JOURNAL OF AIRCRAFT
Vol. 33, No. 5, September–October 1996

In-Flight Demonstration of a Real-Time Flush Airdata Sensing System

Stephen A. Whitmore,* Roy J. Davis,† and John Michael Fife‡
NASA Dryden Flight Research Center, Edwards, California 93523

A prototype real-time flush airdata sensing system has been developed and flight tested at the NASA Dryden Flight Research Center. This system uses a matrix of pressure orifices on the vehicle nose to estimate airdata parameters in real time using nonlinear regression. The algorithm is robust to sensor failures and noise in the measured pressures. The real-time flush airdata sensing (FADS) system has been calibrated using inertial trajectory measurements that were bootstrapped for atmospheric conditions using meteorological data. Mach numbers as high as 1.6 and angles of attack greater than 45 deg have been tested. The system performance has been evaluated by comparing the real-time FADS to the ship system airdata computer measurements to give a quantitative evaluation relative to an accepted measurement standard. Nominal agreements of approximately 0.003 in Mach number and 0.20 deg in angle of attack and angle of sideslip have been achieved.

Introduction

SPECIALIZED requirements of advanced vehicles make use of conventional intrusive airdata measurement systems[1] highly undesirable. For example, on the X-31 Enhanced Maneuverability Fighter aircraft, the presence of the airdata noseboom caused unsteadiness in the primary forebody vortex cores and induced lateral instabilities at high angles of attack. These instabilities resulted in degraded aircraft handling qualities and, in the worst cases, induced aircraft departures.[2] In other applications, such as hypersonic aircraft, the hostility of the hypersonic environment mandates the use of nonintrusive airdata systems. On stealth vehicles where a minimal radar cross section is required, such as the B-2 bomber or the F-22 fighter, conventional intrusive systems are highly visible. Eliminating these systems from the basic vehicle design is desirable.

The flush airdata sensing (FADS) system concept was developed as a means of circumventing many of the aforementioned difficulties with intrusive airdata systems. Using this concept, airdata are inferred from nonintrusive surface pressure measurements. The original system prototype was developed for the X-15 program and used a hemispherical nose that was actively steered into the local relative wind vector to measure stagnation pressure and flow incidence angles.[3] The mechanical design of this system was extremely complicated, and the steered-nose concept was abandoned after the X-15 program ended.

A more modern approach, the Shuttle entry airdata system, was developed at the NASA Langley Research Center for the Space Shuttle program.[4] This approach used a matrix of fixed static pressure measurements and no mechanical actuation of the nose was required. The Shuttle entry airdata system technique later was adapted to aeronautical applications, and several demonstration programs were performed in the early 1980s at the NASA Dryden Flight Research Center.[5,6] For these early programs, measurement and presentation of individual pressure coefficient data and their empirical relationships to the airdata parameters were emphasized. These tests verified the feasibility of the fixed-orifice concept, but did not attempt to derive algorithms for estimating the airdata from the pressure measurements.

A more advanced program, the high-angle-of-attack flush airdata sensing (HI-FADS) system, was developed and has recently concluded flight testing at the NASA Dryden Flight Research Center.[7,8] The system was required for the F-18 High Alpha Research Vehicle flight tests because the noseboom installation altered the flow characteristics on the aircraft nose at high incidence angles and adversely affected the vehicle flight dynamics. The HI-FADS design, as with the earlier fixed-orifice FADS systems, used a matrix of flush static-pressure orifices arranged on the nose of the vehicle. This design, however, incorporated the pressure measurements into an overdetermined estimation algorithm where all surface pressure observations were used simultaneously to infer the airdata parameters using nonlinear regression.

For the High Alpha Research Vehicle flight tests, the HI-FADS computations were performed postflight using pressure data telemetered to the ground. To allow autonomous operation as part of an actual flight system, the HI-FADS algorithm was integrated into a real-time system that included pressure sensors, computational hardware, onboard program data storage, and interface to the aircraft instrumentation system. This system, the real-time flush airdata sensing (RT-FADS) system, was flight tested on the NASA Dryden Flight Research Center's F-18 Systems Research Aircraft (SRA). This article describes the RT-FADS measurement system, including the basic measurement hardware, the airdata parameter estimation algorithm, and redundancy management schemes that ensure algorithm tolerance to sensor failures. System calibration methods and evaluations of the system performance for subsonic, high angle of attack, and supersonic flight regimes are presented.[9]

Presented as Paper 95-3433 at the AIAA Flight Simulation Technologies Conference, Baltimore, MD, Aug. 7–10, 1995; received Jan. 28, 1996; revision received May 17, 1996; accepted for publication May 28, 1996.

*Vehicle Dynamics Group Leader, Aerodynamics Branch, M/S D-2033, P.O. Box 273. e-mail: whitmore@wilbur.dfrc.nasa.gov. Senior Member AIAA.

†Research Engineer; currently at Lockheed-Sanders Corporation, P.O. Box 868, Nashua, NH 03061-0868. e-mail: rjdavis1@rapnet.sanders.lockheed.com.

‡Student Intern, Aerodynamics Branch, M/S D-2033; currently Graduate Student, Massachusetts Institute of Technology, M/S 37-427, 77 Massachusetts Avenue, Cambridge, MA 02139. e-mail: fife@orville.mit.edu. Student Member AIAA.

Real-Time Flush Airdata Sensing System Hardware

Figure 1 shows a basic system overview. The various hardware components, including the flight test radome, the RT-FADS pressure port matrix, and the measurement transducers, are described in the following subsections.

Radome Configuration

The transducers and the electrical interface unit are mounted on a palette in the SRA radome. The radome for the RT-FADS system is a preproduction unit with a composite matrix nosecap with 11 pressure orifices molded into the structure and attached at the radome tip and faired flush to the surface. The locations of the nosecap ports are defined using cone and clock coordinate angles, ϕ and λ, respectively. Figure 2 shows the clock and cone angles and the pressure port locations on the radome.

