Instit Für Flugmechanik, Braunschweig, Federal Republic of Germany, Aug. 1978.

Gauss, Karl Friedrich: Theory of the Motion of the Heavenly Bodies Moving About the Sun in Conic Sections; a Translation of Theoria Motus. Translated and with an Appendix by Charles Henry Davis in 1857. Dover Publications, New York, 1963.

Gerlach, O.H.: Determination of Performance, Stability and Control Characteristics from Measurements in Non-Steady Manoeuvers. Stability and Control, Part 1, AGARD-CP-17, pp. 499-523, Sept. 1966.

Gerlach, O.H.: Determination of Performance and Stability Parameters from Nonsteady Flight Test Maneuvers. SAE National Business Aircraft Meeting, Wichita, Kansas, SAE Paper 700236, March 18-20, 1970.

Gerlach, O.H.: The Determination of Stability Derivatives and Performance Characteristics from Dynamic Manoeuvres. Flight Test Techniques, AGARD CP-85, Paper 16, 1971.

Gessow, Alfred; and Myers, Garry C.: Aerodynamics of the Helicopter. Frederick Ungar Publishing Co., New York, 1967.

Geyser, Lucille C.; and Lehtinen, Bruce: Digital Program for Solving the Linear Stochastic Optimal Control and Estimation Problem. NASA TN D-7820, Mar. 1975.

Gibson, J.C.: Flying Qualities and the Fly-By-Wire Aeroplane. Stability and Control, AGARD-CP-260, Paper 22, May 1979.

Gill, Philip E.; and Murray, Walter: Nonlinear Least Squares and Nonlinearly Constrained Optimization. Dundee Conference on Numerical Analysis, 1975; Numerical Analysis, pp. 134-147, Springer-Verlag, Berlin, 1976. (Lecture Notes in Mathematics, No. 506.)

Gilyard, Glenn B.: Determination of Propulsion-System-Induced Forces and Moments of a Mach 3 Cruise Aircraft. Parameter Estimation Techniques and Application in Aircraft Flight Testing, NASA TN D-7647, pp. 369-374, 1974.

Gilyard, Glenn B.; and Belte, Daumants: Flight-Determined Lag of Angle-of-Attack and Angle-of-Sideslip Sensors in the YF-12A Airplane From Analysis of Dynamic Maneuvers. NASA TN D-7819, 1974.

Golub, G.H.: Matrix Decompositions and Statistical Calculations. In R.D. Milton and John A. Nelder (eds.), Statistical Computation, pp. 365-397, Academic Press, New York, 1969.

Golub, G.H.; and Pereyra, V.: The Differentiation of Pseudo-Inverses and Nonlinear Least Squares Problems Whose Variables Separate. SIAM J. Numer. Anal., vol. 10, no. 2, pp. 413-432, Apr. 1973.

Goodson, R.E.; and Polis, M. (eds.): Identification of Parameters in Distributed Systems. Proc. Joint Automatic Control Conference, Austin, Texas, June 17-21, p. 155, 1974.

Goodwin, Graham C.: An Overview of the System Identification Problem — Part 4: Experiment Design. Sixth IFAC Symposium on Identification and System Parameter Estimation, Washington, D.C., June 1982, G.A. Bekey and G.N. Saridis, eds., Pergamon Press, Oxford, 1983.

Goodwin, Graham C.; and Payne, Robert L.: Dynamic System Identification: Experiment Design and Data Analysis. Academic Press, New York, 1977.

Goodwin, G.C.; Payne, R.L.; and Murdoch, J.C.: Optimal Test Signal Design for Linear Single Input-Single Output Closed Loop Identification. Conference on Computer Aided Control System Design, University of Cambridge, IEE Publication No. 96, Institution of Electrical Engineers, London, 1973.

Gould, D.G.; and Hindson, W.S.: Estimates of the Lateral-Directional Stability Drivatives of a Helicopter from Flight Measurements. National Research Council Canada, Aeronautical Report LR-572, NRC-13882, Dec. 1973.

Gould, D.G.; and Hindson, W.S.: Estimates of the Stability Derivatives of a Helicopter from Flight Measurements. Ninth Congress of the International Council of the Aeronautical Sciences, Haifa, Israel, Aug. 25-30, IGAS Paper 74-49, 1974.

Gould, D.G.; and Hindson, W.S.: Estimates of the Stability Derivatives of a Helicopter and a V/STOL Aircraft from Flight Data. Methods for Aircraft State and Parameter Identification, AGARD-CP-172, pp. 23-1 to 23-9, May 1975.

Gracey, William: Measurement of Aircraft Speed and Altitude. NASA RP-1046, 1980.

Graupe, Daniel: On Identifying Stochastic Closed-Loop Systems. IEEE Trans. Automat. Control, vol. AC-20, pp. 553-555, Aug. 1975.

Greenberg, Harry: A Survey of Methods for Determining Stability Parameters of an Airplane from Dynamic Flight Measurements. NACA TN-2340, 1951.

Greenstadt, John: On the Relative Efficiencies of Gradient Methods. Math. Comp., vol. 21, no. 99, pp. 360-367, July 1967.

Grove, R.D.; Bowles, R.L.; and Mayhew, S.C.: A Procedure for Estimating Stability and Control Parameters from Flight Test Data by Using Maximum Likelihood Methods Employing a Real-Time Digital System. NASA TN D-6735, 1972.

Gupta, Narendra K.: New Frequency Domain Methods for System Identification. Proc. Joint Automatic Control Conference, San Francisco, June 22-24, 1977, vol. 2, pp. 804-808, Institute of Electrical and Electronic Engineers, New York, 1977.

Gupta, Narendra K.; and Hall, W. Earl, Jr.: Design and Evaluation of Sensor Systems for State and Parameter Estimation. J. Guid. Control, vol. 1, no. 6, pp. 397-403, Nov.-Dec. 1978.

Gupta, Naren K.; and Iliff, Kenneth W.: Identification of Aerodynamic Indicial Functions Using Flight Data. AIAA Paper 82-1375, Aug. 1982.

Gupta, N.K.; and Mehra, R.K: Computational Aspects of Maximum Likelihood Estimation and Reduction in Sensitivity Function Calculations. IEEE Trans. Automat. Control, vol. AC-19, no. 6, pp. 774-783, Dec. 1974.

Gupta, Narendra K.; Hall, W. Earl, Jr.; and Trankle, Thomas L.: Advanced Methods of Model Structure Determination from Test Data. J. Guid. Control, vol. 1, no. 3, pp. 197-204, May-June 1978.

Hafer, X.: Wind Tunnel Testing of Dynamic Derivatives in W. Germany. Dynamic Stability Parameters, AGARD-CP-235, Paper 5, 1978.

Hajdasinski, A.K.; Eykhoff, P.; Damen, A.A.H.; and van den Boom, A.J.W.: The Choice and Use of Different Model Sets for System Identification. Sixth IFAC Symposium on Identification and System Parameter Identification, Washington, D.C., June 7-11, 1982, pp. 47-55, Pergamon Press, Oxford and New York, 1983.

Hall, W.E.; and Bryson, A.E.: Inclusion of Rotor Dynamics in Controller Design for Helicopters. J. Aircraft, vol. 10, no. 4, pp. 200-206, Apr. 1973.

Hall, W.E., Jr.; Gupta, N.K.; and Hansen, R.S.: Rotorcraft System Identification Techniques for Handling Qualities and Stability and Control Evaluation. Preprint No. 78-30, Proc. 34th Annual Forum of the American Helicopter Society, Washington, D.C., May 1978.

Hall, W.E., Jr.; Gupta, N.K.; and Tyler, J.S.: Model Structure Determination and Parameter Identification for Nonlinear Aerodynamic Flight Regimes. Methods for Aircraft State and Parameter Identification, AGARD-CP-172, Paper 21, 1974.

Hamel, P.G.: A Systems Analysis View of Aerodynamic Coupling. J. Aircraft, vol. 7, no. 6, pp. 567-569, Nov.-Dec. 1970.

Hamel, P.G.: Aircraft Parameter Identification Methods and Their Applications – Survey and Future Aspects. Parameter Identification, AGARD-LS-104, Paper 1, 1979.

Hamel, P.G.: Determination of Aircraft Dynamic Stability and Control Parameters from Flight Testing. Dynamic Stability Parameters, AGARD-LS-114, Paper 10, May 1981.

Hamel, P.G.; and Krag, B.: Dynamic Windtunnel Simulation of Active Control Systems. Stability and Control, AGARD-CP-260, Paper 16A, May 1979.

