Parameter Uncertainty for Aircraft Aerodynamic Modeling using Recursive Least Squares

Jared A. Grauer* and Eugene A. Morelli†

NASA Langley Research Center, Hampton, Virginia 23681

A real-time method was demonstrated for determining accurate uncertainty levels of stability and control derivatives estimated using recursive least squares and time-domain data. The method uses a recursive formulation of the residual autocorrelation to account for colored residuals, which are routinely encountered in aircraft parameter estimation and change the predicted uncertainties. Simulation data and flight test data for a subscale jet transport aircraft were used to demonstrate the approach. Results showed that the corrected uncertainties matched the observed scatter in the parameter estimates, and did so more accurately than conventional uncertainty estimates that assume white residuals. Only small differences were observed between batch estimates and recursive estimates at the end of the maneuver. It was also demonstrated that the autocorrelation could be reduced to a small number of lags to minimize computation and memory storage requirements without significantly degrading the accuracy of predicted uncertainty levels.

Nomenclature

a_x, a_y, a_z	body-axis accelerometer output, g
b	wing span, ft
C_l, C_m, C_n	body-axis moment coefficients
C_X, C_Y, C_Z	body-axis force coefficients
$\bar{c}$	mean aerodynamic chord, ft
cov(.)	covariance
$\mathbf{D}$	dispersion matrix
$E[.]$	expectation operator
h	altitude, ft
$\mathbb{I}$	identity matrix
$I_{..}$	moment of inertia, slug·ft^2
J	cost function
$\mathbf{K}$	gain matrix
m	mass, slug
N	number of samples
p, q, r	body-axis angular rate, rad/s
$\bar{q}$	dynamic pressure, lbf/ft^2
S	wing reference area, ft^2
$S_{..}$	power spectral density
t	time, s
V	true airspeed, ft/s
$\mathbf{X}$	regressor matrix
$\mathbf{y}$	modeled output
$\mathbf{z}$	measured output
α	angle of attack, rad
β	angle of sideslip, rad
Δ	perturbation
δ	control surface deflection, rad
$\boldsymbol{\epsilon}$	measurement equation error
$\boldsymbol{\theta}$	model parameter vector
σ	standard deviation

Subscripts

0	bias or trim value
a, e, r	aileron, elevator, rudder
cm	center of mass
T	thrust

Superscripts

$^{-1}$	inverse
^{T}	transpose
$\hat{}$	estimated value
$\dot{}$	time derivative

*Research Engineer, Dynamic Systems and Control Branch, MS 308. Senior Member AIAA.

†Research Engineer, Dynamic Systems and Control Branch, MS 308. Associate Fellow AIAA.

I. Introduction

LEAST-squares parameter estimation is a commonly-used method for building models from measured data because of a number of appealing qualities including simplicity, efficiency, and extensibility. The recursive form of this approach is an attractive option for estimating aircraft dynamic models in real time as flight test data are being collected. Such a real-time analysis can enable rapid flight envelope expansion, control law verification, input type or amplitude selection, or aerodynamic database identification. Real-time estimates could instead update adaptive control laws during aircraft faults or loss-of-control type scenarios. These algorithms are simple enough to run on low-cost systems aboard unmanned air vehicles.

However, recursive least squares is generally not favored and has instead been superseded by other methods, such as in Ref. [1], for real-time estimation. One reason is that recursive least squares produces inaccurate parameter uncertainty estimates, which is a primary indicator of estimation quality. One source of error is colored modeling residuals, which violate the underlying theory and result in parameter covariances that are too small. This error causes a misleadingly optimistic interpretation of the estimation results. For these reasons, parameter uncertainty calculations are delayed until the end of the maneuver[2] or are omitted entirely.[3]

A residual sequence is said to be colored when it is correlated in time with itself. This occurs routinely with aircraft flight test data because of model structure error, which can arise from numerous sources. For example, turbulence, unsteady aerodynamics, and aeroelastic effects all create responses that are not considered in typical modeling applications. Other neglected dynamics can include nonlinearities, cross-axis coupling, and interactions with the propulsion system. Steps taken in the data reduction and analysis can also potentially color the residuals.

Reference [4] identified colored residuals as the source of the discrepancy between predicted parameter covariance and observed scatter in flight test results. The phenomenon was well-known at the time, and the standard solution approach was to multiply the parameter stanadrd errors by an arbitrary "fudge factor" of 5 to 10. In Ref. [5], a batch correction for the uncertainty levels based on the coloring of the model residuals was presented for maximum likelihood estimators, with least squares as a special case. This correction was made recursive for real-time estimation in Ref. [6], and was applied to frequency response estimation from input and output measurements.

This paper investigates the real-time method presented in Ref. [6] but for a different application, namely, estimating accurate parameter uncertainties of aircraft stability and control derivatives. Estimation of stability and control derivatives from flight test data is a common task, and accurate results are essential because they are used in important applications.[7] The aerodynamic models containing these parameters are described in Section II, along with the batch and recursive least squares algorithms. That section then explains the recursive correction for colored residuals that enables accurate real-time parameter uncertainty prediction. Section III describes the T-2 subscale jet transport aircraft used to study the methods, as well as a linearized flight dynamics simulation model. Results using simulation data and flight test data for the T-2 aircraft are then presented in Section IV, followed by conclusions in Section V.

Software for the input design, real-time smooth differentiation, and batch least-squares estimation used in this work are available in a MATLAB® toolbox called System IDentification Programs for AirCraft, or SIDPAC.[2]

II. Theoretical Development

II.A. Aerodynamic Modeling

Nondimensional aerodynamic force and moment coefficients cannot be directly measured in flight, and are instead computed from flight test data and known quantities. Under typical simplifying assumptions,

the nondimensional aerodynamic coefficients can be computed as[2]

$$C_X = \frac{1}{\bar{q}S}(ma_x - X_T) \tag{1a}$$

$$C_Y = \frac{1}{\bar{q}S}(ma_y) \tag{1b}$$

$$C_Z = \frac{1}{\bar{q}S}(ma_z) \tag{1c}$$

$$C_l = \frac{1}{\bar{q}Sb}\left[I_{xx}\dot{p} - I_{xz}(\dot{r} + pq) + (I_{zz} - I_{yy})qr\right] \tag{1d}$$

$$C_m = \frac{1}{\bar{q}S\bar{c}}\left[I_{yy}\dot{q} + (I_{xx} - I_{zz})pr + I_{xz}(p^2 - r^2)\right] \tag{1e}$$

$$C_n = \frac{1}{\bar{q}Sb}\left[I_{zz}\dot{r} - I_{xz}(\dot{p} - qr) + (I_{yy} - I_{xx})pq\right] \tag{1f}$$

which are expressed in the body axes and retain the full nonlinearity of the aircraft motion. These equations require the aircraft geometry and mass properties, thrust, and measurements of the motion. Time histories of the rotational accelerations are often obtained by smoothly differentiating angular rate data.[2,8]

The aerodynamic coefficients can be modeled using linear expansions in the aircraft states and controls, assuming small excursions from a reference flight condition and short durations of time. These models can be simplified to include only the most dominant on-axis effects, for example[9]

$$C_X = C_{X_0} + C_{X_\alpha}\Delta\alpha + C_{X_{\delta_e}}\Delta\delta_e \tag{2a}$$

$$C_Y = C_{Y_0} + C_{Y_\beta}\Delta\beta + C_{Y_r}\frac{b}{2V}r + C_{Y_{\delta_r}}\Delta\delta_r \tag{2b}$$

$$C_Z = C_{Z_0} + C_{Z_\alpha}\Delta\alpha + C_{Z_{\delta_e}}\Delta\delta_e \tag{2c}$$

$$C_l = C_{l_0} + C_{l_\beta}\Delta\beta + C_{l_p}\frac{b}{2V}p + C_{l_r}\frac{b}{2V}r + C_{l_{\delta_a}}\Delta\delta_a + C_{l_{\delta_r}}\Delta\delta_r \tag{2d}$$

$$C_m = C_{m_0} + C_{m_\alpha}\Delta\alpha + C_{m_q}\frac{\bar{c}}{2V}q + C_{m_{\delta_e}}\Delta\delta_e \tag{2e}$$

$$C_n = C_{n_0} + C_{n_\beta}\Delta\beta + C_{n_r}\frac{b}{2V}r + C_{n_{\delta_r}}\Delta\delta_r \tag{2f}$$

The bias terms, for example C_{Z_0}, subsume all the steady portions of the signals, including aerodynamic biases and non-zero trim contributions. These terms, sometimes referred to as "nuisance parameters," are estimated but are not of primary interest.[9] The quantities multiplying the normalized state and control variables are the nondimensional stability and control derivatives.