Data collection and algorithm computations are performed by two commercially available Motorola 68040-based single-board computers inserted in a flight-ruggedized Versa-Module Eurocard (VME) chassis (Fig. 1). The two processor boards act as slave and master. The master processor acts as the system controller (SC) and manages the data flow through the system. The SC services the slave RT-FADS processor, which in turn communicates to measurement transducers and performs the airdata calculations. Outputs from the RT-FADS processor are passed to the master processor through the VME backplane. The SC processor communicates with the aircraft data system through a 1553 bus controller card and a dedicated bus to the vehicle instrumentation system.

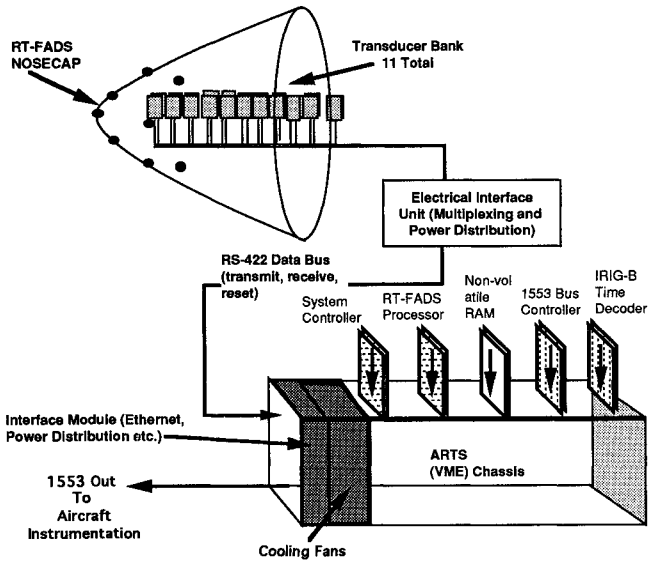

Fig. 1 RT-FADS/ARTS system hardware schematic.

Measurement Transducers

Pressures at the FADS ports are sensed by 11 miniaturized, digital absolute-pressure transducers. Each absolute-pressure transducer incorporates a four-arm active strain-gauge bridge for data sensing, internal signal conditioning with a 20-Hz anti-aliasing low-pass filter and output signal amplification, and a 20-bit A/D conversion. The transducers have a repeatability that exceeds +0.01% of full scale and have a measurement range from 1.50 to 40.00 psia. The digital transducers use non-standard asynchronous serial communication, logically similar to the industry-standard RS-485/422/232 protocol.[10] This protocol allows all of the transducers to be connected through a single common-data bus.

System Interface

As previously mentioned, the RT-FADS algorithm is loaded into the modified RT-FADS slave processor and communicates with the SC processor through the backplane. Each computational cycle is performed as rapidly as transducer communications and airdata computations allow. At the end of each computational cycle, the RT-FADS processor sets a flag on the backplane and waits for acknowledgment from the SC processor. Acknowledgment instructs the algorithm to proceed with a new set of computations, the transducers are polled again, and a new set of computations are performed. The SC processor performs time synchronization with the data rate currently variable from 25 to 100 samples/s. FADS throughput rates as high as 50 samples/s were achieved.

Aerodynamics Model

A pressure model is used to relate the pressure measurements to airdata quantities. The model prescribes measured pressure in terms of four airdata parameters: 1) dynamic pressure q_c, 2) angle of attack α, 3) angle of sideslip β, and 4) static pressure P_∞. Using these four basic parameters, other airdata quantities of interest may be directly calculated. As previously reported,[4] the measured pressure data at the ith pressure port is related to the desired airdata parameters by the simple model

$$P_i = q_c[\cos^2(\theta_i) + \varepsilon \sin^2(\theta_i)] + P_\infty \tag{1}$$

In Eq. (1), θ_i is the flow incidence angle between the surface at the ith port and the velocity vector, and ε is an empirically

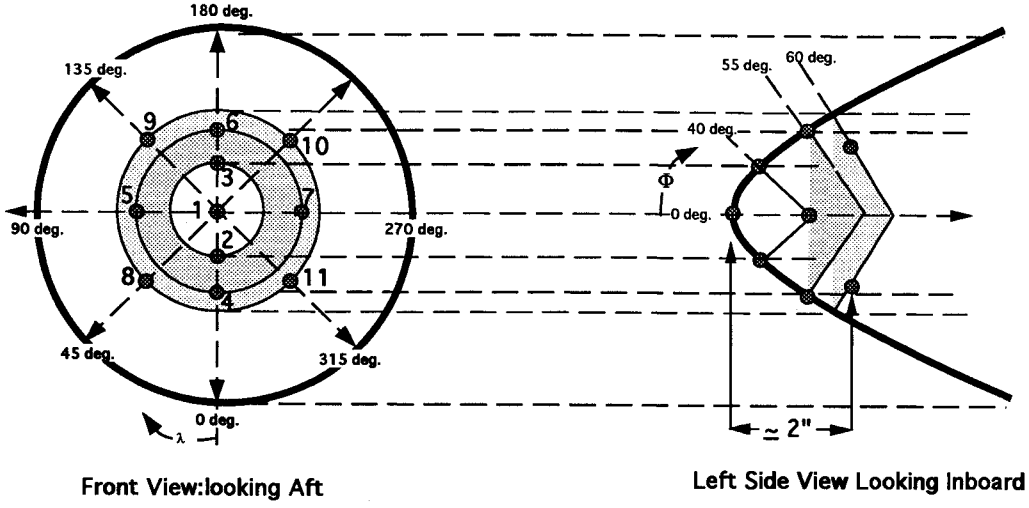

Fig. 2 Schematic of RT-FADS nosecap showing coordinate definitions and port locations.