Herrington, R.M.; Shoemacher, R.E.; Bartlett, E.P.; and Dunlap, E.W.: Flight Test Engineering Handbook. AFFTC-TR-6273, Edwards, California, AD-636392, Jan. 1966.

Hodge, Ward F.; and Bryant, Wayne H.: Monte Carlo Analysis of Inaccuracies in Estimated Aircraft Parameters Caused by Unmodeled Flight Instrumentation Errors. NASA TN D-7712, 1975.

Hodgkinson, J.; LaManna, W.J.; and Heyde, J.L.: Handling Qualities of Aircraft with Stability and Control Augmentation Systems — A Fundamental Approach. Aeronaut. J., vol. 80, no. 782, pp. 75-81, Feb. 1976.

Hohenemser, K.H.; Banerjee, D.; and Yin, S.K.: Rotor Dynamic State and Parameter Identification from Simulated Forward Flight Transients. Part 1: NASA CR-137963; Part 2: NASA CR-137964, June 1976.

Hohenemser, K.H.; and Banerjee, D.: Application of System Identification to Analytical Rotor Modeling from Simulated and Wind Tunnel Dynamic Test Data, Part 2. NASA CR-152023, 1977.

Hohenemser, K.H.; and Crews, Sam T.: Unsteady Hovering Wake Parameters Identified from Dynamic Model Tests, Part 1. NASA CR-152022, 1977.

Hohenemser, K.H.; and Yin, S.K.: On the Use of First Order Rotor Dynamics in Multiblade Coordinates. Preprint No. 831, Proc. 30th Annual National Forum of the American Helicopter Society, Washington, D.C., May 7-9, 1974.

Holleman, Euclid C.: Summary of Flight Tests to Determine the Spin and Controllability Characteristics of a Remotely Piloted, Large-Scale (3/8) Fighter Airplane Model. NASA TN D-8052, 1976.

Horsten, J.J.; Jonkers, H.L.; and Mulder, J.A.: Flight Path Reconstruction in the Context of Nonsteady Flight Testing. Technische Hogeschool, Delft (Netherlands), Report LR-280, May 1979.

Hosman, R.J.A.W.: A Method to Derive Angle of Pitch, Flight-Path Angle and Angle of Attack from Measurements in Nonsteady Flight. Technische Hogeschool, Delft (Netherlands), Report VTH-156, Apr. 1971.

Hosman, R.J.A.W.: Advanced Flight Test Instrumentation: Design and Calibration. Technische Hogeschool, Delft (Netherlands), Report VTH-M-222, Oct. 1974.

Hosman, R.J.A.W.: Advanced Flight Test Instrumentation: Design and Calibration. Methods for Aircraft State and Parameter Identification, AGARD-CP-172, Paper 6, May 1975.

Howard, J.: The Determination of Lateral Stability and Control Derivatives from Flight Data. Can. Aeronautics Space J., vol. 13, no. 3, pp. 126-134, Mar. 1967.

Hsia, Tien C.: System Identification: Least-Squares Methods. D.C. Heath and Co., Lexington, Massachusetts, 1977.

Huber, H.: Parametric Trends and Optimization — Preliminary Selection of Configuration — Prototype Design and Manufacture. Helicopter Aerodynamics and Dynamics, AGARD-LS-63, Paper 7, Apr. 1973.

Huber, H.: Effect of Torsion-Flap-Lag Coupling on Hingeless Rotor Stability. 29th Annual National Forum of the American Helicopter Society, Washington, D.C., AHS Reprint 731, May 9-11, 1973.

Iliff, Kenneth W.: Identification and Stochastic Control with Application to Flight Control in Turbulence. Ph.D. Dissertation, Univ. of California, Los Angeles, May 1973.

Iliff, Kenneth W.: Maximum Likelihood Estimation of Lift and Drag from Dynamic Aircraft Maneuvers. J. Aircraft, vol. 14, no. 12, pp. 1175-1181, Dec. 1977.

Iliff, Kenneth W.: Identification and Stochastic Control of an Aircraft Flying in Turbulence. J. Guid. Control, vol. 1, no. 2, pp. 101-108, Mar.-Apr. 1978 (a).

Iliff, Kenneth W.: Estimation of Aerodynamic Characteristics from Dynamic Flight Test Data. Dynamic Stability Parameters, AGARD-CP-235, Paper 15, Nov. 1978 (b).

Iliff, Kenneth W.: Aircraft Identification Experience. Parameter Identification, AGARD-LS-104, Paper 6, 29 Oct.-2 Nov. 1979.

Iliff, Kenneth W.; and Maine, Richard E.: Practical Aspects of Using a Maximum Likelihood Estimator. Methods for Aircraft State and Parameter Identification, AGARD-CP-172, Paper 16, May 1975.

Iliff, Kenneth W.; and Maine, Richard E.: Further Observations on Maximum Likelihood Estimates of Stability and Control Characteristics Obtained from Flight Data. AIAA Paper 77-1133, Hollywood, Florida, Aug. 1977.

Iliff, Kenneth W.; and Maine, Richard E.: NASA Dryden's Experience in Parameter Estimation and Its Use in Flight Test. AIAA Paper 82-1373, AIAA AFM Conference, San Diego, California, Aug. 1982.

Iliff, Kenneth W.; and Taylor, Lawrence W., Jr.: Determination of Stability Derivatives from Flight Data Using a Newton-Raphson Minimization Technique. NASA TN D-6579, 1972.

Iliff, Kenneth W.; Maine, Richard E.; and Shafer, Mary: Subsonic Stability and Control Derivatives for an Unpowered, Remotely Piloted 3/8-scale F-15 Airplane Model Obtained from Flight Test. NASA TN D-8136, 1976.

Iliff, Kenneth W.; Maine, Richard E.; and Steers, Sandra Thornberry: Flight-Determined Stability and Control Coefficients of the F-111A Airplane. NASA TM-72851, 1978.

Iliff, Kenneth W.; Maine, Richard E.; and Montgomery, T.D.: Important Factors in the Maximum Likelihood Analysis of Flight Test Maneuvers. NASA TP-1459, 1979.

Iliff, Kenneth W.; Maine, Richard E.; and Cooke, Douglas R.: Selected Stability and Control Derivatives from the First Space Shuttle Entry. AIAA Paper 81-2451. AIAA/SETP/SFTE/SAE First Flight Testing Conference, Nov. 1981.

Jategaonkar, Ravindra; and Plaetschke, Ermin: Maximum Likelihood Parameter Estimation from Flight Test Data for General Non-Linear Systems. DFVLR-FB 83-14, AD-B076999, Apr. 1983.

Jazwinski, Andrew H.: Stochastic Processes and Filtering Theory. Academic Press, New York, 1970.

Jeglum, P.M.: Air Force Flight Test Center Experience in the Identification of Stability and Control Parameters from Dynamic Flight Test Maneuvers. Dynamic Stability Parameters, AGARD-CP-235, Paper 14, Nov. 1978.

Johnston, D.E.: Identification of Key Maneuver-Limiting Factors in High-Angle-of-Attack Flight. Dynamic Stability Parameters, AGARD-CP-235, Paper 36, Nov. 1978.

Jonkers, H.L.: Application of the Kalman Filter to Flight Path Reconstruction from Flight Test Data Including Estimation of Instrumental Bias Error Corrections. Delft Technische Hogeschool, Department of Aerospace Engineering, Delft, Report VTH-162, 1976.

Jonkers, H.L.; and Mulder, J.A.: New Developments and Accuracy Limits in Aircraft Flight Testing. AIAA Aircraft Systems and Technology Meeting, Dallas, Texas, AIAA Paper 76-897, Sept. 1976 (a).

Jonkers, H.L.; and Mulder, J.A.: Accuracy Limits in Nonsteady Flight Testing. Tenth Congress of the International Council of the Aeronautical Sciences, Ottawa, ICAS Paper 76-46, Oct. 1976 (b).

Jonkers, H.L.; Mulder, J.A.; and Horsten, J.J.: Introduction to State Reconstruction of Dynamic Flight Test Manoeuvres. Control and Dynamic Systems, Advances in Theory and Applications, Vol. 21, C.T. Leondes, ed., Academic Press, New York, 1984.

Jonkers, H.L.; Mulder, J.A.; and Woerkom van, K.: Measurements in Nonsteady Flight: Instrumentation and Analysis. Proc. Seventh International Aerospace Instrumentation Symposium, Cranfield, England, March 20-23, 1972, pp.21.1-21.10; Peter Peregrinus, Ltd., Stevonage, Herts, 1972.

Joy, D.: Airspeed and Direction Measurement by Vortex Detection. Air Data Symposium Proceedings, Naval Postgraduate School, Session IV: Helicopter/VSTOL Sensors, Monterey, California, June 1976.