II.B. Least Squares Parameter Estimation

In the equation-error formulation, which is essentially the least-squares problem, the dependent or response variables are the aerodynamic coefficients computed using Eqs. (1), the regressor or explanatory variables are the normalized state and control measurements in Eqs. (2), and the unknown parameters to be estimated are the corresponding stability and control derivatives. The estimation is decoupled in that each aerodynamic coefficient can be modeled separately, instead of all together. Each analysis has the form

$$\begin{aligned} \mathbf{z} &= \mathbf{y} + \boldsymbol{\epsilon} \\ &= \mathbf{X}\boldsymbol{\theta} + \boldsymbol{\epsilon} \end{aligned} \tag{3}$$

where $\mathbf{z}$ and $\mathbf{y}$ are the $N \times 1$ time histories of the measured and modeled dependent variable, respectively, $\mathbf{X}$ is a $N \times n_\theta$ matrix of independent variable or regressor time histories, $\boldsymbol{\theta}$ is an $n_\theta \times 1$ array of unknown model parameters, and $\boldsymbol{\epsilon}$ is an $N \times 1$ array of the measurement equation errors. For example using the vertical force coefficient,

$$\mathbf{z} = \frac{mg}{\bar{q}S}\begin{bmatrix} a_{z_1} \\ a_{z_2} \\ \vdots \\ a_{z_N} \end{bmatrix}, \qquad \mathbf{X} = \begin{bmatrix} \mathbf{x}_1^T \\ \mathbf{x}_2^T \\ \vdots \\ \mathbf{x}_N^T \end{bmatrix} = \begin{bmatrix} 1 & \alpha_1 & \delta_{e_1} \\ 1 & \alpha_2 & \delta_{e_2} \\ \vdots & \vdots & \vdots \\ 1 & \alpha_N & \delta_{e_N} \end{bmatrix}, \qquad \boldsymbol{\theta} = \begin{bmatrix} C_{Z_0} \\ C_{Z_\alpha} \\ C_{Z_{\delta_e}} \end{bmatrix} \tag{4}$$

When the regressors are all measured, contain no error, and form an adequate model structure, least squares can be applied to match the aerodynamic coefficient time histories and accurately estimate the stability and control derivatives.

The least-squares cost function

$$J(\boldsymbol{\theta}) = \frac{1}{2}\left(\mathbf{z}-\mathbf{y}\right)^T\left(\mathbf{z}-\mathbf{y}\right) \tag{5}$$

represents the total generalized distance between the measured and modeled dependent variable. This cost function is minimized by equating the cost gradient to zero and solving for the optimal model parameters as

$$\begin{aligned}\hat{\boldsymbol{\theta}} &= \left(\mathbf{X}^T\mathbf{X}\right)^{-1}\mathbf{X}^T\mathbf{z}\\ &= \mathbf{D}\mathbf{X}^T\mathbf{z}\end{aligned} \tag{6}$$

where $\mathbf{D}$ is the dispersion matrix and is related to the information content in the data. The difference between the dependent variable and the model output

$$\mathbf{v} = \mathbf{z} - \mathbf{y} \tag{7}$$

are called the model residuals.

The uncertainty in the model parameter estimates is quantified by the parameter covariance matrix, which simplifies in this case as

$$\begin{aligned}\mathrm{cov}(\hat{\boldsymbol{\theta}}) &= E\left[(\hat{\boldsymbol{\theta}}-\boldsymbol{\theta})(\hat{\boldsymbol{\theta}}-\boldsymbol{\theta})^T\right]\\ &= E\left[\left(\mathbf{X}^T\mathbf{X}\right)^{-1}\mathbf{X}^T(\mathbf{z}-\mathbf{y})(\mathbf{z}-\mathbf{y})^T\mathbf{X}\left(\mathbf{X}^T\mathbf{X}\right)^{-1}\right]\\ &= \mathbf{D}\mathbf{X}^T E[\mathbf{v}\mathbf{v}^T]\mathbf{X}\mathbf{D}\end{aligned} \tag{8}$$

because the regressors are assumed to be deterministic. The square root of the diagonal entries of the parameter covariance are the Cramér-Rao bounds, which provide a lower bound on the parameter standard error. The middle term in Eq. (8) is the autocorrelation matrix of the residual, defined as

$$E[\mathbf{v}\mathbf{v}^T] = \begin{bmatrix} \mathcal{R}(0) & \mathcal{R}(1) & \dots & \mathcal{R}(N-1)\\ \mathcal{R}(1) & \mathcal{R}(0) & \dots & \mathcal{R}(N-2)\\ \vdots & \vdots & \ddots & \vdots\\ \mathcal{R}(N-1) & \mathcal{R}(N-2) & \dots & \mathcal{R}(0)\end{bmatrix} \tag{9}$$

which is symmetric about the diagonal. Entries in this matrix are approximated by the discrete sample autocorrelation function of the residual

$$\hat{\mathcal{R}}(i) = \frac{1}{N}\sum_{j=1}^{N-i} v_{i+j}v_j, \qquad i = 0, 1, \dots, N-1 \tag{10}$$

taken at different lag indices i. The autocorrelation function is even, so that $\mathcal{R}(i) = \mathcal{R}(-i)$. The definition used in Eq. (10) is biased because it does not consider the statistical degrees of freedom. However, this form can have less mean square error than other definitions, and the bias is negligible for typical flight test sampling rates and record lengths. The parameter covariance in Eq. (8) can alternatively be expressed using summations as

$$\mathrm{cov}(\hat{\boldsymbol{\theta}}) = \mathbf{D}\left[\sum_{i=1}^{N}\sum_{j=1}^{N}\mathbf{x}_i\mathcal{R}(i-j)\mathbf{x}_j^T\right]\mathbf{D} \tag{11}$$

If the model residuals can be assumed to have constant variance σ^2 and are uncorrelated in time, which is the conventional assumption of white residuals, then

$$\mathcal{R}(i) = \begin{cases}\sigma^2, & i = 0\\ 0, & \text{otherwise}\end{cases} \tag{12}$$

and the parameter covariance reduces to

$$\begin{aligned}\mathrm{cov}(\hat{\boldsymbol{\theta}}) &= \mathbf{D}\left[\sum_{i=1}^{N}\sum_{j=1}^{N}\mathbf{x}_i(\sigma^2\mathbb{I})\mathbf{x}_j^T\right]\mathbf{D} \\ &= \mathbf{D}\left[\sigma^2\left(\mathbf{X}^T\mathbf{X}\right)\right]\mathbf{D} \\ &= \sigma^2\mathbf{D} \end{aligned} \tag{13}$$

It is usually necessary to estimate the residual variance from the data because accurate values from prior or repeated experiments are difficult to obtain. One technique is to use the fit error variance

$$\hat{\sigma}^2 = \frac{1}{N}\mathbf{v}^T\mathbf{v} \tag{14}$$

which is based on the model residuals and is therefore dependent upon the accuracy of the model structure and parameter estimates. When the residuals are significantly colored, Eq. (13) will underpredict the parameter uncertainties, and the more general expression in Eq. (11) is needed for accurate results.[5]

II.C. Recursive Least Squares

The previous formulation of least squares is for batch estimation, after all the data from an experiment has been collected. In the recursive version, estimates are updated as each new set of measurements are obtained, viz.[2]

$$\mathbf{K}_k = \mathbf{D}_{k-1}\mathbf{x}_k\left[1+\mathbf{x}_k^T\mathbf{D}_{k-1}\mathbf{x}_k^T\right]^{-1} \tag{15a}$$

$$\mathbf{D}_k = \left[\mathbb{I}-\mathbf{K}_k\mathbf{x}_k^T\right]\mathbf{D}_{k-1} \tag{15b}$$

$$\hat{\boldsymbol{\theta}}_k = \hat{\boldsymbol{\theta}}_{k-1}+\mathbf{K}_k\left[z_k-\mathbf{x}_k^T\hat{\boldsymbol{\theta}}_{k-1}\right] \tag{15c}$$

where the subscript k is used to denote the sample index. With each sample, a new set of regressors $\mathbf{x}_k$ and a new measurement z_k are obtained. This information is used to compute a gain matrix $\mathbf{K}_k$ and the dispersion matrix $\mathbf{D}_k$. The parameter estimate is then updated, based upon the previous estimate and the weighted residual.