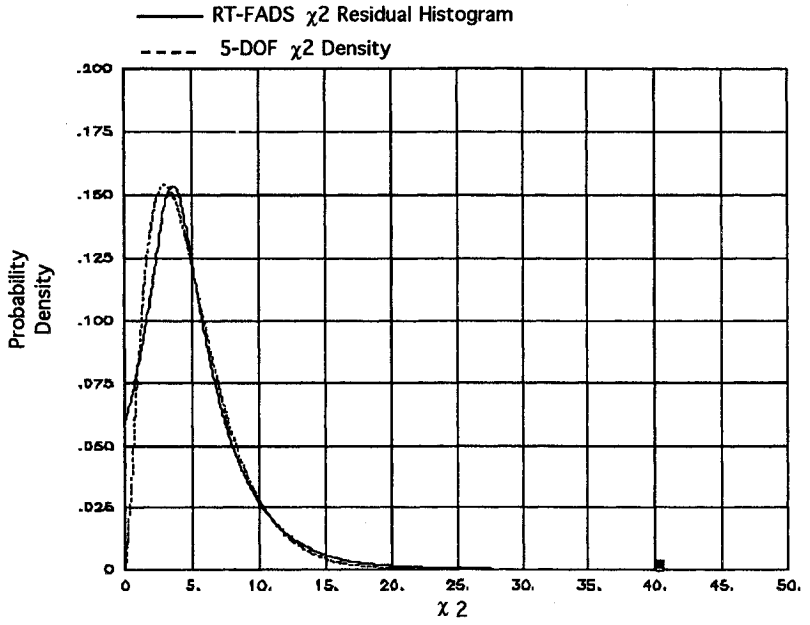

Fig. 3 Residual distribution χ^2 for converged algorithm.

determined calibration parameter that adjusts for the static source error caused by the presence of the aircraft afterbody and/or the supersonic bow shock wave. The incidence angle is related to the local angle of attack and angle of sideslip by

$$\cos(\theta_i) = \cos(\beta)[\cos(\alpha)\cos(\phi_i) + \sin(\alpha)\cos(\lambda_i)\sin(\phi_i)] + \sin(\beta)\sin(\lambda_i)\sin(\phi_i) \quad (2)$$

The parameters α and β are local flow angles whose values are influenced by the aircraft-induced wash. Their relationships to the freestream values must also be calibrated. Values for the calibration parameters are presented in the Flight Test Results section.

Nonlinear Regression Algorithm

Because the aerodynamic model is nonlinear and cannot be directly inverted to allow calculation of airdata as a function of the measured pressures, the measurements must be used to indirectly infer the airdata state using a nonlinear least-squares regression. Within each computational cycle, the algorithm is linearized about a starting airdata value for each port location. Linearizing for all 11 pressure observations, the system may be written as a matrix equation:

$$\begin{bmatrix} \delta P_1 \\ \cdots \\ \delta P_{11} \end{bmatrix} = \begin{bmatrix} \frac{\partial P_1}{\partial \alpha} & \frac{\partial P_1}{\partial \beta} & \frac{\partial P_1}{\partial q_c} & \frac{\partial P_1}{\partial p_\infty} \\ & \cdots\cdots\cdots & & \\ \frac{\partial P_{11}}{\partial \alpha} & \frac{\partial P_{11}}{\partial \beta} & \frac{\partial P_{11}}{\partial q_c} & \frac{\partial P_{11}}{\partial p_\infty} \end{bmatrix} \begin{bmatrix} \delta\alpha \\ \delta\beta \\ \delta q_c \\ \delta p_\infty \end{bmatrix} + \begin{bmatrix} \nu_1 \\ \cdots \\ \nu_{11} \end{bmatrix} \quad (3)$$

In Eq. (3), $\nu_1 \ldots \nu_{11}$, represent the unmodeled errors in the aerodynamic model. This overdetermined system of perturbation equations is solved using the weighted least-squares technique.[11] At the end of least-squares regression, the resulting perturbation is added to the starting value and the system is relinearized about the resulting update. The iteration is repeated until algorithm convergence is reached, typically in two cycles, but as many as eight cycles are allowed. The specific criterion used to determine algorithm convergence is discussed later in the redundancy management module description. At the beginning of each new computational cycle, the system of equations is relinearized about the result of the previous cycle and the iteration is repeated using new pressure data. Extensive development of the regression algorithm, and methods for starting up the regression, have previously been reported.[7-9]

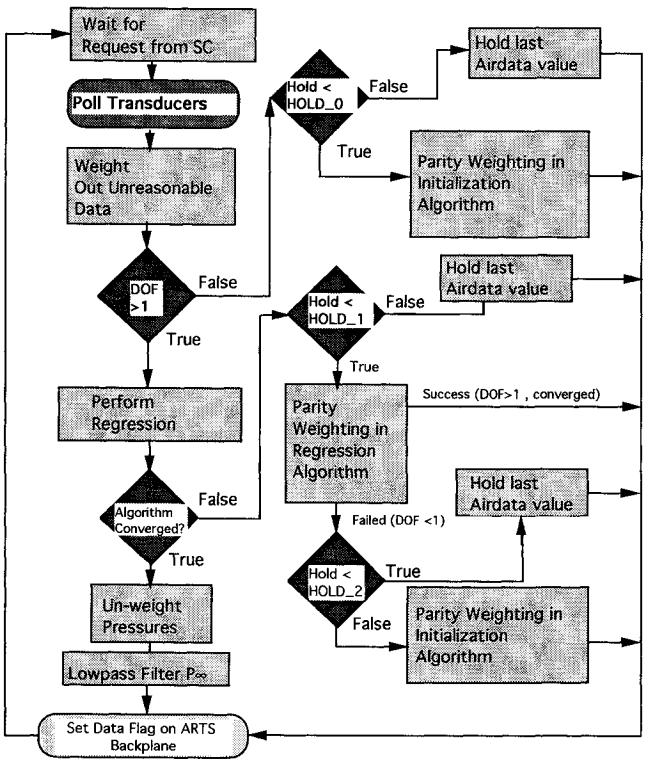

Fig. 4 Fault-tolerant data flow through RT-FADS computational cycle.