Kagiwada, Harriet H.: System Identification, Methods and Applications. Addison-Wesley Publishing Co., Reading, Massachusetts, 1974.

Kailath, Thomas: Linear Systems. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1980.

Kailath, Thomas: Lectures on Wiener and Kalman Filtering. New, Revised Edition. Springer-Verlag, New York, 1981.

Kailath, T.; and Lyung, L.: Asymptotic Behavior of Constant-Coefficient Riccati Differential Equations. IEEE Trans. Automat. Control, vol. AC-21, pp. 385-388, June 1976.

Kaletka, J.: Rotorcraft Identification Experience. Parameter Identification, AGARD-LS-104, Paper 7, 1979.

Kaletka, J.; and Langer, H.J.: Correlation Aspects of Analytical, Wind Tunnel and Flight Test Results for a Hingeless Rotor Helicopter. Group/Flight Test Techniques and Correlation, AGARD-CP-339, Paper 16, 1982.

Kaletka, J.; and Rix, O.: Aspects of System Identification of Helicopters. Third European Rotorcraft and Powered Lift Aircraft Forum, Aix-en-Provence, France, Paper 15, Sept. 1977.

Kalman, R.E.; and Bucy, R.S.: New Results in Linear Filtering and Prediction Theory. Trans. ASME, Ser. D, J. Basic Eng., vol. 63, pp. 95-107, 1961.

Kaminski, P.G.; Bryson, A.E.; and Schmidt, S.F.: Discrete Square Root Filtering — A Survey of Current Techniques. IEEE Trans. Automat. Control, vol. 16, no. 6, pp. 727-736, Dec. 1971.

Kanning, G.; and Biggers, J.C.: Application of a Parameter Identification Technique to a Hingeless Helicopter Rotor. NASA TN D-7834, 1974.

Kashyap, R.L.: Maximum Likelihood Identification of Stochastic Linear Systems. IEEE Trans. Automat. Control, vol. AC-15, no. 1, pp. 25-34, Feb. 1970.

Klein, V.: Parameter Identification Applied to Aircraft. Cranfield Report Aero No. 26, Cranfield Inst. of Technology, 1974 (a).

Klein, V.: Longitudinal Aerodynamic Derivatives of a Slender Delta-Wing Research Aircraft Extracted from Flight Data. Cranfield Report Aero No. 27, CIT-FI-74-023, Cranfield Inst. of Technology, July 1974 (b).

Klein, V.: On the Adequate Model for Aircraft Parameter Estimation. Cranfield Report Aero No. 28, Cranfield Inst. of Technology, March 1975.

Klein, V.: Determination of Longitudinal Aerodynamic Derivatives From Steady-State Measurement of an Aircraft. AIAA Paper 77-1123, Hollywood, Florida, Aug. 1977.

Klein, V.: Aircraft Parameter Estimation in Frequency Domain. AIAA Paper 78-1344, 1978.

Klein, V.: Determination of Stability and Control Parameters of a Light Airplane From Flight Data Using Two Estimation Methods. NASA TP-1306, 1979.

Klein, V.: Identification Evaluation Methods. Parameter Identification, AGARD-LS-104, Paper 2, 29 Oct.-2 Nov. 1979.

Klein, Vladislav: Maximum Likelihood Method for Estimating Airplane Stability and Control Parameters From Flight Data in Frequency Domain. NASA TP-1637, 1980.

Klein, V.; and Batterson, J.G.: Determination of Airplane Aerodynamic Parameters From Flight Data at High Angles of Attack. Proc. Thirteenth Congress of the International Council of the Aeronautical Sciences, vol. I, B. Laschka and R. Staufenbiel, eds., pp. 467-474, American Institute of Aeronautics and Astronautics, New York, 1982.

Klein, Vladislav; and Batterson, James G.: Determination of Airplane Model Structure From Flight Data Using Splines and Stepwise Regression. NASA TP-2126, 1983.

Klein, Vladislav; and Batterson, James G.: Aerodynamic Characteristics of a Fighter Airplane Determined From Flight and Wind Tunnel Data. NASA TP-2483, 1985.

Klein, V.; and Keskar, A.D.: Frequency Domain Identification of a Linear System Using Maximum Likelihood Estimation. Fifth IFAC Symposium on Identification and System Parameter Estimation, Darmstadt, Sept. 1979. Pergaman Press, Oxford, 1980.

Klein, V.; and Schiess, J.R.: Compatibility Check of Measured Aircraft Responses Using Kinematic Equations and Extended Kalman Filter. NASA TN D-8514, 1977.

Klein, V.; and Williams, D.A.: On Some Problems Related to the Identification of Aircraft Parameters. Proc. Third IFAC Symposium on Identification and System Parameter Estimation, The Hague, June 12-15, 1973, pp. 435-444, North Holland, Amsterdam, 1973.

Klein, Vladislav; Batterson, James G.; and Murphy, Patrick C.: Determination of Airplane Model Structure From Flight Data by Using Modified Stepwise Regression. NASA TP-1916, Oct. 1981.

Klein, V.; Batterson, J.G.; and Smith, P.L.: On the Determination of Airplane Model Structure From Flight Data. Sixth IFAC Symposium on Identification and System Parameter Estimation, vol. 2, George A. Bekey and George N. Saridis, eds., pp. 1034-1039, Peragamon Press, Oxford, 1982.

Kleingeld, H.W.: Design and Evaluation of a Symmetric Flight Test Manoeuver for the Estimation of Longitudinal Performance and Stability and Control Characteristics. Methods for Aircraft State and Parameter Identification, AGARD-CP-172, pp. 9-1 to 9-6, 1974.

Klopfenstein, H.B.: Obtaining Airplane Drag Data from Nonsteady Flight. AIAA Paper 65-211, 1965.

Kloster, M.; Kaletka, J.; and Schaufele, H.: Parameter Identification of a Hingeless Rotor Helicopter in Flight Conditions with Increased Instability. Sixth European Rotorcraft and Powered Lift Aircraft Forum, Bristol, England, Sept. 16-19, 1980, Conference Papers, Part 2, University of Bristol, Bristol, 1980.

Koehler, R.; and Marchand, M.: Open/Closed Loop Identification of Stability and Control Characteristics of Combat Aircraft. Stability and Control, AGARD-CP-260, Paper 16, May 1979.

Koehler, R.; and Wilhelm, K.: Auslegung von Eingangssignalen für die Kennwertermittlung. IB 154-77/40, DFVLR Institut für Flugmechanik, Braunschweig, Federal Republic of Germany, Dec. 1977.

Koehler, R.; and Wilhelm, K.: Closed Loop Aspects of Aircraft Identification. Parameter Identification, AGARD-LS-104, Paper 10, 29 Oct.-2 Nov. 1979.

Krag, B.: The Wind Tunnel Behaviour of a Scaled Model with a Gust Alleviation System in a Deterministic Gust Field. Symposium Dynamic Analysis of Vehicle Ride and Manoeuvring Characteristics, London, Preprint pp. 149-166, Nov. 1978 (a).

Krag, B.: Active Control Technology for Gust Alleviation. Von Karman Institute for Fluid Dynamics, Brussels, Active Control Technology, The Institute, Rhode St. Genese, Belgium (Von Karman Institute for Fluid Dynamics, Lecture Series 1978-1, vol. 2), Dec. 4-8, 1978 (b).

Kurz, H.; and Isermann, R.: Methods for On-Line Process Identification in Closed Loop. Proc. International Federation of Automatic Control Sixth Triennial World Congress, Boston and Cambridge, Massachusetts, Aug. 24-30, 1975, Part 2, pp. 11.3 1 - 11.3 14, Instrument Society of America, Pittsburg, 1975.

Kurzhals, P.R., ed.: Active Controls in Aircraft Design. AGARD-AG-234, Nov. 1978.

Kushner, Harold J.: Introduction to Stochastic Control. Holt, Rinehart and Winston, New York, 1971.

Kwakernaak, Hubert; and Sivan, Raphael: Linear Optimal Control Systems. John Wiley & Sons, New York, 1972.

Lainiotis, Demetrios G., ed.: Estimation Theory. American Elsevier Publishing Co., New York, 1974.

Lawson, Charles L.; and Hanson, Richard J.: Solving Least Squares Problems. Prentice-Hall, Englewood Cliffs, New Jersey, 1974.

Langer, H.J.; and Stricker, R.: Some Results of Dynamic Measurements with a Model Hingeless Rotor. Fifth European Rotorcraft and Powered Lift Aircraft Forum, Amsterdam, Netherlands, Sept. 1979.