Assuming white residuals, the parameter covariance in Eq. (13) is updated using the recursive dispersion matrix and a recursive estimate of the fit error variance

$$\begin{aligned}\hat{\sigma}_k^2 &= \frac{1}{k}\left(v_1^2+v_2^2+\ldots+v_{k-1}^2+v_k^2\right) \\ &= \left(\frac{k-1}{k-1}\right)\left(\frac{1}{k}\right)\left(v_1^2+v_2^2+\ldots+v_{k-1}^2\right)+\frac{1}{k}v_k^2 \\ &= \left(\frac{k-1}{k}\right)\hat{\sigma}_{k-1}^2+\frac{1}{k}v_k^2 \end{aligned} \tag{16}$$

There is a subtle difference in how the batch and recursive methods compute the parameter uncertainties. The batch algorithm uses all the data to estimate one estimate of the model parameters, from which the residuals and parameter uncertainties are computed. The recursive algorithm uses the most recent sample to update the parameter estimates, but does not recompute the residual history using these new estimates. Although the parameter estimates will match at the end of the maneuver using either algorithm, the residuals, and thus parameter uncertainties, are not expected to be the same. Additionally, when there is little information content in the data, such as early in the data record, parameter estimates are poor, which then make the residuals and parameter uncertainties also incorrect. Accuracy is sacrificed here for the sake of real-time estimation. Colored residuals will also increase the error in the recursive parameter uncertainties computed in this way.

II.D. Recursive Parameter Covariance With Colored Residuals

The batch algorithm for parameter covariances considering colored residuals given in Eq. (11) can be made recursive, as described in Ref. [6]. The first step is to write the residual autocorrelation in a recursive

form, as

$$\hat{\mathcal{R}}_k(i) = \left(\frac{k-1}{k}\right)\hat{\mathcal{R}}_{k-1}(i) + \left(\frac{1}{k}\right)v_{k-i}v_k \tag{17}$$

Equation (17) produces the same values as the analogous equation given in Ref. [6], but is expressed in a different, and perhaps simpler, form which was derived in a similar manner as Eq. (16). The residual autocorrelation function depends on past autocorrelation values and past residuals, which are both functions of past parameter estimates. This expression is not expected to give the same numerical values as a batch estimate, although they should be similar once parameter estimates begin to converge.

The next step is to write the parameter covariance considering colored residuals in a recursive form. By expanding the summand and regrouping terms, Eq. (11) can be rewritten as

$$\mathrm{cov}(\hat{\boldsymbol{\theta}}) = \mathbf{D}\left[\sum_{i=0}^{N-1}\mathcal{R}(i)\boldsymbol{\Lambda}(i)\right]\mathbf{D} \tag{18}$$

where

$$\boldsymbol{\Lambda}(i) = \begin{cases} \displaystyle\sum_{j=1}^{N}\mathbf{x}_j\mathbf{x}_j^T, & i = 0 \\ \\ \displaystyle\sum_{j=1}^{N-i}\mathbf{x}_{i+j}\mathbf{x}_j^T + \mathbf{x}_j\mathbf{x}_{i+j}^T, & i > 0 \end{cases} \tag{19}$$

where the index i here represents the relative lags between samples. The first definition for $\boldsymbol{\Lambda}(0)$ represents the diagonal products involving the zero-lag autocorrelation, and the remaining definition is for the products involving non-zero lag indices. The recursive parameter covariance update is then

$$\mathrm{cov}(\hat{\boldsymbol{\theta}}_k) = \mathbf{D}_k\left[\sum_{i=0}^{k-1}\mathcal{R}_k(i)\boldsymbol{\Lambda}_k(i)\right]\mathbf{D}_k \tag{20}$$

where

$$\boldsymbol{\Lambda}_k(i) = \begin{cases} \boldsymbol{\Lambda}_{k-1}(i) + \mathbf{x}_k\mathbf{x}_k^T, & i = 0 \\ \\ \boldsymbol{\Lambda}_{k-1}(i) + \mathbf{x}_{k-i}\mathbf{x}_k^T + \mathbf{x}_k\mathbf{x}_{k-i}^T, & i > 0 \end{cases} \tag{21}$$

and can be used for accurate real-time estimation of the parameter uncertainty levels.

Note that as more data are collected, the summation in Eq. (20) operates over a growing number of indices. More computations and more memory storage are required as the autocorrelation matrix increases in size. The autocorrelation function can, however, be truncated after a finite number of lag terms because only proximate residuals are correlated in typical flight test data. Computation and memory storage requirements would then be reduced and bounded. This approximation can be visualized as cropping the autocorrelation function around the zero-lag index, or equivalently by thinning the autocorrelation matrix in Eq. (9) about its diagonal. Retaining only the first n_τ lags, the truncated parameter covariance is

$$\mathrm{cov}(\hat{\boldsymbol{\theta}}_k) = \mathbf{D}_k\left[\sum_{i=0}^{n_\tau}\hat{\mathcal{R}}_k(i)\boldsymbol{\Lambda}_k(i)\right]\mathbf{D}_k \tag{22}$$

Retaining only one term, i.e. $n_\tau = 0$, corresponds to the conventional uncertainty estimate which assumes white residuals. There is no procedure for selecting how many lags to maintain. For batch analysis, $N/5$ lags was suggested.[2] For the real-time analysis, Ref. [6] used 11 lags. This number also varies according to aircraft scale, dynamics of interest, and sampling rate.

III. T-2 Aircraft

The estimation methods were applied to the airplane known as the T-2, which is a 5.5% dynamically scaled version of a generic transport aircraft. The T-2 has retractable tricyle landing gear, two jet engines

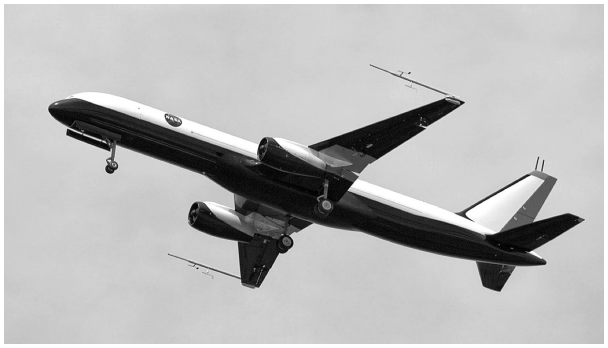

Figure 1. T-2 airplane (credit: NASA Langley Research Center)

mounted under the wings, and a conventional tail configuration. Figure 1 shows a photograph of the airplane in flight, and Table 1 lists airplane geometry and nominal mass properties.

Control surfaces for the T-2 are left and right ailerons, left and right inboard and outboard elevators, upper and lower rudders, left and right inboard and outboard trailing-edge flaps, and left and right inboard and outboard spoilers. This configuration amounts to 16 control surfaces, all of which can be moved independently. For the maneuvers examined here, only the elevators, ailerons, and rudders were deflected. The individual elevator surfaces were moved together as a single elevator surface, and similarly for the rudders, whereas the ailerons were moved asymmetrically, in the conventional way. Trailing-edge down is considered positive deflection for wing and elevator surfaces, and trailing-edge left is positive for rudder surfaces. Definitions of control surface deflections are

$$\delta_e = \frac{1}{4}\left(\delta_{e_{lo}} + \delta_{e_{li}} + \delta_{e_{ri}} + \delta_{e_{ro}}\right) \tag{23a}$$

$$\delta_a = \frac{1}{2}\left(\delta_{a_r} - \delta_{a_l}\right) \tag{23b}$$

$$\delta_r = \frac{1}{2}\left(\delta_{r_u} + \delta_{r_l}\right) \tag{23c}$$

The T-2 was outfitted with a variety of research-quality hardware. A micro-INS provided tri-axial translational accelerometer measurements, angular rate gyroscope measurements, estimated Euler angles, and GPS-derived position and velocity. Air data probes on both wing tips measured angle of attack, sideslip angle, static pressure, and dynamic pressure. Measurements from static pressure sensors and outside air temperature sensors were used to compute air density and altitude. Engine speeds were measured and used in a thrust model which was identified from ground test data and augmented with adjustments for ram drag identified from flight data. Potentiometers on the rotation axes of the control surfaces measured control surface deflections. Mass properties were computed based on measured fuel flow, pre-flight weight and balance, and body-axis inertia measurements done on the ground for the aircraft without fuel. Measurements were sampled at 200 Hz.