Sensor Failure Detection and Redundancy Management Using χ^2 Analysis

Because the RT-FADS algorithm is nonlinear and the solution is based on small perturbations to a current-state estimate, one true minimum (the physical solution) exists and multiple false minima may exist for each airdata grouping. If a large false perturbation is input to the algorithm, as may happen in the event of a sensor measurement failure, then the algorithm can be disturbed so far away from the true minimum that it

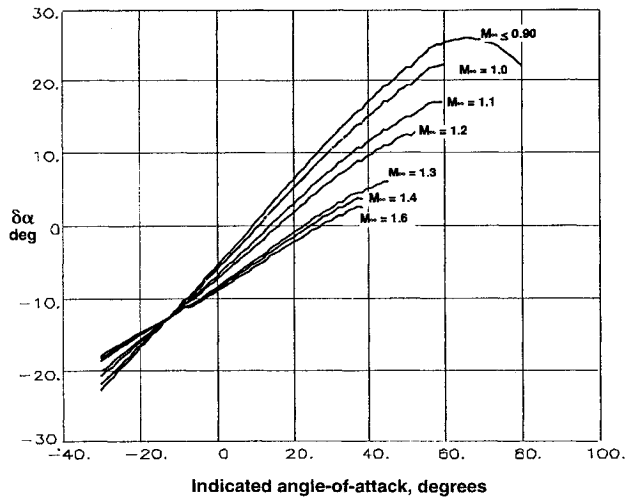

Fig. 5 RT-FADS upwash calibration curve.

converges to a false minimum or diverges altogether. In the former case, the algorithm computes erroneous results. In the latter case, the algorithm fails entirely. If a false minimum has been reached, the algorithm will not reliably return to the true minimum without reinitializing with a new starting condition.

Redundancy management techniques to identify and remove errors that could cause catastrophic instabilities have been developed using the methods of χ^2 analysis.[11,12] At the end of each iteration, the rms value of the data residuals represents a quantitative measure of the algorithm performance for that iteration. For a converged iteration, the pressure residuals are a subset of a much larger random population whose statistical properties are approximately zero mean and Gaussian distributed.

Because the RT-FADS residuals are Gaussian distributed, the sum squares of the RT-FADS residuals is a variable whose random distribution is closely approximated by the χ^2 distribution for five degrees of freedom.[11,12] Five-degrees-of-freedom results from the 11 independent transducer measurements and the fact that the model residuals are related by six param-

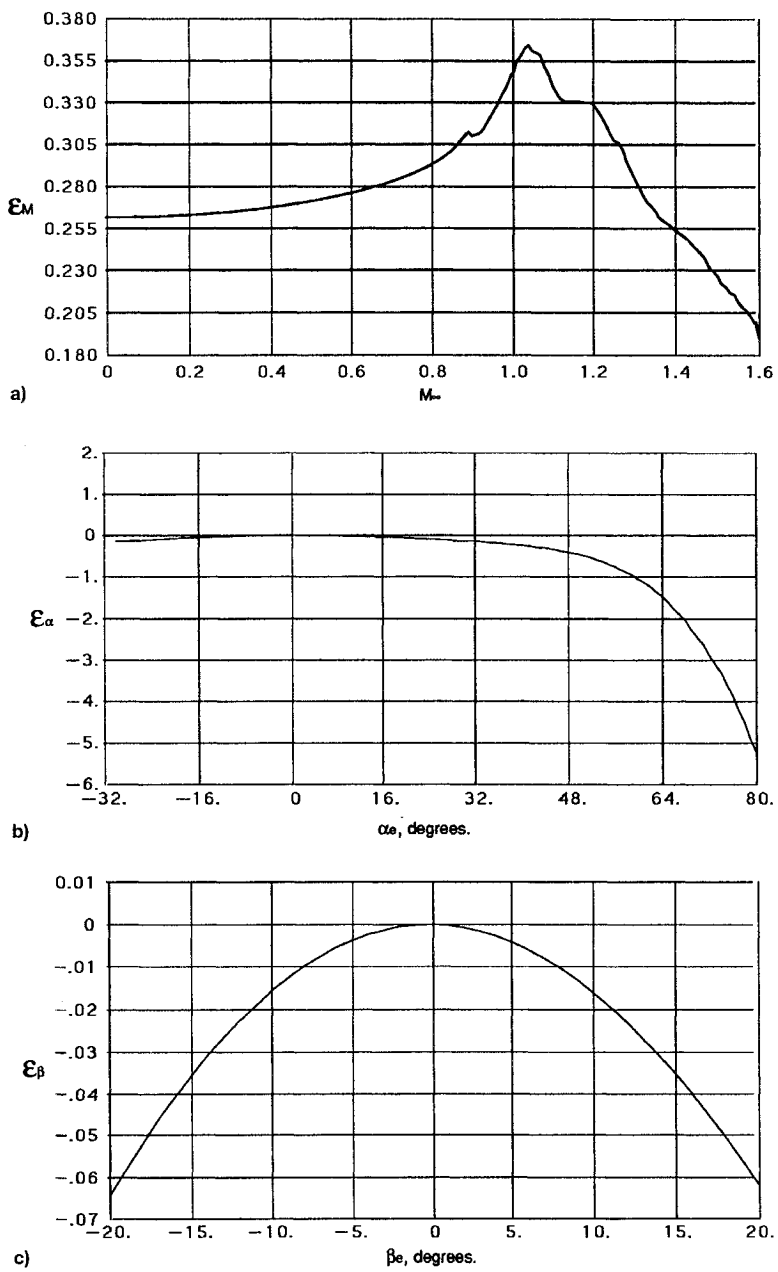

Fig. 6 Static source error calibration data. Variation with a) Mach number, b) angle of attack, and c) angle of sideslip.

eters estimated from the pressure data, the sample variance σ; the calibration parameter ε, and the four airdata parameters q_c, P_∞, α, and β.