Levadi, Victor S.: Design of Input Signals for Parameter Estimation. IEEE Trans. Automat. Control, vol. 11, no. 2, pp. 205-211, 1966.

Levan, N.: Systems and Signals. Optimization Software, Inc., New York, 1983.

Levin, Morris J.: Optimal Estimation of Impulse Response in the Presence of Noise. IRE Trans. Circuit Theory, vol. CT-7, no. 1, pp. 50-56, Mar. 1960.

Lindberger, N.A.: Stochastic Modelling of Computer-Regulated Linear Plants in Noisy Environments. Int. J. Control, vol. 16, no. 6, pp. 1009-1019, 1972.

Lindberger, N.A.: Stochastic Identification of Computer-Regulated Linear Plants in Noisy Environments. Int. J. Control, vol. 17, no. 1, pp. 65-80, 1973.

Liptser, Robert S.; and Shiryayev, A.N.: Statistics of Random Processes. Translated by A.B. Aries, vol. 1: General Theory, Springer-Verlag, New York, 1977.

Litmann, S.; and Huggins, W.H.: Growing Exponentials as a Probing Signal for System Identification. Proc. IEEE, vol. 51, pp. 917-923, June 1963.

Liusternik, Lazar A.; and Sobolev, V.J.: Elements of Functional Analysis. Translated from the Russian by A.E. Labarre, Jr., Herbert Izbicki, and H. Ward Crowley, Frederick Ungar Publishing Co., New York, 1961.

Ljung, L.; Gustavsson, I.; and Soderstrom, T.: Identification of Linear, Multivariable Systems Operating Under Linear Feedback Control. IEEE Trans. Automat. Control, vol. AC-19, pp. 836-840, Dec. 1974.

Lobbia, R.N.; and Saridis, G.N.: On-Line Identification of Multivariable Stochastic Feedback Systems. Thirteenth Joint Automatic Control Conference, Stanford, California, Aug. 16-18, 1972, Preprints of Technical Papers, pp. 802-810, American Institute of Astronautics, New York, 1972,

Luenburger, David G.: Optimization by Vector Space Methods. John Wiley & Sons, New York, 1969.

Luenberger, David G.: Introduction to Linear and Nonlinear Programming. Addison-Wesley, Reading, Massachusetts, 1973.

Lumsden, R.B.: Estimation of Longitudinal Aerodynamic Derivatives of the Comet 3B from Flight Test Data Using an Automatic Model Matching Method. RAE TR 74103, 1974.

Maine, Richard E.: Aerodynamic Derivatives for an Oblique Wing Aircraft Estimated from Flight Data by Using a Maximum Likelihood Technique. NASA TP-1336, 1978.

Maine, Richard E.: User's Manual for SYNC, a FORTRAN Program for Merging and Time Synchronizing Data. NASA TM-81355, 1981 (a).

Maine, Richard E.: Programmer's Manual for MMLE3, A General FORTRAN Program for Maximum Likelihood Parameter Estimation. NASA TP-1690, 1981 (b).

Maine, Richard E.; and Iliff, Kenneth W.: A FORTRAN Program for Determining Aircraft Stability and Control Derivatives From Flight Data. NASA TN D-7831, 1975.

Maine, Richard E.; and Iliff, Kenneth W.: Maximum Likelihood Estimation of Translational Acceleration Derivatives from Flight Data. J. Aircraft, vol. 16, no. 10, pp. 674-679, Oct. 1979.

Maine, Richard E.; and Iliff, Kenneth W.; User's Manual for MMLE3, A General FORTRAN Program for Maximum Likelihood Parameter Estimation. NASA TP-1563, 1980.

Maine, Richard E.; and Iliff, Kenneth W.: Formulation and Implementation of a Practical Algorithm for Parameter Estimation with Process and Measurement Noise. SIAM J. Appl. Math., vol. 41, no. 3, pp. 558-579, Dec. 1981 (a).

Maine, Richard E.; and Iliff, Kenneth W.: The Theory and Practice of Estimating the Accuracy of Dynamic Flight-Determined Coefficients. NASA RP-1077, 1981 (b).

Maine, Richard E.; and Iliff, Kenneth W.: Identification of Dynamic Systems. AGARD-AG-300, 1984.

Marchand, M.: Der Einfluss von Messfehlern auf die Bestimmbarkeit der Kennwerte der Längsbewegung. DFVLR IB 154-74/31, 1974.

Marchand, M.: Untersuchung der Bestimmbarkeit der flugmechanischen Derivative des CCV-Versuchsträgers F-104 G. IB 154-77/12, DFVLR Institut für Flugmechanik, Braunschweig, Federal Republic of Germany, Mar. 1977.

Marchand, M.; and Fu, K.-H.: Frequency Domain Parameter Estimation of Aeronautical Systems Without and With Time Delay. Seventh IFAC/IFORS Symposium on Identification and System Parameter Estimation, York, England, July 3-7, 1985.

Marchand, M.; and Koehler, R.: Determination of Aircraft Derivatives by Automatic Parameter Adjustment and Frequency Response Methods. Methods for Aircraft State and Parameter Identification, AGARD-CP-172, pp. 17-1 to 17-18, May 1975.

Margason, R.J.; and Lamar, J.B.: Vortex Lattice FORTRAN Program for Estimating Subsonic Aerodynamic Characteristics of Complex Planforms. NASA TN D-6142, 1971.

Marquardt, Donald W.: An Algorithm for Least Squares Estimation of Nonlinear Parameters. J. SIAM, vol. 11, pp. 431-441, 1963.

Marquardt, Donald W.: Generalized Inverses, Ridge Regression, Biased Linear Estimation, and Nonlinear Estimation. Technometrics, vol. 12, no. 3, pp. 591-612, Aug. 1970.

Matheny, Neil W.; and Pangeas, George N.: HiMAT Aerodynamic Design and Flight Test Experience. AIAA Paper 81-2433, Nov. 1981.

McCracken, J.R.; Kaletka, J.; Meyer, H.; Rix, O.; and Gartung, B.: CASA C-212 Flight Test and Parameter Identification. Spain-Federal Republic of Germany Cooperative Research Program, DFVLR IB 154-76/16, 1976.

McLaren, I.: Flight Test Instrumentation Series, Open and Closed Loop Accelerometers. AGARD-AG-160, vol. 6, July 1974.

Meditch, J. S.: Stochastic Optimal Linear Estimation and Control. McGraw-Hill Book Co., New York, 1969.

Mehra, R.K.: Identification of Stochastic Linear Dynamic Systems Using Kalman Filter Representation. AIAA J., vol. 9, no. 1, pp. 28-31, Jan. 1971.

Mehra, R.K.: Optimal Inputs for Linear System Identification. Thirteenth Joint Automatic Control Conference, Stanford, California, Aug. 16-18, 1972, Preprints of Technical Papers, pp. 811-820, American Institute of Aeronautics and Astronautics, New York, 1972.

Mehra, R.K.: Identification in Control and Econometrics; Similarities and Differences. Division of Engineering and Applied Physics, Harvard University, Cambridge, Massachusetts, Technical Report 647, AD-767393, 1973.

Mehra, R.K.: Frequency-Domain Synthesis of Optimal Inputs for Linear System Parameter Estimation. Division of Engineering and Applied Physics, Harvard University, Cambridge, Massachusetts, Technical Report 645, July 1973. (Also published in ASME Trans. Ser. G, J. Dynamic Systems, Measurement and Control, vol. 98, June 1976, pp. 130-138. ASME Paper 76-AUT-V.)

Mehra, R.K.: Optimal Input Signals for Parameter Estimation in Dynamic Systems, Survey and New Results. IEEE Trans. Automat. Control, vol. AC-19, no. 6, pp. 753-768, Dec. 1974.

Mehra, R.K.; and Gupta, N.K.: Status of Input Design for Aircraft Parameter Identification. Methods for Aircraft State and Parameter Identification, AGARD-CP-172, Paper 12, May 1975.

Mehra, Raman K.; and Lainiotis, Dimitri G. (eds.): System Identification: Advances and Case Studies. Academic Press, New York, 1976.

Mendel, Jerry M.: Discrete Techniques of Parameter Estimation: The Equation Error Formulation. Marcel Dekker, Inc., New York, 1973.

Moler, C.B.; and Stewart, G.W.: An Algorithm for Generalized Matrix Eigenvalue Problems. SIAM J. Numer. Anal., vol. 10, no. 2, pp. 241-256, Apr. 1973.

Moler, Cleve; and Van Loan, Charles: Nineteen Dubious Ways to Compute the Exponential of a Matrix. SIAM Rev., vol. 20, no. 4, pp. 801-836, Oct. 1978.