The aircraft was flown using the NASA Langley AirSTAR (Airborne Subscale Transport Aircraft Research) flight test facility. This capability includes a mobile control station, which houses telemetry equipment, computational hardware, a pilot station, and stations for flight test personnel and researchers. A research pilot flies the aircraft from inside the control station, using synthetic vision drawn from telemetry data and a database of the local terrain. Control deflections are generated by the pilot and the ground-based flight control system. The AirSTAR flight control system has the capability to inject arbitrary inputs at the actuators, just before position and rate limiters. These inputs were engaged by the pilot pressing and holding a button on the thrust levers. Telemetered data, downsampled to 50 Hz, is available to researchers inside the control station through the internal ethernet network. For more information on AirSTAR, see for example Ref. [10].

For the results examined here, the excitation inputs applied were orthogonal phase-optimized multisines. This type of input was developed at NASA Langley[11,12] and has been used for highly-efficient excitation of

different types of vehicles in a variety of unusual flight conditions. Each multisine input has the form

$$u(t) = a \sum_k a_k \sin\left(\frac{2\pi k}{T} t + \phi_k\right) \tag{24}$$

where a is the aggregate amplitude, a_k are the normalized relative sinusoid amplitudes, T is the excitation record length, and ϕ_k are the phase angles. Choice of the record length determines the available excitation frequencies $\omega_k = 2\pi k/T$. Selection of the sinusoid amplitudes and harmonic indices k can be used to tailor the power spectrum of the inputs. The first few harmonic indices are usually discarded to ensure that the data record contains repeated cycles of each excitation frequency. Phase angles are optimized for minimum relative peak factor (RPF) of the waveform to create small perturbation responses. The aggregate amplitude is used to scale the resulting waveform to achieve good signal-to-noise ratios. Multisine input parameters used here are given in Table 2.

IV. Results

IV.A. Simulation Results

A simulation model of the T-2 longitudinal short-period dynamics was first used to test the method. The state-space perturbation model was

$$\begin{bmatrix} \Delta\dot{\alpha}(t) \\ \Delta\dot{q}(t) \end{bmatrix} = \begin{bmatrix} \frac{\bar{q}S}{mV} C_{Z_\alpha} & 1 \\ \frac{\bar{q}S\bar{c}}{I_{yy}} C_{m_\alpha} & \frac{\bar{q}S\bar{c}}{I_{yy}} \frac{\bar{c}}{2V} C_{m_q} \end{bmatrix} \begin{bmatrix} \Delta\alpha(t) \\ \Delta q(t) \end{bmatrix} + \begin{bmatrix} \frac{\bar{q}S}{mV} C_{Z_{\delta_e}} \\ \frac{\bar{q}S\bar{c}}{I_{yy}} C_{m_{\delta_e}} \end{bmatrix} \Delta\delta_e(t) \tag{25a}$$

$$\begin{bmatrix} \Delta\alpha(t) \\ \Delta q(t) \\ \Delta a_z(t) \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ \frac{\bar{q}S}{mg} C_{Z_\alpha} & 0 \end{bmatrix} \begin{bmatrix} \Delta\alpha(t) \\ \Delta q(t) \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \frac{\bar{q}S}{mg} C_{Z_{\delta_e}} \end{bmatrix} \Delta\delta_e(t) \tag{25b}$$

The reference flight condition simulated was straight and level flight with 134 ft/s airspeed, 1370 ft altitude, and 4.8 deg angle of attack. Values for the stability and control derivatives were taken from a separate analysis using flight test data in calm air, and are given as the true values in Table 3. The simulation model was excited using the elevator multisine specified by Eq. (24) and given in Table 2. Angular accelerations were computed using a fixed-lag smooth differentiation method.[8,2]

The colored noise added to the measurements was simulated by summing together two noise sequences: a wide-band noise contribution from 0 Hz to the 25 Hz Nyquist frequency, and a band-limited contribution from 0–2 Hz. The wide-band component was realized using a normally-distributed pseudorandom number generator. Signal-to-noise ratios for the elevator, angle of attack, pitch rate, and vertical acceleration were 40, 12, 30, and 40, respectively. These values were also determined from flight test data in calm air. The band-limited component was realized by passing a different white Gaussian noise sequence through a fifth-order Chebyshev filter with a 2 Hz corner frequency. This method of generating colored noise is known to be representative of typical flight test data.[4,5] This range is appropriate for this scale of aircraft because it extends slightly beyond the rigid body dynamics at 1.1 Hz. The colored noise therefore contains the dynamics of interest and the bandwidth of the input, and is expected to color the modeling residuals. Simulations were run for band-limited noise set between 0% and 20%, in 5% increments, where the percentage pertains to the root mean square of the signal variation. Note that 0% band-limited noise means that only wide-band Gaussian white noise was added the measurements, and that the model residuals should then be white. Band-limited noise levels greater than 20% were not attempted because these are generally not representative of aircraft system identification scenarios. Simulations were repeated 250 times for each level of band-limited noise using unique noise sequences each time.

A summary of parameter estimation results, obtained at the end of the maneuver and averaged over all 250 simulation runs for each level of band-limited noise, is given in Table 3. For 0% band-limited noise, the parameter estimates are accurate, parameter uncertainties are small, and results are in statistical agreement with the true values. As the level of band-limited noise increases, the mean estimates become more biased and the mean parameter uncertainties increase. For all cases with band-limited noise, the corrected uncertainties match the observed scatter in the parameter estimates well, and does so more closely than the conventional uncertainty estimates that assume white residuals.

In the remaining discussion, results will be illustrated for the 20% band-limited noise case, because it represents a somewhat exaggerated scenario of what is usually encountered. Of these results, the C_{Z_α} estimates will be highlighted, because it is an important parameter that is usually estimated well. For this case, the conventional standard error was too small by a factor of four, whereas the corrected standard error was consistent with the observed scatter in the estimate. The measurements from the last simulation run are shown in Fig. 2.

The recursive least-squares fits for C_Z and C_m in this simulation example are shown in Fig. 3. The fits had coefficients of determination of 0.93 and 0.87, respectively, which indicate good fits considering this level of noise. The model residuals are shown in Fig. 4, first as time histories and then as power spectra. The time histories show that the residual is small, and contains both coherent and incoherent content. The power spectra show the the larger band-limited noise, its roll off at 2 Hz, and the wide-band noise floor at higher frequencies. These characteristics are particularly clear for the vertical force coefficient. Some amplification of the residual around 10 Hz is visible in the pitching moment residual power spectrum, which is due to the real-time differentiation of angular rate data. Overall, these plots agree with those shown later for the flight test results and are similar to those reported previously,[4,5] which indicates a realistic simulation.

The autocorrelation estimates for the C_Z coefficient residual are shown in Fig. 5. For clarity, only positive lags are shown since autocorrelation is an even function. Figure 5(a) shows the time history, with a decreased temporal resolution for additional clarity. Figure 5(b) shows the final recursive estimate, along with the batch estimate and the $\pm 2\sigma$ uncertainty bound of the autocorrelation, given[2] as $\hat{\mathcal{R}}(0)/\sqrt{N}$. If the residual were white, the autocorrelation would have a large peak at the zero-lag index equal to σ^2 and be zero otherwise. Because there are lagged autocorrelations outside the $\pm 2\sigma$ bound, this is not the case and the residual is colored. Most of the residual coloring is due to the large autocorrelations near the zero-lag index. Note the similarity between the final recursive estimates and the batch estimate, particularly for the higher autocorrelations. This indicates that not much accuracy was scarified by using a recursive method.

The evolving time history of estimation results for the model parameter C_{Z_α}, averaged over all 250 simulation runs, is shown in Fig. 6(a). On average, the parameter estimate converged by 2 s, which was 1.5 s after the elevator excitation began. By 6 s, there was enough data information content and good enough parameter estimates that the corrected parameter uncertainty envelope enclosed the observed scatter in the parameter estimate. This means that on average, the recursive estimate is slightly conservative, which is safer in modeling applications than having it too optimistic. The conventional estimate of the parameter uncertainty was too small by a factor of 3.3, and was not in statistical agreement with the true value. Similar time histories were observed for the other model parameters. Batch and recursive estimates, obtained at the end of the maneuver, are compared for conventional and corrected parameter uncertainties in Fig. 6(b). In both cases, the corrected uncertainties are in statistical agreement with the true value, whereas the conventional estimates are not. There was a 1% difference in the size of the corrected error bounds, again meaning that not much accuracy is sacrificed by using a recursive technique.