Figure 3 shows a sample histogram of the sum-square residuals normalized by the sample population standard deviation (derived from flight data for converged computations) compared to the normalized χ^2 distribution for five degrees of freedom. This histogram validates the observation that the RT-FADS residuals are normally distributed, and gives a good quantitative means of evaluating the health of the RT-FADS algorithm.

Algorithm Convergence and Fault Detection Using the χ^2 Criterion

At the end of each iteration, comparison of this sample χ^2 variable with percentage points of χ^2 distribution gives a reliable statistical test of whether the algorithm converged. Because χ^2 is a relative probability indicator, a small value of χ^2 corresponding to high probability in the tables indicates that convergence is likely. A large χ^2 value corresponding to a low probability in the tables indicates that convergence is unlikely. Because of its quantitative indication of convergence probabilities, the χ^2 of the residuals is used as the convergence criterion for the algorithm. When the 90% residual confidence level is reached for a particular degree of freedom, then convergence is implicit and one additional iteration is performed.

Redundancy Management and Fault-Tolerant Processing Options

Figure 4 shows the redundancy management schemes employed in the RT-FADS algorithm. Because sensor failures are likely to occur, and because the RT-FADS processing algorithm uses all of the data simultaneously to compute the airdata values, an easy mechanism for removing failed sensors from the algorithm is implemented using the weighted least-squares technique. If a weight is set to zero, then the pressure reading corresponding to the weight has no influence on the estimation. A series of tests are performed for each transducer, and the result of each test is either one (passed) or zero (failed). If a transducer reading fails any of the fault processing tests, then

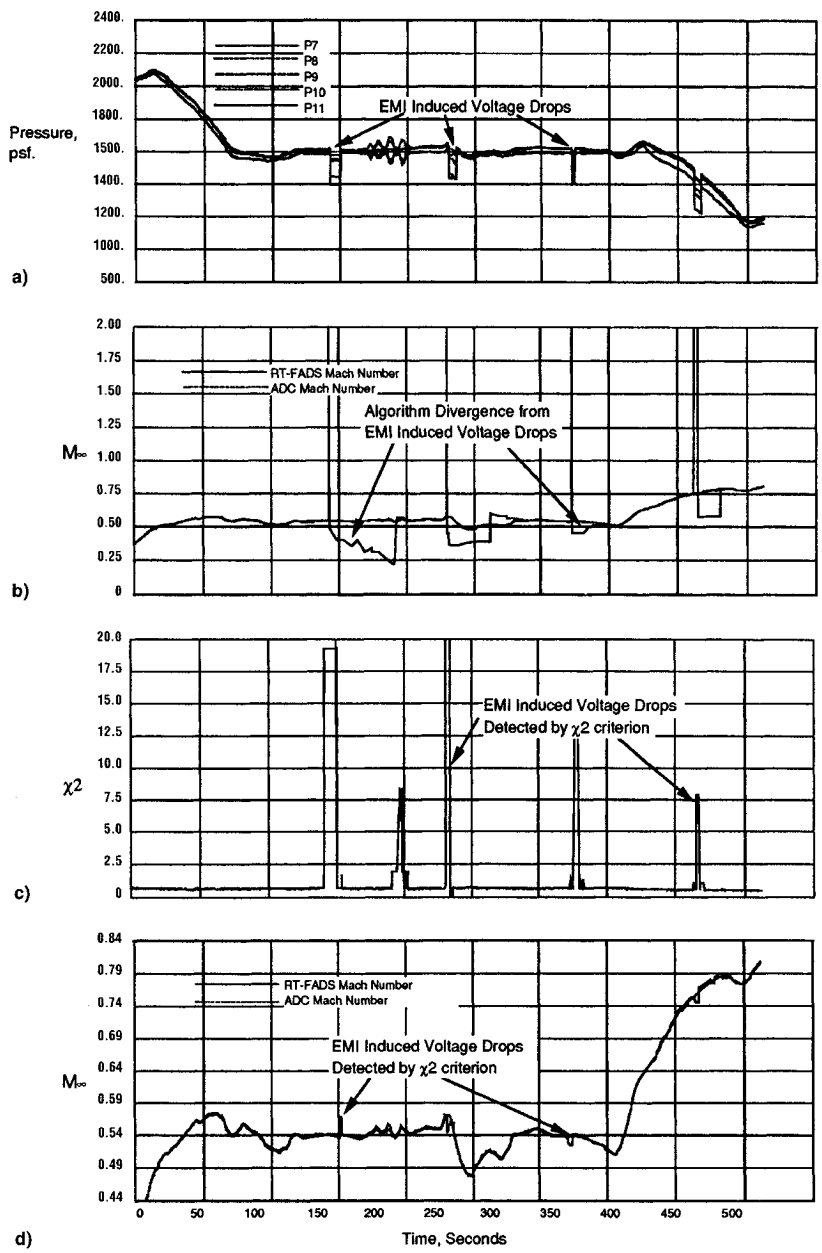

Fig. 7 RT-FADS failure detection and fault management example. a) Corrupted pressure data; b) Mach number, fault management inactive; c) χ^2 parameter, fault management active; and d) Mach number, fault management active.

it is weighted out of the algorithm and the degrees of freedom of the system are reduced by one. This weighting-out allows up to four transducers to be eliminated from the algorithm while still giving an identifiable system. The main benefit of the χ^2 methods is that they allow a nominal operation that devotes very little of its processing time to data checking. This benefit allows a fast nominal throughput. Only when the χ^2 test fails are the time-consuming data checking schemes used. This operation is a unique approach to fault-tolerant systems. More detailed theoretical development of the RT-FADS fault identification and redundancy management schemes have previously been presented.[8,9]

Flight-Test Results

This section presents flight results for the RT-FADS system. Generally, the integrated RT FADS/ARTS performed well, with more than 14 flight-hours completed. The system was flight demonstrated from takeoff to landing and throughout the entire nominal flight envelope of the F-18 airplane (to a maximum of 45-deg α, +25-deg β, and Mach 1.6).