Molusis, J.A.: Helicopter Stability Derivative Extraction and Data Processing Using Kalman Filtering Techniques. 28th Annual National Forum of the American Helicopter Society, Washington, D.C., May 17-19, 1972. AHS Preprint No. 641, 1972.

Molusis, J.A.: Analytical Study to Define a Helicopter Stability Derivative Extraction Method, vol. 1: Final Report. NASA CR-132371, 1973.

Molusis, J.A.: Helicopter Stability Derivative Extraction from Flight Data Using the Bayesian Approach to Estimation. J. Am. Helicopter Soc., vol. 18, no. 2, pp. 12-23, Apr. 1973.

Molusis, J.A.: Rotorcraft Derivative Identification from Analytical Models and Flight Test Data. Rotorcraft Parameter Identification, Session V, AGARD-CP-172, Paper 24, May 1975.

Mulder, J.A.: Aircraft Performance Measurements in Nonsteady Flights. Proc. Third International Federation of Automatic Control Symposium on Identification and System Parameter Estimation, Part 2, Delft, The Hague, pp. 1131-1145, June 12-15, 1973.

Mulder, J.A.: Some Aspects of Performance Measurements in Non-steady Flight. Bestimming Von Flugzengkenndaten ans Flugverschen, DLGR-Mitt-73-25 und Raumfahrt, Cologne, 1973. (Translated as Determination of Aircraft Characteristics From Flight Tests, European Space Agency, Paris, pp. 171-190, Nov. 1974.

Mulder, J.A.: Estimation of the Aircraft State in Non-steady Flight. Methods for Aircraft State and Parameter Identification, AGARD-CP-172, Paper 19, May 1975.

Mulder, J.A.: Estimation of Thrust and Drag in Non-steady Flight. Proc. Fourth International Federation of Automatic Control Symposium on Identification and System Parameter Estimation, Tbilisi, Georgian, USSR, Sept. 21-27, 1976.

Mulder, J.A.: The Measurement of Performance, Stability and Control Characteristics of a High Subsonic Jet Aircraft in Non-steady Flight Conditions Employing High Accuracy Instrumentation Techniques. Aeronaut. J., vol. 81, pp. 247-258, June 1977.

Mulder, J.A.; and den Hollander, J.G.: Status of Dynamic Flight Test Technology – Model Identification for Flight Simulation. SAE Business Aircraft Meeting and Exposition, Wichita, Kansas, SAE Paper 810597, Apr. 1981. (Also published in Trans. SAE, Sept. 1982.

Mulder, J.A.; and Hosman, R.J.A.W.: The Application of High Accuracy Instrumentation Techniques to Aircraft Flight Testing. Shell Aviation News, no. 425, 1974.

Mulder, J.A.; and van Sliedregt, J.M.: Estimation of Drag and Thrust of Jet-Propelled Aircraft by Nonsteady Flight Test Maneuvers. AGARD-CP-222, pp. 11-1 to 11-30, 1977.

Mulder, J.A.; Jonkers, H.L.; Horsten, J.J.; Breeman, J.H.; and Simons, J.H.: Analysis of Aircraft Performance, Stability and Control Measurements, Part I: Theoretical Aspects in Parameter Identication. Parameter Identification, AGARD-LS-104, 1979.

Mulder, J.A.; den Hollander, J.G.; and Binkhorst, H.: Aerodynamic Model Identification from Dynamic Flight Test Data and Wind Tunnel Experiments. Ground/Flight Test Techniques and Correlation, AGARD-CP-339, Paper 14, Jan. 1983. (Also published as Report VLR-361, Delft University of Technology, Technische Hogeschool, Delft, Netherlands, Oct. 1982.)

Nahi, N.E.; and Napjus, G.A.: Design of Optimal Probing Signals for Vector Parameter Estimation. Proc. IEEE Conference on Decision and Control, Miami Beach, Florida, Dec. 15-17, 1971, IEEE, New York, pp. 162-168, 1971.

Nahi, N.E.; and Wallis, D.E., Jr.: Optimal Inputs for Parameter Estimation in Dynamic Systems with White Observation Noise. Tenth American Automatic Control Council Joint Automatic Control Conference, University of Colorado, Boulder, Colorado, Preprints of Technical Papers, pp. 506-513, Aug. 5-7, 1969.

Nering, Evar D.: Linear Algebra and Matrix Theory. Second ed. John Wiley & Sons, New York, 1970.

Neveu, Jacques: Mathematical Foundations of the Calculus of Probability. Holden-Day, Inc., San Francisco, 1965.

Nguyen, Luat T.; Gilbert, William P.; Gera, Joseph; Iliff, Kenneth W.; and Enevoldson, Einar K.: Application of High-Alpha Control System Concepts to a Variable-Sweep Fighter Airplane. AIAA Atmospheric Flight Mechanics Conference, Danvers, Massachusetts, Aug. 11-13, 1980, AIAA Paper 80-1582, 1980.

Nippress, K.R.: Performance Estimation from Non-Steady Maneuvers. First AIAA, SETP, SFTE, SAE, ITEA and IEEE Flight Testing Conference, Las Vegas, Nevada, Nov. 11-13, 1981, AIAA Paper 81-2424, 1981.

Norton, Harry N.: Handbook of Transducers for Electronic Measuring Systems. Prentice Hall, Englewood Cliffs, New Jersey, 1969.

Ogata, Katsuhiko: Modern Control Engineering. Prentice-Hall, Englewood Cliffs, New Jersey, 1970.

O'Leary, C.O.: Wind-Tunnel Measurement of Lateral Aerodynamic Derivatlves Using a New Oscillatory Rig and Comparisons for a Gnat Aircraft. RAE TR 77159, 1977.

Oppenheim, A.V.; and Shafer, R.W.: Digital Signal Processing. Prentice-Hall, Englewood Cliffs, New Jersey, 1975.

Orlik-Rückemann, K.J.: Aerodynamic Coupling between Lateral and Longitudinal Degrees of Freedom. AIAA J., vol. 15, no. 12, pp. 1792-1799, Dec. 1977.

Orlik-Rückemann, K.J.: Techniques for Dynamic Stability Testing in Wind Tunnels. Dynamic Stability Parameters, AGARD-CP-235, Paper 1, 1978.

Paige, Lowell J.; Swift, J. Dean; and Slobko, Thomas A.: Elements of Linear Algebra. Second ed. Xerox College Publishing, Lexington, Massachusetts, 1974.

Panich, Yu. V.; and Trachevskii, M.L.: Identification of a Class of Closed-Loop Control Systems. Automation and Remote Control, vol. 34, no. 9, pt. 2, pp. 1457-1466, Feb. 15, 1974.

Papoulis, Athanasios: Probability, Random Variables, and Stochastic Processes. McGraw-Hill Book Co., New York, 1965.

Park, Gary D.: Determination of Tail-Off Aircraft Parameters, Systems Identification. Proc. AIAA Third Atmospheric Flight Mechanics Conference, Arlington, Texas, pp. 128-136, June 7-9, 1976.

Park, Gary D.: Parameter Identification Technology Used in Determining In-Flight Airload Parameters. J. Aircraft, vol. 14, pp. 251-256, March 1977.

Penrose, R.: A Generalized Inverse for Matrices. Proc. Cambridge Philosophical Society 51, pt. 3, pp. 406-413, 1955.

Perkins, Courtland D.; and Hage, Robert E.: Airplane Performance Stability and Control. John Wiley & Sons, New York, 1949.

Pitman, Edwin. J.G.: Some Basic Theory for Statistical Inference. Chapman and Hall, London, 1979.

Plaetschke, E.: Parameter Identification: Its Purpose, Methods, and Resulting Requirements for Flight Test Instrumentation. AGARD/Cranfield/DFVLR Short Course on Flight Test Instrumentation, 1977.

Plaetschke, E.; and Schulz, G.: Practical Input Signal Design. Parameter Identification, AGARD-LS-104, Paper 3, 29 Oct.-2 Nov. 1979.

Plaetschke, E.; Mulder, J.A.; and Breeman, J.H.: Flight Test Results of Five Input Signals for Aircraft Parameter Identification. Proc. Sixth IFAC Symposium on Identification and System Parameter Estimation, Washington, D.C., Jan. 7-11, 1982, vol. 2, pp. 1149-1154, Pergamon Press, Oxford and New York, 1983.

Plaetschke, E.; Mulder, J.A.; and Breeman, J.H.: Results of Beaver Aircraft Parameter Identification. DFVLR-FB 83-10, 1983.