Figure 7(a) shows the impact of reducing the number of lags retained in the residual autocorrelation. Using only the zero lag, the parameter uncertainty is too small. Only two lags were required for statistical agreement with the true value in this example, although more lags were needed in other runs having the same level of band-limited noise. For roughly 10–200 reatined lags, the parameter covariance overshoots its final value by about 12%. A safe number of lags to retain would be about 50. This would reduce computation and memory, safeguard against selecting too few lags, and provide a conservative estimate of the parameter covariance. Figure 7(b) illustrates the autocorrelation time history with 50 lags. A timing test conducted on a 2.6 GHz Intel Core i7 laptop computer using MATLAB® R2015a showed that 20% of the available 50 Hz frame rate was needed to compute the full autocorrelation function in real time for this 12 s maneuver, but only 4% of the frame rate was needed when the autocorrelation was truncated to 50 lags. Note that this number is dependent upon the aircraft, sampling rate, and the dynamics of interest. For example, a full scale aircraft and/or lower sampling rates would require many fewer lag terms, and thus less computation, to achieve the same level of accuracy.

These results were obtained using degraded residuals from intermediate parameter estimates that were not fully converged, and an approximated recursive correction for colored residuals. In the past, these two factors have been primary obstacles to successfully applying recursive least squares in aircraft system identification. The method presented in Ref. [6] and used here is a simple and effective method of mitigating those previous challenges.

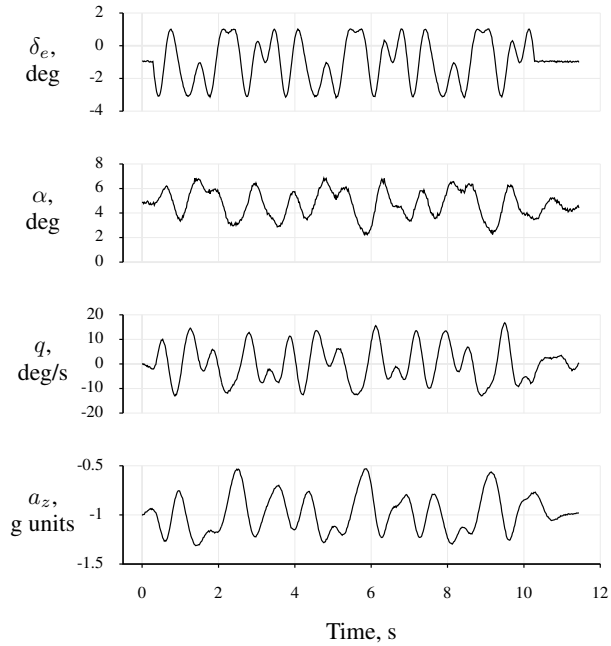

Figure 2. Measurements from simulation data (20% band-limited noise, Run 250)

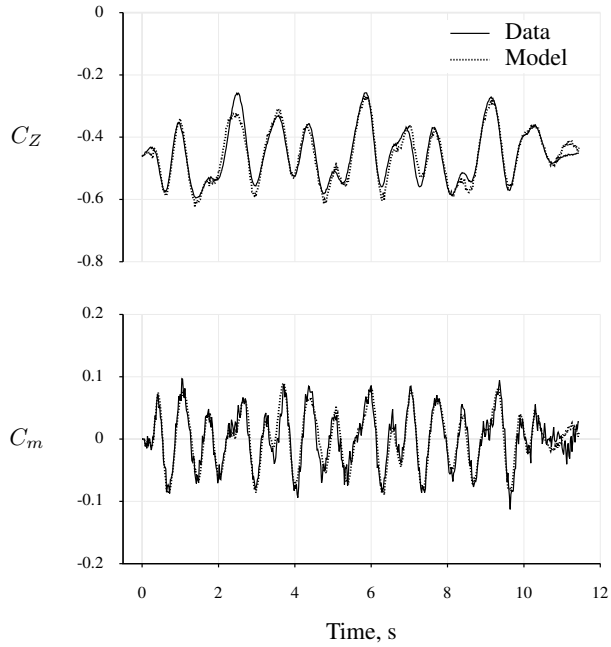

Figure 3. Recursive fitting using simulation data (20% band-limited noise, Run 250)

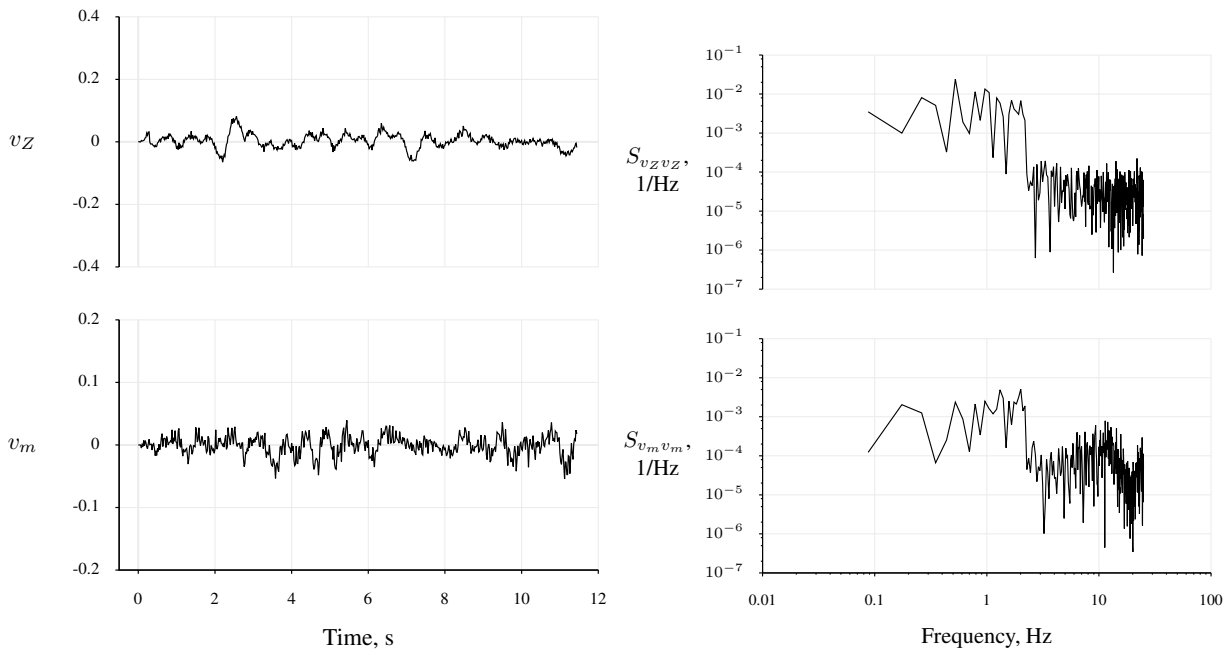

(a) time series (b) power spectrum

Figure 4. Model residuals using simulation data (20% band-limited noise, Run 250)

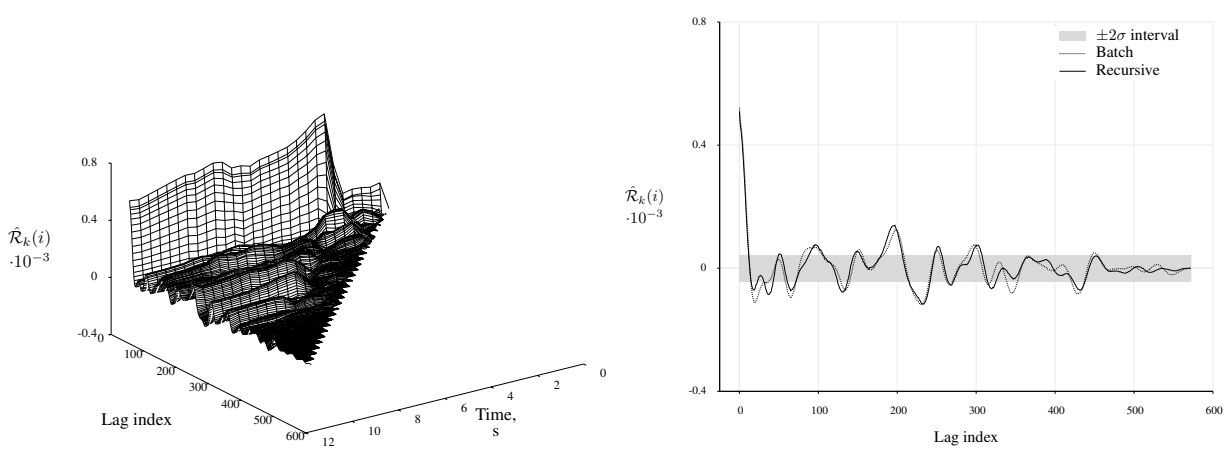

(a) time history (b) final estimate

Figure 5. Autocorrelation of C_Z residual using simulation data (20% band-limited noise, Run 250)