Aerodynamic Calibration Procedures and Results

The calibration effort, based on the reference airdata methods,[7,13] consists primarily of two tasks: identifying the static source error represented by the parameter ε; and identifying the induced wash parameters $\delta\alpha$ and $\delta\beta$. The reference airdata values were generated by merging complementary information from multiple data sources, including the onboard inertial navigation system attitudes, rates, and accelerations; radar tracking velocity and position data[14]; and rawinsonde weather balloon sounding data.[15]

The weather balloon sounding data were verified in-flight by flying 360-deg level turn at constant airspeed before and after each maneuver. When the indicated airspeed from the RT-FADS system is averaged over the course of the 360-deg turn, the effects of the winds are eliminated. The difference between the averaged airspeed and the averaged radar-derived ground speed is the velocity error for that airspeed caused by the static source error.

The local wind direction and speed are evaluated by adding this static source velocity error to the indicated airspeed reading and then plotting the groundspeed and corrected airspeed as a function of time. Velocity data are converted to Mach number using temperature values obtained from the rawinsonde balloon soundings and radar-derived geometric altitude. Local ambient pressure values are evaluated using rawinsonde balloon soundings and radar-derived geometric altitude. The calibration parameters were estimated by substituting the reference airdata into the aerodynamic model and comparing the pressure predictions of the model to the pressures that were actually measured. Systematic trends in the calibration parameters were identified by plotting the estimated calibration parameters as a function of flight variables and visually inspecting the results. Once trends were identified, they were curve-fit

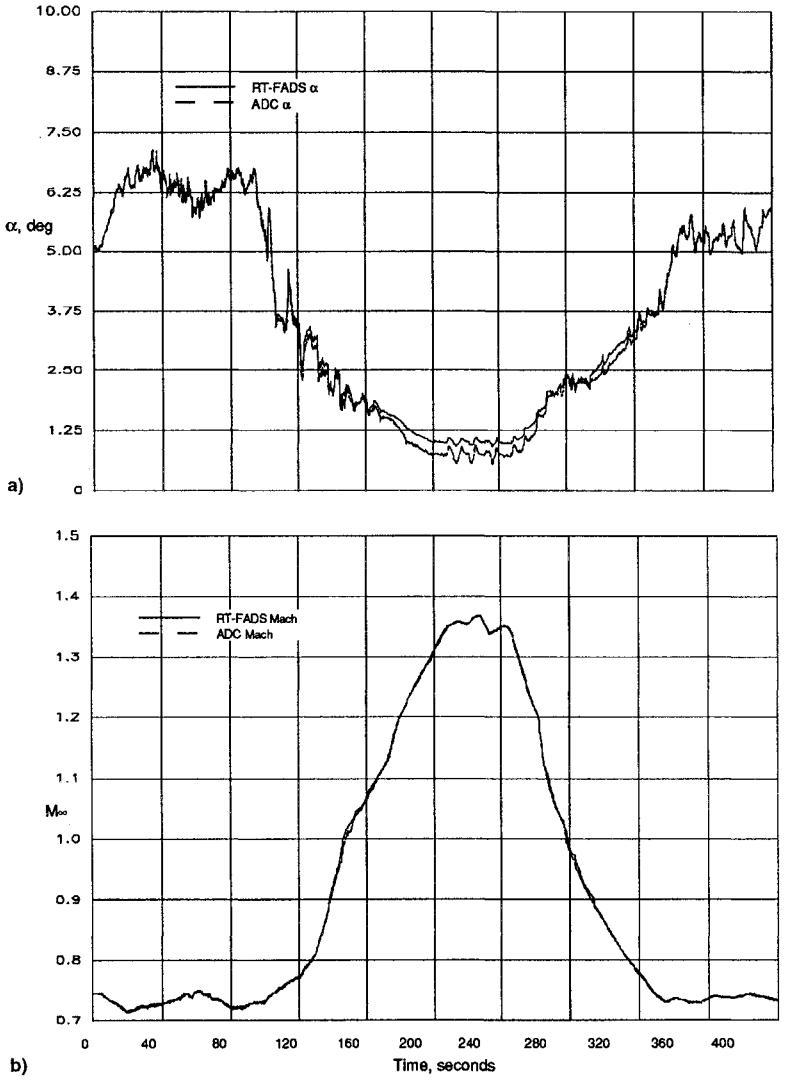

Fig. 8 Comparison of RT-FADS airdata to ADC for Mach 1.35 accel/deccel maneuver. a) Angle-of-attack and b) Mach number time histories comparisons.

and interpolated to generate a series of tabular break points that were hard-coded into the algorithm. As an example, Fig. 5 shows the resulting angle-of-attack calibration trend.

It was empirically determined that the static source error ε can be decomposed into three distinct components: 1) ε_M, which varies as a function of Mach number; 2) ε_α, which varies as a function of indicated angle of attack; and 3) ε_β, which varies as a function of indicated angle of sideslip. The total static source error is given by the summation of the three components. Figure 6a shows ε_M plotted on the ordinate axis and Mach number plotted on the abscissa. The curve rises steeply through the transonic flight regime, levels off, and then plunges rapidly at the high supersonic Mach numbers. Figure 6b shows ε_α plotted on the ordinate axis and angle of attack plotted on the abscissa. Figure 6c shows a similar plot for ε_β. The calibration parameters ε_α and ε_β are adjustments for the flow expansion over the nosecap at higher incidence angles. These parameters are typically negative in magnitude and offset the compression caused by the vehicle static source error.