Polak, Elijah: Computational Methods in Optimization: A Unified Approach. Academic Press, New York, 1971.

Pool, A.; and Bosman, D.: Basic Principles of Flight Test Instrumentation Engineering. AGARD Flight Test Instrumentation Series, AGARD-AG-160, vol. 1, 1974.

Potter, James E.: Matrix Quadratic Solutions. SIAM J. Appl. Math., vol. 14, no. 3, pp. 496-501, May 1966.

Poulter, R.L.: Measurement of the Yawing Moment of Inertia of an Aircraft (HP 115) in Flight. RAE A.R.C. R. & M. 3691, 1972.

Pouwels, H.: Instrumentation for the Determination of Aircraft Performance from Dynamic Maneuvers. Proc. 25th International Instrumentation Symposium, Anaheim, California, May 7-10, 1979, Part 2, pp. 611-621, Instrument Society of America, 1979.

Powell, M.J.D.: An Efficient Method for Finding the Minimum of a Function of Several Variables Without Calculating Derivatives. Comput. J., vol. 7, no. 2, pp. 155-162, July 1974.

Powell, J. David; and Tyler, James S., Jr.: Application of the Kalman Filter and Smoothing to VTOL Parameter and State Identification. 11th American Automatic Control Council, Joint Automatic Control Conference, Georgia Institute of Technology, Atlanta, June 22-26, 1970, Paper 18-G, pp. 449-450, American Society of Mechanical Engineering, New York, 1970.

Powers, Bruce 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.

Prasad, R.M.; Sinha, A.K.; and Mahalanabis, A.K.: Two-Stage Identification of Closed-Loop Systems. IEEE Trans. Automat. Control, vol. AC-22, no. 6, pp. 987-988, Dec. 1977.

Priestley, M.B.: Estimation of Transfer Functions in Closed Loop Stochastic Systems. Automatica, vol. 5, no. 5, pp. 623-632, Sept. 1969.

Ramachandran, S.; Schneider, H.; Mason, J.D.; and Stalford, H.L.: Identification of Aircraft Aerodynamic Characteristics at High Angles of Attack and Sideslip Using the Estimation Before Modelling (EBM) Technique. AIAA Paper 77-1169, 1977.

Ramachandran, S.; and Wells, W.R.: Estimation of Nonlinear Aerodynamic Derivatives of a Variable Geometry Fighter Aircraft from Flight Data. AIAA Paper 74-790, 1974.

Rampy, John M.; and Berry, Donald T.: Determination of Stability Derivatives from Flight Test Data by Means of High Speed Repetitive Operation Analog Matching. FTC-TDR-64-8, Edwards, California, May 1964.

Rao, C.R.: Linear Statistical Inference and Its Applications. John Wiley & Sons, New York, 1965.

Rao, S.S.: Optimization, Theory and Applications. Wiley Eastern Limited, New Delhi, 1978.

Reichert, G.: Basic Dynamics of Rotors: Control and Stability of Rotary Wing Aircraft: Aerodynamics and Dynamics of Advanced Rotary-Wing Configurations. Helicopter Aerodynamics, AGARD-LS-63, Paper 3, Apr. 1973.

Reid, G.E.A.: Validation of Kinematic Compatibility of Flight Data Using Parameter Estimation Methodology. RAE TR 81020, Mar. 1981.

Reid, G.E.A.: Identification of the Aerodynamic Stability and Control Derivatives of an Unpowered Free-Flight Aircraft Model. RAE TR 81146, Jan. 1981.

Renz, Ronald R.L.; Clarke, Robert; Mosser, Mark A.; Roskam, Jan; and Rummer, Dale: Development of a Simple, Self-Contained Flight Test Data Acquisition System. SAE Business Aircraft Meeting and Exposition, Wichita, Kansas, April 7-10, 1981, SAE Paper 810596, 1981.

Rix, O.; and Hanke, D.: In-Flight Measured Characteristics of Combined Flap-Spoiler Direct Lift Controls. Aerodynamic Characteristics of Controls, AGARD-CP-262, Paper 16, Sept. 1979.

Rix, O.; Huber, H.; and Kaletka, J.: Parameter Identification of a Hingeless Rotor Helicopter. Proc. 33rd Annual Forum of the American Helicopter Society, Washington, D.C., May 9-11, AHS 77-33-42, 1977.

Rödder, P.: Statistischer Fehler bei der Systemidentifizierung im geschlossenen Regelkreis. Regelungstechnik und Prozess Datenverarbeitung, vol. 22, pp. 282-283, 1974.

Rosenbrock, H.H.: An Automatic Method for Finding the Greatest or Least Value of a Function. Comput. J., vol. 3., pp. 175-184, 1960.

Ross, A. Jean: Determination of Aerodynamic Derivatives From Transient Responses in Manoeuvring Flight. Tehnical Memorandum Aero 1598, Royal Aircraft Establishment, 1974. (Also published in Methods for Aircraft State and Parameter Identification, AGARD-CP-172, May 1975.)

Ross, A. Jean. Determination of Aerodynamic Derivatives from Transient Response in Manoeuvring Flight. Methods for Aircraft State and Parameter Identification, AGARD-CP-172, Paper 14, 1974.

Ross, A. Jean: Application of Parameter Identification Techniques to the Analysis of Flight Data. Prog. Aerospace Sciences, vol. 18, no. 4, pp. 325-349, 1979.

Ross, A. Jean: Identification Experience in Extreme Flight Regimes. Parameter Identification, AGARD-LS-104, Paper 8, Nov. 1979.

Ross, A. Jean: Revised Values of the Lateral Aerodynamic Deviations of the Gnat Aircraft From Flight Tests. RAE TR 80099, 1980.

Ross, A. Jean; and Foster, G.W.: Fortran Programs for the Determination of Aerodynamic Derivatives From Transient Longitudinal or Lateral Responses of Aircraft. RAE ARC CP-1344, 1976.

Ross, A.J.; and Thomas, H.H.B.M.: A Survey of Experimental Data on the Aerodynamics of Controls in the Light of Future Needs. Aerodynamic Characteristics of Controls, AGARD-CP-262, Paper 2, Sept. 1979.

Ross, A. Jean; Foster, G.W.; and Turvey, T.: An Investigation of Dutch Roll and Wing Rock Oscillations of a Gnat Trainer Aircraft: Flight Tests and Linear Analysis. RAE TR 78032, Mar. 1978.

Rotier, D.J.; and Ferrin, F.J.: Ultrasonic Wind Vector Sensor. Air Data Symposium Proceedings, Naval Postgraduate School, Monterey, California, Session V: Helicoper/VSTOL Sensors, June 1976.

Royden, H.L.: Real Analysis. Second ed. MacMillan, London, 1968.

Rudin, Walter: Real and Complex Analysis. Second ed. McGraw-Hill, New York, 1974.

Rynaski, Edmund G.; Andrisani, Dominick, II; and Weingarten, Norman C.: Identification of the Stability Parameters of An Aeroelastic Airplane. AIAA Paper 78-1328, Aug. 1978.

Sage, Andrew P.; and Melsa, James L.: Estimation Theory with Application to Communications and Control. McGraw-Hill, New York, 1971.

Sage, Andrew P.; and Melsa, James L.: System Identification. Academic Press, New York, 1971.

Schmidt, E.: Die AVA-Derivativwaage. DLR-Mitt-74-32, July 1974.

Schuetz, A.J.: Low Angle-of-Attack Longitudinal Aerodynamic Parameters of Navy T-2 Trainer Aircraft Extracted from Flight Data: A Comparison of Identification Techniques; Volume I-Data Acquisition and Modified Newton-Raphson Analysis. NADC-74181-30-VOL-1 AD-A013181, Naval Air Development Center, Warminster, Pennsylvania, June 23, 1975.

Schulz, G.: Maximum Likelihood Identification Using Kalman Filtering Least-Squares Estimation. A Comparison for the Estimation of Stability Derivatives Considering Gust Disturbances. DLR-FB 75-94, Feb. 1975. (English Translation: European Space Agency TT 258.)

Schulz, G.: Entwurf optimaler Eingangssignale für die System-identifizierung unter Beruecksichtigung von Mess - und Systemrauschen. Regelungstechnik, Heft 10, pp. 324-330, 1977.

Schulze, H.: Anwendung von Schaetzverfahren für die Kenngrössen von Regelstrecken aufgrund von Messungen im geschlossenen Regelkreis. Regelungstechnik und Prozess-Datenverarbeitung, vol. 19, pp. 113-119, 1971.