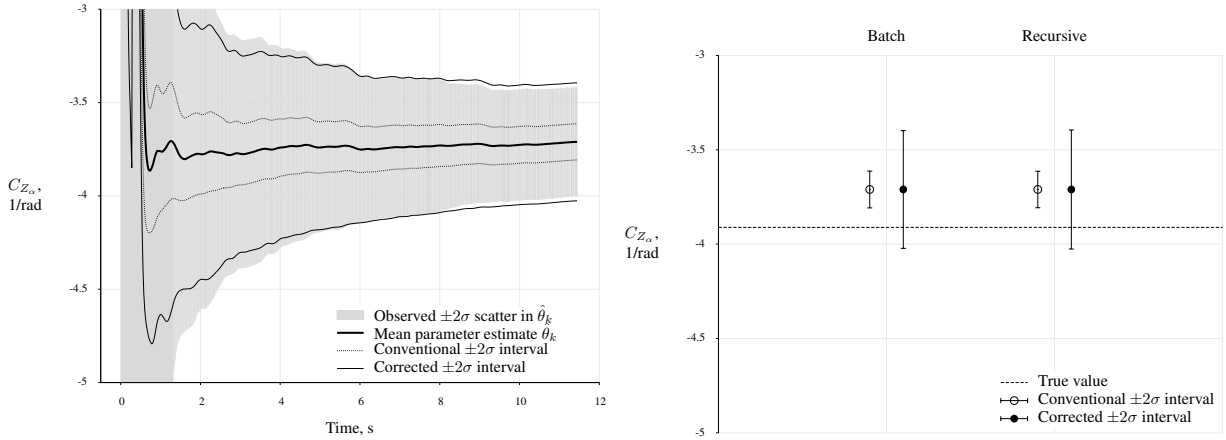

(a) time history (b) final estimates

Figure 6. Mean estimation results of C_{Z_α} using simulation data (20% band-limited noise)

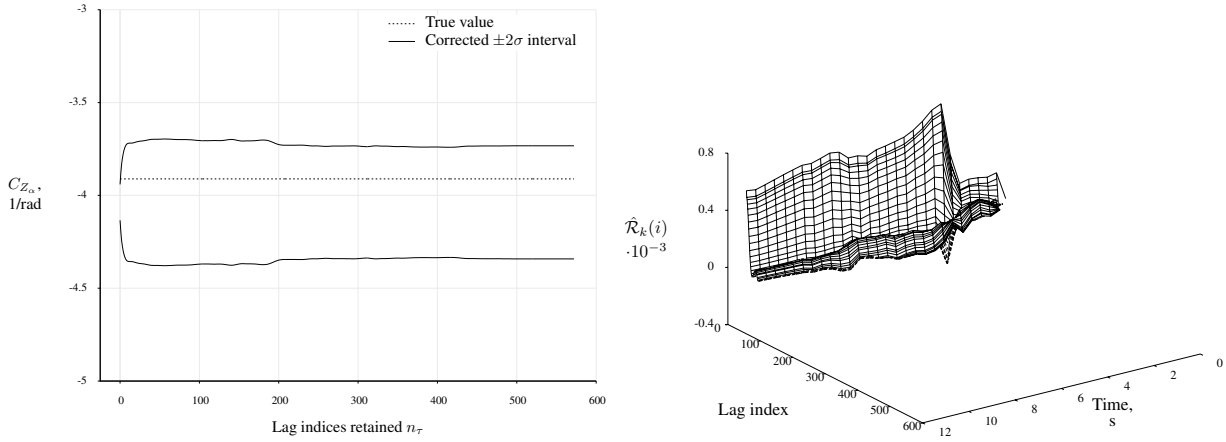

(a) Effect on C_{Z_α} parameter covariance (b) autocorrelation of C_Z residual time history, n_τ=50

Figure 7. Truncation of C_Z residual autocorrelation on C_{Z_α} estimate (20% band-limited noise, Run 250)

IV.B. Flight Test Results

Flight test data for the T-2 airplane was also used to demonstrate the method. A set of 14 repeated maneuvers were analyzed from Flight Number 15, conducted on 23 April 2010 at Blackstone Army Airfield in Virginia. These maneuvers were conducted under severe turbulence. As turbulence manifests as an unmeasured input and excites unmodeled dynamics, the residuals were expected to be colored. The trim condition was approximately the same as used for the simulation, and the same excitation input was commanded of the actuators. The measured flight test data from Maneuver 14, which is highlighted as a representative example, is shown in Fig. 8. Some attenuation of the higher frequencies is seen in the elevator deflection due to the actuators and flight control software. The responses contain higher frequency content, due to the turbulence, but amplitudes are about the same as in the simulation.

Recursive least-squares fits to the aerodynamic coefficients, shown in Fig. 9, were generally good and had coefficients of determination of 0.98 and 0.90 for C_Z and C_m, respectively. Model residuals are shown in Fig. 10, again as time series and power spectra. The flight data residuals have approximately the same amplitude wide-band noise, but less band-limited noise than used in the simulation example. Residuals were again small relative to the measurements, and displayed the same spectral characteristics. The recursive autocorrelation for this data is shown in Fig. 11.

Parameter estimation results, obtained at the end of the maneuver, are shown for all 14 maneuvers in Fig. 12 for both conventional and corrected uncertainties. The corrected parameter uncertainties exhibit better statistical agreement than the conventional uncertainties. Only 50 lags were retained for the corrected parameter uncertainties, based on the simulation analysis. This truncation resulted in a 13% error in the parameter standard error. However, this truncation enabled faster computation for real-time execution, and did not significantly impact the results because they remained in agreement with the observed scatter of the parameter estimates. The time history of the C_{Z_α} estimate, using 50 lags, is shown in Fig. 13. Conventional estimates at the end of the maneuver were smaller by a factor of 3.

A summary of mean results is given in Table 4, where each table entry was averaged over all 14 maneuvers. The columns show the mean estimates, the uncertainties using the conventional recursive method, the uncertainties using the corrected recursive method (both for all the available lags and for only 50 lags), and the observed scatter in the parameter estimates. The corrected uncertainties were larger than the conventional uncertainties, and were more accurate. Not much error was incurred from approximating the autocorrelation with 50 lag terms, as seen by comparing the two columns of corrected results. The corrected uncertainties agreed with the observed scatter in the parameter estimates, except for the angle of attack derivatives, which were too low by a factor of two. However, these were still more accurate than the conventional estimates, which were too low by a factor of four. This discrepancy could be due to the fact that the turbulence most directly affects angle of attack, and that only a small ensemble of maneuvers was used. However, the overall approach is verified here using the flight test data. Notice that the values of the mean estimates, particularly the C_Z derivatives, are significantly different than those estimated from flight test data in calm air. This does not disprove the method, because the estimated uncertainty matches the observed scatter. For example, the elevator lift slope $C_{Z_{\delta_e}}$ decreases from +0.215 in calm air data to +0.031 in the turbulent data. This parameter typically has low sensitivity, making it harder to estimate accurately. In addition, the effects of the elevator on lift and the turbulence (which is not measured) on lift are no longer easily distinguishable to the estimator. This is a problem with the experiment and model structure, and not with the recursive uncertainty correction presented.

V. Conclusions

This paper demonstrated a recursive correction to conventional parameter covariances for recursive least-squares estimates of aircraft stability and control parameters considering colored residuals. Colored residuals are not assumed in the underlying theory of ordinary least squares, but are routinely encountered in aircraft system identification from various sources of model structure error. The method was demonstrated using simulation runs and flight test data from repeated maneuvers using the T-2 airplane.