Evaluation of Code Stability and Fault-Tolerance Methods

During the early phases of the RT-FADS flight tests, electromagnetic interference from the aircraft forward transmitter caused erroneously low readings in several of the RT-FADS pressure sensors whenever the pilot keyed the onboard communications and navigation radio. (The exact cause of this interference was never identified, but the problem was eventually remedied by disconnecting the forward transmitter.) Figure 7a shows the corrupted pressure data time histories. This partial failure of the RT-FADS measurement system offered the opportunity to evaluate the performance of the failure detection and fault management techniques. When the algorithm is run off-line with the fault processing options deactivated, the algorithm diverges momentarily, recovers to an erroneous solution, and finally recovers to the correct solution when the pressures return to their correct values. Figure 7b shows this failure where the RT-FADS Mach number estimate is compared to the ship system airdata computer (ADC) value. When the fault detection is activated, the χ^2 values jump rapidly at the instance of measurement failure. The failures are clearly identified and weighted out of the system. Figure 7c shows this χ^2 time history. Figure 7d shows that when the fault detection mechanisms are activated, only minor fluctuations in the Mach number solution are noted. The fault processing options worked and algorithm divergence was avoided. Other fault-tolerance examples have previously been published.[8,9]

Evaluation of the System Accuracy

Quantitative RT-FADS system performance was evaluated by comparing the calibrated outputs to the ship system ADC outputs. Although the ADC outputs are subject to the same types of measurement errors as the RT-FADS system, the ADC provides an accepted standard for comparison. Figure 7 shows sample time history comparisons for a Mach 1.35 supersonic acceleration/deceleration maneuver. Figure 8a shows time history comparisons for angle of attack. Figure 8b shows Mach number comparisons. The comparisons are excellent and are typical of results achieved.

Assuming that the ADC is the truth set and ignoring angles of attack greater than 25 deg (where the ADC data begin to lose accuracy because of design limitations), the statistical accuracy of the RT-FADS was evaluated as a function of Mach number. This evaluation was performed by taking root-mean-squared residuals between the ADC and RT-FADS measurements and graphing the residuals on a scatter plot as a function of Mach number for all of the data points gathered in the Phase I flight tests, a database of approximately 1,000,000 data points. Starting at Mach 0.20 and extending to Mach 1.60 at intervals of Mach 0.20, residual boundaries were drawn so that more than 99.9% of the residuals in each Mach number interval are included. These residual boundaries establish the 3-σ error bounds for the RT-FADS (relative to the ADC reference) parameters as a function of Mach number. Figure 9 shows the residual scatter plots and 3-σ error boundaries for Mach number and angle of attack.

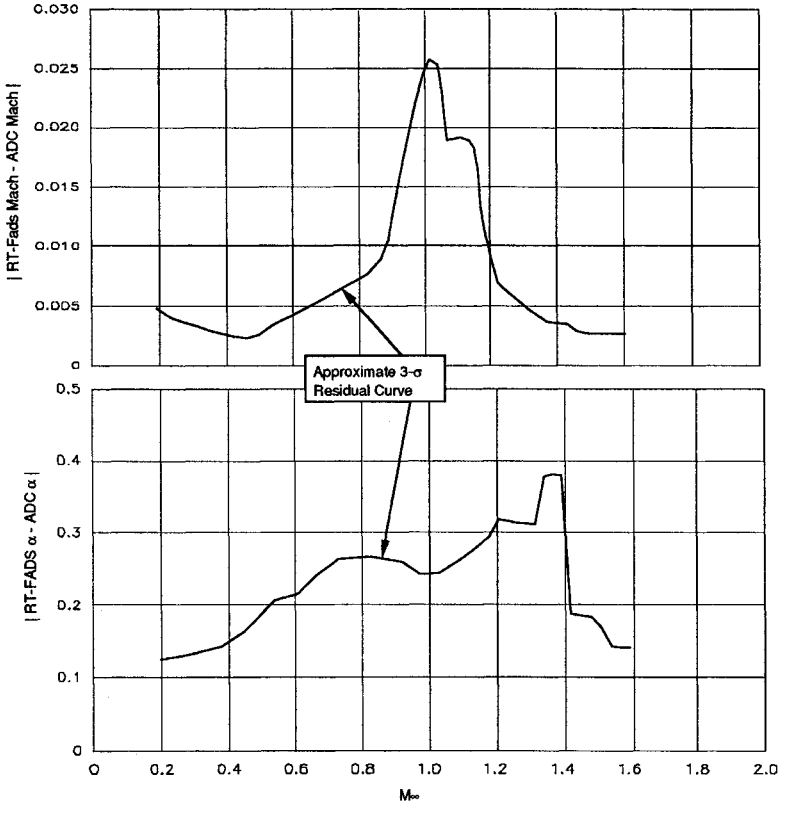

Fig. 9 3-σ residual scatter plot for RT-FADS system.

Table 1 3-σ residual values for RT-FADS parameters

M_∞	δM_∞	$\delta\alpha$, deg	δq_c, psf	δP_∞, psf
0.2	0.0045	0.125	3.50	1.00
0.4	0.0025	0.145	2.00	1.00
0.6	0.0035	0.210	2.00	3.25
0.8	0.0070	0.260	2.00	3.50
1.0	0.0240	0.240	12.00	16.50
1.2	0.0100	0.300	6.00	3.00
1.4	0.0030	0.380	5.00	2.00
1.6	0.0025	0.150	3.00	1.00

For angles of attack below 25 deg, the RT-FADS measurement accuracies were nearly independent of angles of attack and sideslip. Because no ship system value was available for angle of sideslip, no statistical evaluation was performed. Because the angles of attack and sideslip were calibrated using the same reference source, inferring that the accuracy levels are similar, is reasonable. Table 1 presents 3-σ residual values as functions of Mach number for various RT-FADS airdata parameters. Because the ship ADC is believed to have the same general error magnitudes as the RT-FADS, the error values shown in Table 1 are conservative. Even if the ADC reference is assumed to be a perfect source, the 3-σ residuals are well within accepted standards for airdata measurement accuracies.