Schwalm, D.: Bestimmung des Frequenzganges eines linearen rueckgekoppelten Systems mit zwei unabhaengigen stochastischen Eingangssignalen. Regelungstechnik und Prozess-Datenverarbeitung, vol. 18, pp. 453-455, 1970.

Schweppe, Fred C.: Uncertain Dynamic Systems. Prentice-Hall, Englewood Cliffs, New Jersey, 1973.

Shafer, Mary F.: Stability and Control Drivatives of the T-37B Airplane. NASA TM X-56036, Sept. 1975.

Shafer, Mary F.: Flight-Determined Correction Terms for Angle of Attack and Sideslip. AIAA Paper 82-1374, Aug. 1982 (a).

Shafer, Mary F.: Low-Order Equivalent Models of Highly Augmented Aircraft Determined from Flight Data. J. Guid. Control Dynam., vol. 5, no. 5, pp. 504-511, Sept.-Oct. 1982 (b).

Shinbrot, Marvin: A Least-Squares Curve Fitting Method With Application to the Calculation of Stability Coefficients From Transient Response Data. NACA TN-2341, Apr. 1951.

Siemers, P.M., III; Wolf, H.; and Flanagan, P.F.: Shuttle Entry Air Data System Concepts Applied to Space Shuttle Orbiter Flight Pressure Data to Determine Air Data — STS 1-4. AIAA Paper 83-0118, Jan. 1983.

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

Sisk, T.R.: A Technique for the Assessment of Fighter Aircraft Precision Controllability. AIAA Paper 78-1364, Aug. 1978.

Skudridakis, Jean: Systemmodellierung Und Identifizie rung von 'äquivalenten' Systemen mit und ohne Ersatztotzeit feur reglergesteutzte Flugeuge. DFVLR-FB 83-036, Nov. 1983. (Also published as translation ESA-TT-852.)

Smith, H.J.: Flight-Determined Stability and Control Derivatives for an Executive Jet Transport. NASA TM X-56034, July 1975.

Smith, R.H.: A Theory for Longitudinal Short-Period Pilot Induced Oscillations. AFFDL-TR-77-57, AD-A056982, June 1977.

Smith, B.T.; Boyle, J.M.; Dongarra, J.J.; Garbow, B.S.; Ikebe, Y.; Klema, V.C.; and Moler, C.B.: Matrix Eigensystem Routines-EISPack Guide. Second ed. Springer-Verlag, Berlin, 1976.

Söderström, T.; Gustavsson, I.; and Ljung, L.: Identifiability Conditions for Linear Systems Operating in Closed Loop. Int. J. Control, vol. 21, no. 2, pp. 243-255, Feb. 1975.

Söderström, T.; Ljung, L.; and Gustavsson, I.: Identifiability Conditions for Linear Multivariable Systems Operating Under Feedback. IEEE Trans. Automat. Control, vol. AC-21, pp. 837-840, Dec. 1976,

Sorensen, John A.: Analysis of Instrumentation Error Effects on the Identification Accuracy of Aircraft Parameters. NASA CR-112121, 1972.

Sorenson, Harold W.: Parameter Estimation; Principles and Problems. Marcel Dekker, Inc., New York, 1980.

Stalford, H.L.: Application of the Estimation-Before-Modelling (EBM) System Identification Method to the High Angle of Attack/Sideslip Flight of the T-2C Jet Trainer Aircraft. Three vols., AD-A079924, Dynamics Research Corporation, System Division, Wilmington, Massachusetts, Nov. 1979.

Stalford, H.L.: High-Alpha Aerodynamic Model Identification of T-2C Aircraft Using the EBM Method. J. Aircraft, vol. 18, no. 10, pp. 801-809, Oct. 1981.

Steers, Sandra Thornberry; and Iliff, Kenneth W.: Effects of Time-Shifted Data on Flight-Determined Stability and Control Derivatives. NASA TN D-7830, 1975.

Stepner, D.E.; and Mehra, R.K.: Maximum Likelihood Identification and Optimal Input Design for Identifying Aircraft Stability and Control Derivatives. NASA CR-2200, 1973.

Stewart, G.W., III: A Modification of Davidon's Minimization Method to Accept Difference Approximation of Derivatives. J. ACM, vol. 14, no. 1, pp. 72-83, Jan. 1967.

Strang, Gilbert: Linear Algebra and Its Applications. Second ed. Academic Press, New York, 1980.

Subke, H.: Test Installations to Investigate the Dynamic Behaviour of Aircraft with Scaled Models in Wind Tunnels. Symposium Dynamic Analysis of Vehicle Ride and Manoeuvering Characteristics, London, Preprint pp. 137-148, Nov. 28-30, 1978.

Suit, William T.: Aerodynamic Parameters of the Navion Airplane Extracted from Flight Data. NASA TN D-6643, 1972.

Tanner, Russel R.; and Montgomery, Terry D.: Stability and Control Derivative Estimates Obtained From Flight Data for the Beech 99 Aircraft. NASA TM-72863, 1979.

Taylor, L.W., Jr.: A New Criterion for Modeling Systems. Parameter Estimation Techniques and Application in Aircraft Flight Testing, NASA TN D-7647, pp. 291-313, 1974.

Taylor, L.W., Jr.: Application of a New Criterion for Modelling Systems. Methods for Aircraft State and Parameter Identification, AGARD-CP-172, Paper 4, May 1974.

Taylor, Lawrence W., Jr.; and Iliff, Kenneth W.: A Modified Newton-Raphson Method for Determining Stability Derivatives from Flight Data. Second International Conference on Computing Methods in Optimization Problems, San Remo, Italy, Sept. 9-13, 1968, pp. 353-364, Academic Press, New York, 1969.

Taylor, Lawrence W., Jr.; and Iliff, Kenneth W.: Systems Identification Using a Modified Newton-Raphson Method: A FORTRAN Program. NASA TN D-6734, 1972.

Thomas, H.H.B.M.; and Edwards, Geraldine: Mathematical Models of Aircraft Dynamics for Extreme Flight Conditions (Theory and Experiment). Dynamic Stability Parameters, AGARD-CP-235, Paper 27, Nov. 1978.

Thomas, H.H.B.M.; and Ross, A. Jean: The Role of Theoretical Studies of Flight Dynamics in Relation to Flight Testing. Stability and Control, AGARD-CP-119, Paper 22, Nov. 1972.

Tomaine, R.L.: The Effect of Pilot Control Input Shape on the Identification of Six Degree-of-Freedom Stability and Control Derivatives of a Transport Helicopter. Master's Thesis, George Washington University, Washington, D.C., Dec. 1976.

Tomaine, R.L.: Flight Data Identification of Six Degree-of-Freedom Stability and Control Derivatives of a Large "Crane" Type Helicopter. NASA TM X-73958, 1976.

Tomaine, R.L.; Bryant, W.H.; and Hodge, W.F.: VALT Parameter Identification Flight Test. Fourth European Rotorcraft and Powered Lift Aircraft Forum, Stresa, Italy, Paper 73, Sept. 13-15, 1978.

Trankle, T.L.; Vincent, J.H.; and Franklin, S.N.: System Identification of Nonlinear Aerodynamic Models. Advances in the Techniques and Technology of the Application of Nonlinear Filters and Kalman Filters, AGARD-AG-256, Paper 7, Mar. 1982.

Trenkle, F.; and Reinhardt, M.: In Flight Temperature Measurements. AGARD-AG-160, vol. 2, 1973.

Twisdale, T.R.; Jones, G.L.; and Ashurst, T.A.: A Mission Oriented Flight Test Technique for Identifying Aircraft and Flight Control System Transfer Functions. Flight Test Techniques, AGARD-CP-223, Paper 13, Apr. 1977.

Van den Boom, A.J.W.: On the Relation Between Weighted Least-Squares Estimators and Instrumental Variable Estimators. Fourth IFAC Symposium on Identification and System Parameter Estimation, Tbilisi, U.S.S.R., Sept. 21-27, 1976, Naum S. Raibman, ed., North Holland, Amsterdam and New York, 1978.

Van der Linen, J.C.; and Mensink, H.A.: Linear and Angular Position Measurement of Aircraft Components. AGARD Flight Test Instrumentation Series, AGARD-AG-160, vol. 8, 1977.

Van der Wilt, M.: Flight-Path Reconstruction of Symmetric Unsteady Flights. National Aerospace Laboratory NLR-TR-76133-U, The Netherlands, May 1976.

Van Trees, H. L.: Detection, Estimation and Modulation Theory. Volume 1. John Wiley & Sons, New York, 1969.

Van Woerkom, K.: Design and Evaluation of an Instrumentation System for Measurements in Nonsteady Symmetrical Flight Conditions with the Hawker Hunter Mk VII. Delft University of Technology, Department of Aerospace Engineering, Delft, The Netherlands, Report LR-308, Jan. 1981.