The discrepancy between predicted uncertainty levels and observed scatter in parameter estimates was identified in Ref. [4]. A solution was developed in Ref. [5] for batch analysis. Reference [6] extended this method for recursive estimation, and employed it in frequency response identification. This paper demonstrated that method for estimating stability and control derivatives, which is a common application of aircraft parameter estimation. A comparison was also made with batch and conventional estimates, and

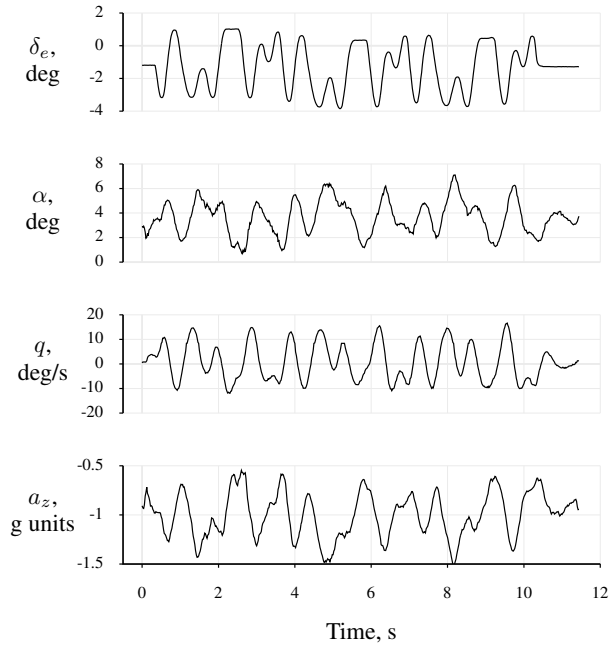

Figure 8. Measurements from flight test data (Maneuver 14)

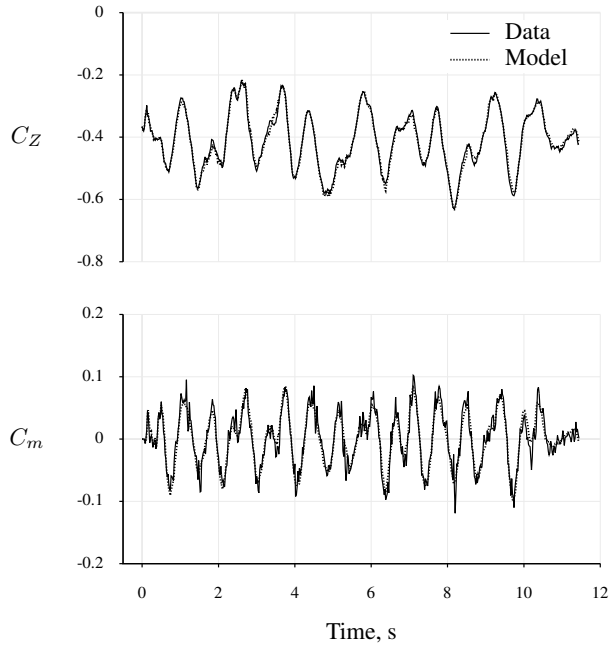

Figure 9. Recursive fitting using flight test data (Maneuver 14)

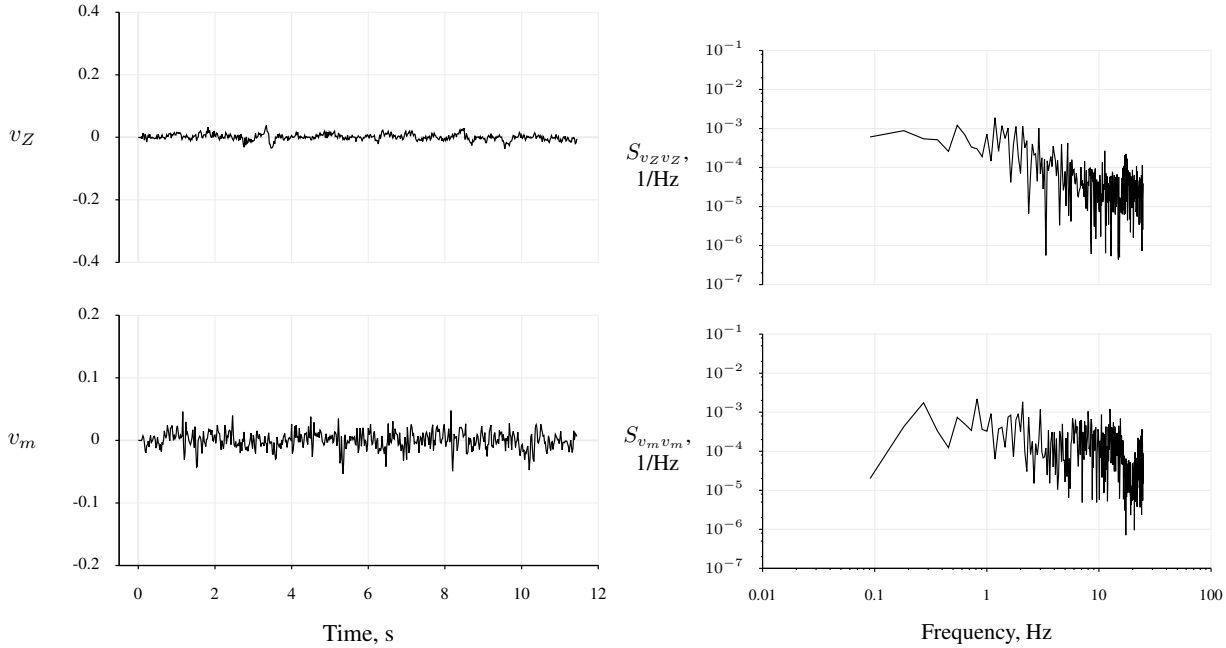

(a) time series (b) power spectrum

Figure 10. Model residuals using flight test data (Maneuver 14)

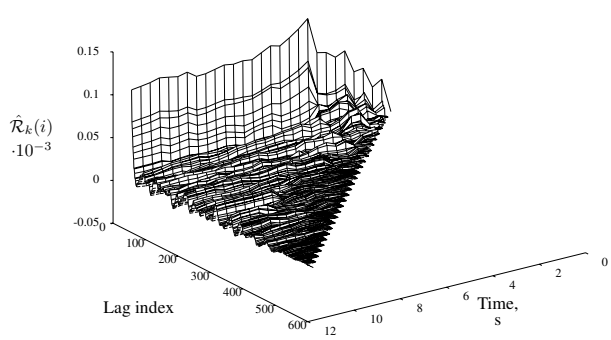

Figure 11. Autocorrelation of C_Z residuals using flight test data (Maneuver 14)

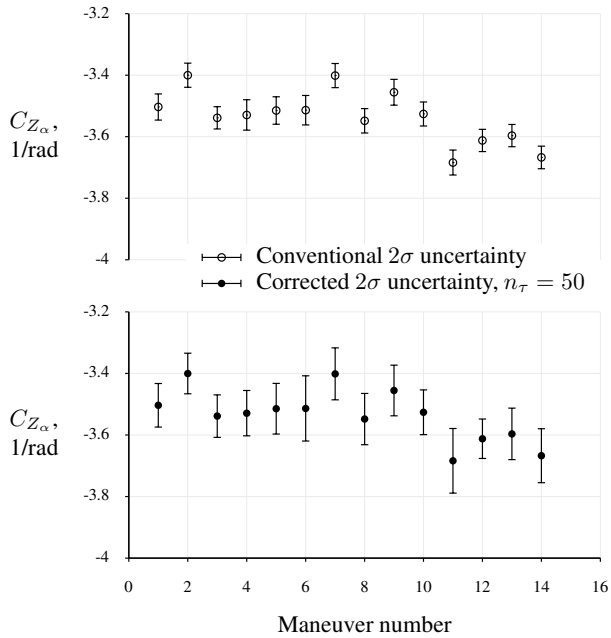

Figure 12. Final estimation results for C_{Z_α} using flight test data

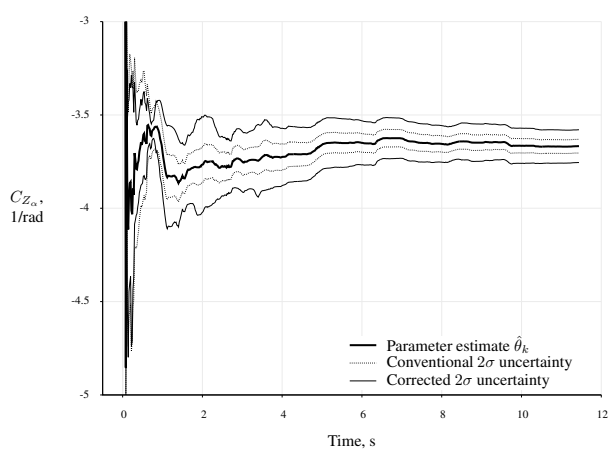

Figure 13. Estimation time history results for C_{Z_α} using flight test data (Maneuver 14)

the effects of approximating the autocorrelation function by truncating the summation were explored.