Concluding Remarks

A novel nonintrusive airdata sensing system was developed and flight tested at the NASA Dryden Flight Research Center. This system uses a matrix of flush pressure orifices arranged on the aircraft forebody to determine the vehicle airdata, critical measurement parameters for flight control and research analyses. The system eliminates the need for external probes that are sensitive to vibration and alignment error, are easily damaged, may alter the flying qualities of the aircraft at high angles of attack, and are unacceptable for hypersonic or stealthy vehicles.

The RT-FADS system is significantly more advanced than earlier nonintrusive airdata concepts. The development of several innovations has resulted in a system that is capable of operating autonomously in real time. The system incorporates an overdetermined algorithm in which all surface pressure observations are used simultaneously to infer the airdata parameters using nonlinear regression. This innovation provides a system that is robust to noise in the measured pressure data and allows multiple sensor losses without significantly degrading the airdata computations. Generally, the integrated system performed well, with more than 14 operational flight hours completed. The system was flight demonstrated from takeoff to landing and over the entire nominal flight envelope of the F-18 airplane (to a maximum of 45-deg α, +25 β, and Mach 1.6).

Flight calibrations were performed using reference airdata values generated using data sources that include the ship system inertial navigation system, radar-tracking velocity and position data, and rawinsonde weather balloon sounding data. The calibration parameters were estimated by substituting the reference airdata into the aerodynamic model and using a nonlinear regression to identify the calibration parameters.

A χ^2 fault detection and redundancy management scheme was developed to stabilize the algorithm in the presence of measurement errors or sensor failures. These methods worked well for all of the Phase I test flights. This weighting-out allows up to four transducers to be eliminated from the algorithm while still giving an identifiable system. The main benefit of the χ^2 methods are they allow a nominal operation that devotes very little of its processing time to data checking. This benefit allows a fast nominal throughput. Only when the χ^2 test fails are the time-consuming data checking schemes used. This operation is a unique approach to fault-tolerant systems.

The RT-FADS equals the performance of the F-18 ADC at low angles of attack and extends the accurate measurement range to considerably high angles of attack. The statistical accuracy of the RT-FADS was evaluated by taking root-squared residuals between the ADC and RT-FADS measurements and graphing the residuals as a function of Mach number. Residual boundaries were drawn so that more than 99.9% of the residuals in each Mach number interval are included. These residual boundaries established a 3-σ error bound for the RT-FADS parameters as a function of Mach number. Because the ship ADC is believed to have the same general error magnitudes as the RT-FADS, the calculated 3-σ residual boundaries are conservative estimates of the true 3-σ error bounds. Even so, the resulting accuracy estimates are well within accepted standards for airdata measurement accuracies.

References

[1]Gracey, W., "Measurement of Aircraft Speed and Altitude," NASA RP-1046, May 1980.

[2]Cobleigh, B. R., and Del Frate, J., "Water Tunnel Flow Visualization Study of a 4.4% Scale X-31 Forebody," NASA TM-104276, Oct. 1994.

[3]Cary, J. P., and Keener, E. R., "Flight Evaluation of the X-15 Ball-Nose Flow-Direction Sensor as an Air-Data System," NASA TN-D-2923, April 1965.

[4]Siemers, P. M., III, Wolf, H., and Henry, M. W., "Shuttle Entry Air Data System (SEADS)-Flight Verification of an Advanced Air Data System Concept," AIAA Paper 88-2104, May 1988.

[5]Larson, T. J., Whitmore, S. A., Ehernberger, L. J., Johnson, J. B., and Siemers, P. M., III, "Qualitative Evaluation of a Flush Air Data System at Transonic Speeds and High Angles of Attack," NASA TP-2716, June 1987.

[6]Larson, T. J., Moes, T. R., and Siemers, P. M., III, "Wind-Tunnel Investigation of a Flush Airdata System at Mach Numbers from 0.7 to 1.4," NASA TM-101697, Jan. 1990.

[7]Whitmore, S. A., Moes, T. R., and Larson, T. J., "Preliminary Results from a Subsonic High Angle-of-Attack Flush Airdata Sensing (HI-FADS) System: Design, Calibration, and Flight Test Evaluation," NASA TM-101713, Jan. 1990.

[8]Whitmore, S. A., and Moes, T. R., "Failure Detection and Fault Management Techniques for a Pneumatic High-Angle-of-Attack Flush Airdata Sensing (HI-FADS) System," NASA TM-4335, Jan. 1992.

[9]Whitmore, S. A., Davis, R. J., and Fife, J. M., "In-Flight Demonstration of a Real-Time Flush Airdata Sensing System," NASA TM 104314, Aug. 1995.

[10]Putnam, B. W., *RS-232 Simplified: Everything You Need to Know About Connecting, Interfacing, and Troubleshooting Peripheral Devices*, Prentice-Hall, Englewood Cliffs, NJ, 1987.

[11]Franklin, G. F., and Powell, J. D., *Digital Control of Dynamic Systems*, Addison-Wesley, Reading, MA, 1980.

[12]Bendat, J. S., and Piersol, A. G., *Random Data: Analysis and Measurement Procedures*, Wiley-Interscience, New York, 1971.

[13]Whitmore, S. A., "Reconstruction of the Space Reentry Airdata Using a Linearized Kalman Filter," AIAA Paper 83-2097, Aug. 1983.

[14]Haering, E. A., Jr., and Whitmore, S. A., "FORTRAN Program for Analyzing Ground-Based Radar Data: Usage and Derivations, Version 6.2," NASA TP-3430, Sept. 1994.

[15]Ehernberger, L. J., Haering, E. A., Jr., Lockhart, M. G., and Teets, E. H., "Atmospheric Analysis for Airdata Calibration on Research Aircraft," AIAA Paper 92-0293, Jan. 1992.