Vaughan, David R.: A Nonrecursive Algebraic Solution for the Discrete Riccati Equation. IEEE Trans. Automat. Control, vol. AC-15, pp. 597-599, Oct. 1970.

Velkoff, H.R.: Technical Evaluation Report on the Flight Mechanics Panel Symposium on Rotorcraft Design. AGARD-AG-114, Jan. 1978.

Verbrugge, R.; Rotative Balance of I.M.F. Lille and Associated Experimental Techniques. NASA TM-75886, 1981. (Translation of Balance reotative 'de-71 I.M.F. Lille et Techniques Experimentales Associees'. AAAF Paper NT 80-13, Lille, 1979.)

Verbrugge, R.A.; Charon, W.; and Marchand, M.: Wind Tunnel and Free-Flight Model Identification Experience. Parameter Identification, AGARD-LS-104, Paper 9, Nov. 1979.

Vincent, J.H.; Gupta, N.K.; and Hall, W.E., Jr.: Recent Results in Parameter Identification for High Angle-of-Attack Stall Regimes. AIAA Paper 79-1640, Aug. 1979.

Vleghert, J.P.K.: Measuring Climb Performance of a Propellor Engined Transport Aeroplane Using the Acceleration Technique, AGARD Rep. 127, May 1957.

Von der Decken, Jan; Schmidt, Eberhard; and Schulze, Bernd: On the Test Procedures of the Derivative Balances Used in West Germany. Dynamic Stability Parameters, AGARD-CP-235, Paper 6, 1978.

Vorchik, B.G.: Plant Identification in a Stochastic Closed-Loop System. Automation and Remote Control, vol. 36, pp. 550-565, Sept. 1975.

Vorchik, B.G.; Fetisov, V.N.; and Steinberg, Sh. E.: Identification of a Closed-Loop Stochastic System. Automation and Remote Control, vol. 34, pp. 1069-1081, Dec. 1, 1973.

Waterfall, A.P.: A Technique for the Automatic Digital Analysis of Flight Dynamic Response Data. A.R.C. R. & M. 3699 (RAE Tech. Rep. 70228), 1970.

Weiner, Norbert: Extrapolation, Interpolation, and Smoothing of Stationary Time Series: With Engineering Applications. John Wiley & Sons, New York, 1949.

Wells, W.R.; and Ramachandran, S.: Flight Test Design for Efficient Extraction of Aircraft Parameters. Proc. AIAA Third Atmospheric Flight Mechanics Conference, June 1976, pp. 101-107, AIAA, New York, 1976,

Wellstead, P.E.; and Edmunds, J.M.: Least-Squares Identification of Closed-Loop Systems. Int. J. Control, vol. 21, no. 4, pp. 689-699, Apr. 1975.

Whitmoyer, R.A.: Aerodynamic Interactions on the Fighter CCV Test Aircraft. Dynamic Stability Parameters, AGARD-CP-235, Paper 16, Nov. 1978.

Wiberg, Donald M.: Schaum's Outline of Theory and Probems of State Space and Linear Systems. McGraw-Hill Book Co., New York, 1971.

Wiberg, D.M.; Bellville, J.W.; Brovko, O.; Maine, R.E.; and Tai, T.C.: Modeling and Parameter Identification of the Human Respiratory System. IEEE Trans. Automat. Control, vol AC-24, pp. 716-720, Oct. 1979.

Wilhelm, K.; and Verbrugge, R.A.: Correlation Aspects in the Identification of Dynamic Effects Using Complementary Techniques, Flight in Turbulence: Gust Alleviation. Ground Flight Test Techniques and Correlation, AGARD-CP-339, Paper 17, Feb. 1982.

Wilkinson, James H.: The Algebraic Eigenvalue Problem. Clarendon Press, Oxford, 1965.

Williams, D.A.: The Analysis of Random Data. AGARD-AG-160, vol. 14, 1981.

Williams, James L.; and Suit, William T.: Extraction From Flight Data of Lateral Aerodynamic Coefficients for F-8 Aircraft With SuperCritical Wing. NASA TN D-7749, 1974.

Williams J.E.; and Vukelich, S.R.: The USAF Stability and Control Digital Datcom, Volume 1 User's Manual. AFFDL-TR-79-3032, AD-A086557, Apr. 1979.

Wilson, Donald B.; and Winters, Charles P.: F-15A Approach-to-Stall/Stall/Post-Stall Evaluation. AFFTC TR-75-32, Air Force Flight Test Center, Edwards Air Force Base, California, Jan. 1976.

Wingrove, R.C.: Applications of a Technique for Estimating Aircraft States from Recorded Flight Test Data. AIAA Paper 72-965, Sept. 1972.

Wingrove, R.C.: Quasi-Linearization Technique for Estimating Aircraft States from Flight Data. J. Aircraft, vol. 10, no. 5, pp. 303-307, May 1973.

Wingrove, R.C.: Parameter Estimation of Powered-Lift STOL Aircraft Characteristics Including Turbulence and Ground Effects. Methods for Aircraft State and Parameter Identification, AGARD-CP-172, Paper 28, May 1975.

Wolowicz, Chester H.: Considerations in the Determination of Stability and Control Derivatives and Dynamic Characteristics from Flight Data. AGARD-AR-549, Part 1, 1966.

Wolowicz, Chester H.; and Holleman, Euclid C.: Stability-Derivative Determination from Flight Data. AGARD Rep. 224, Oct. 1958.

Wolowicz, Chester H.; and Yancy, Roxanah B.: Experimental Determination of Airplane Mass and Inertial Characteristics. NASA TR R-433, 1974.

Wong, K.V.; and Polak, E.: Identification of Linear Discrete Time Systems Using the Instrumental Variable Method. IEEE Trans. Automat. Control, vol. AC-12, no. 6, pp. 707-718, Dec. 1967.

Wuest, W.: Pressure and Flow Measurement — Flight Testing. AGARD-AG-160, vol. 11, 1980.

Yazawa, Kenji: Identification of Aircraft Stability and Control Derivatives in the Presence of Turbulence. AIAA Paper 77-1134, Aug. 1977.

Yosida, Kosaku: Functional Analysis. Second ed. Springer-Verlag, Berlin, 1968.

Zacks, Shelemyahu: The Theory of Statistical Inference. John Wiley & Sons, New York, 1971.

Zadeh, Lotfi A.; and Desoer, Charles A.: Linear System Theory. McGraw-Hill Book Co., New York, 1963.

Zverev, A.I.: Handbook of Filter Synthesis. Wiley, New York, 1967.

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
NASA RP-1168		

4. Title and Subtitle	5. Report Date
Application of Parameter Estimation to Aircraft Stability and Control — The Output-Error Approach	June 1986
	6. Performing Organization Code
7. Author(s)	8. Performing Organization Report No.
Richard E. Maine and Kenneth W. Iliff	H-1299
	10. Work Unit No.
9. Performing Organization Name and Address	RTOP 505-61
NASA Ames Research Center Dryden Flight Research Facility P.O. Box 273 Edwards, CA 93523-5000	11. Contract or Grant No.
	13. Type of Report and Period Covered
12. Sponsoring Agency Name and Address	Reference Publication
National Aeronautics and Space Administration Washington, D.C. 20546	14. Sponsoring Agency Code

15. Supplementary Notes

This document was prepared for the AGARD Flight Mechanics Panel, Flight Test Techniques Group, as part of AGARDograph 300, Flight Test Techniques Series

16. Abstract

This document examines the practical application of parameter estimation methodology to the problem of estimating aircraft stability and control derivatives from flight test data. The primary purpose of the document is to present a comprehensive and unified picture of the entire parameter estimation process and its integration into a flight test program. The document concentrates on the output-error method to provide a focus for detailed examination and to allow us to give specific examples of situations that have arisen in our experience.

The document first derives the aircraft equations of motion in a form suitable for application to estimation of stability and control derivatives. It then discusses the issues that arise in adapting the equations to the limitations of analysis programs, using a specific program for an example.

The document then addresses the roles and issues relating to mass distribution data, preflight predictions, maneuver design, flight scheduling, instrumentation sensors, data acquisition systems, and data processing. Finally, the document discusses evaluation and use of the analysis results.

17. Key Words (Suggested by Author(s))	18. Distribution Statement
Aircraft flight testing Aircraft stability and control Maximum likelihood Parameter estimation	Unclassified — Unlimited STAR category 05

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price*
Unclassified	Unclassified	180	A09

For sale by the National Technical Information Service, Springfield, Virginia 22161.