The short period simulation results showed that the method obtains accurate estimates of the parameter covariances in the presence of colored residuals. Parameter uncertainties need to be corrected when colored residuals are present, which is usually the case in practice. Errors in the parameter standard deviations up to a factor of 3.6 were observed using conventional parameter uncertainty estimates when the band-limited noise was 20% of the root mean signal variation. Only 1% error was observed by correcting the errors recursively, instead of in a batch analysis at the end of the maneuver. Timing tests showed that for longitudinal short-period dynamic maneuvers, the method can run in real time in MATLAB® on a common laptop computer. Computations were computed 5 times faster when the autocorrelation function was limited to 50 lag terms, compared to using all available lag terms. This truncation resulted in a 13% error in the parameter uncertainty, but flight test data with repeated maneuvers showed this was still statistically consistent with the observed parameter estimate scatter. Flight test data also showed that the conventional estimates of the parameter covariance were optimistic and inaccurate, as expected.

A significant negative attribute of the recursive least squares algorithm in the past has been inaccuracy of the computed parameter uncertainty. This work successfully addresses that problem, removing a major barrier to using recursive least squares for important practical problems such as adaptive control, fault detection, and simulation database validation and updating from flight data. Retaining a small number of lag indices and computing a recursive autocorrelation function can lead to significant gains in accuracy in terms of the parameter covariance, which is an important indicator of model quality.

Acknowledgments

The efforts of the AirSTAR flight test team at NASA Langley in building and testing the T-2 aircraft and associated systems, carefully calibrating the instrumentation, and conducting flight operations to collect the high-quality flight data used in this study, are gratefully acknowledged. The system identification research was supported by the NASA Aeronautics Research Mission Directorate (ARMD) Learn-To-Fly project.

References

[1]Morelli, E., "Real-Time Parameter Estimation in the Frequency Domain," *Journal of Guidance, Control, and Dynamics*, Vol. 23, No. 5, September–October 2000, pp. 812–818.

[2]Klein, V. and Morelli, E., *Aircraft System Identification: Theory and Practice*, AIAA Education Series, AIAA, Reston, VA, 2006.

[3]Chandler, P., Pachter, M., and Mears, M., "System Identification for Adaptive and Reconfigurable Control," *Journal of Guidance, Control, and Dynamics*, Vol. 18, No. 3, May–June 1995, pp. 516–524.

[4]Maine, R. and Iliff, K., "Use of Cramer-Rao Bounds on Flight Data with Colored Residuals," *Journal of Guidance and Control*, Vol. 4, No. 2, March–April 1981, pp. 207–213.

[5]Morelli, E. and Klein, V., "Accuracy of Aerodynamic Model Parameters Estimated from Flight Test Data," *Journal of Guidance, Control, and Dynamics*, Vol. 20, No. 1, January–February 1997, pp. 74–80.

[6]Holzel, M. and Morelli, E., "Real-Time Frequnecy Response Estimation from Flight Data," *Journal of Guidance, Control, and Dynamics*, Vol. 35, No. 5, September–October 2012, pp. 1406–1417.

[7]Iliff, K. and Maine, R., "Uses of Parameter Estimation in Flight Test," *Journal of Aircraft*, Vol. 20, No. 12, December 1983, pp. 1043–1049.

[8]Morelli, E., "Practical Aspects of the Equation-Error Method for Aircraft Parameter Estimation," No. 2006–6144 in Atmospheric Flight Mechanics Conference, AIAA, Keystone, CO, August 2006.

[9]Maine, R. and Iliff, K., "Application of Parameter Estimation to Aircraft Stability and Control: The Output-Error Approach," Tech. Rep. RP–1168, NASA, Edwards, CA, June 1986.

[10]Jordan, T. and Bailey, R., "NASA Langley's AirSTAR Tetsbed – A Subscale Flight Test Capability for Flight Dynamics and Control System Experiments," No. 2008–6660 in Guidance, Navigation, and Control Conference, AIAA, Honolulu, HI, August 2008.

[11]Morelli, E., "Multiple Input Design for Real-Time Parameter Estimation in the Frequency Domain," No. REG–360 in 13th IFAC Symposium on System Identification, IFAC, Rotterdam, The Netherlands, 2003.

[12]Morelli, E., "Flight Test Experiment Design for Characterizing Stability and Control of Hypersonic Vehicles," *Journal of Guidance, Control, and Dynamics*, Vol. 32, No. 3, May–June 2009, pp. 949–959.

Tables

Table 1. T-2 geometry and nominal mass properties

Parameter	Value	Unit
b	6.849	ft
$\bar{c}$	0.915	ft
S	5.902	ft^2
x_{cm}	56.63	in
y_{cm}	0.000	in
z_{cm}	11.43	in
m	1.585	slug
I_{xx}	1.179	slug·ft^2
I_{yy}	4.520	slug·ft^2
I_{zz}	5.527	slug·ft^2
I_{xz}	0.211	slug·ft^2

Table 2. Multisine input design for the T-2 ($T = 10$ s)

Input	a, deg	RPF	k	f_k, Hz	a_k, deg	ϕ_k, rad
			3	0.3	0.316	2.948
			6	0.6	0.387	0.601
			9	0.9	0.447	3.584
δ_e	1.0	1.03	12	1.2	0.447	4.632
			15	1.5	0.387	2.690
			18	1.8	0.316	2.087
			21	2.1	0.316	3.421
			4	0.4	0.378	1.544
			7	0.7	0.378	4.642
			10	1.0	0.378	1.201
δ_a	1.0	1.15	13	1.3	0.378	1.077
			16	1.6	0.378	3.946
			19	1.9	0.378	3.951
			22	2.2	0.378	3.523
			2	0.2	0.316	2.844
			5	0.5	0.387	2.526
			8	0.8	0.447	2.756
δ_r	1.0	1.14	11	1.1	0.447	5.770
			14	1.4	0.387	5.540
			17	1.7	0.316	2.396
			20	2.0	0.316	5.525

Table 3. Mean estimation results using simulated data and varying levels of band-limited noise

Parameter	True Value	Band-limited noise level 0%	5%	10%	15%	20%
C_{Z_α}	-3.911	-3.888	-3.870	-3.841	-3.774	-3.710
		(0.013)[a]	(0.019)	(0.029)	(0.038)	(0.048)
		[0.015][b]	[0.049]	[0.088]	[0.123]	[0.158]
		{0.012}[c]	{0.049}	{0.085}	{0.129}	{0.146}
$C_{Z_{\delta_e}}$	$+0.215$	$+0.225$	$+0.227$	$+0.232$	$+0.243$	$+0.260$
		(0.011)	(0.017)	(0.025)	(0.034)	(0.043)
		[0.013]	[0.042]	[0.075]	[0.106]	[0.138]
		{0.012}	{0.039}	{0.075}	{0.108}	{0.138}
C_{m_α}	-1.481	-1.521	-1.515	-1.507	-1.486	-1.468
		(0.021)	(0.023)	(0.025)	(0.029)	(0.033)
		[0.014]	[0.027]	[0.044]	[0.060]	[0.079]
		{0.007}	{0.026}	{0.041}	{0.063}	{0.074}
C_{m_q}	-53.25	-54.73	-54.28	-53.15	-51.94	-49.90
		(1.581)	(1.670)	(1.863)	(2.095)	(2.349)
		[1.149]	[2.230]	[3.674]	[5.052]	[6.485]
		{0.598}	{2.009}	{3.447}	{4.806}	{6.478}
$C_{m_{\delta_e}}$	-1.830	-1.912	-1.905	-1.885	-1.864	-1.834
		(0.029)	(0.030)	(0.034)	(0.039)	(0.044)
		[0.021]	[0.041]	[0.067]	[0.092]	[0.119]
		{0.011}	{0.038}	{0.068}	{0.090}	{0.122}

[a]Values in parenthesis are conventional parameter standard errors
[b]Values in brackets are corrected parameter standard errors
[b]Values in braces are parameter estimate observed scatter standard deviations

Table 4. Mean estimation results using flight test data

Parameter	Mean estimate	Standard errors and deviation: Conventional	Corrected, $n_\tau \leq N-1$	Corrected, $n_\tau = 50$	Scatter
C_{Z_α}	-3.535	0.020	0.040	0.047	0.085
$C_{Z_{\delta_e}}$	$+0.031$	0.024	0.049	0.059	0.050
C_{m_α}	-1.275	0.026	0.043	0.046	0.121
C_{m_q}	-44.58	2.389	3.985	4.138	3.829
$C_{m_{\delta_e}}$	-1.782	0.045	0.074	0.078	0.072