WHY SIMULATORS DON'T FLY LIKE THE AIRPLANE - DATA

William G. Schweikhard, Senior Vice President
and
Daryl J. Schueler, Aerospace Engineer
Kohlman Systems Research
319 Perry Street
Lawrence, Kansas

ABSTRACT

The hardware and software technology for simulators and flight training devices have advanced enormously over the past ten years. We have been able to create very realistic visual scenes with high resolution, brightness and field of view; motion systems that provide the cues that give the feeling of actually flying in the airplane; high fidelity sounds that represent the operating environment inside the airplane; and computers that are capable of mathematically modeling the equations that represent the various components and systems being simulated. The quality of the data that is used to mechanize the flight dynamics and systems of the airplane being simulated is lagging. This paper focuses on the traditional approach for generating simulator design and verification data, and then describes a flight test approach for improving the quality of the data. Data developed by the traditional approach are compared with data developed by the flight test approach. Comparisons are made of simulated versus flight test results for operational maneuvers, one employing traditional data and the other employing flight test generated data. The need for high quality flight test data that exceeds those of current Development Test and Evaluation (DT&E) and Operational Test and Evaluation (OT&E) results is emphasized.

INTRODUCTION

The hardware and software technology for simulators and flight training devices have advanced enormously over the past ten years but we tend to take them for granted. It is well to remind ourselves of the things that have been accomplished, the basics of what we are about, and the opportunities for improvement that still remain.

As for the basics we might ask ourselves ... What is a Simulator?... And one might answer somewhat facetiously:

1. A hoax foisted against the pilots/operators to make them believe they are flying/operating the "real thing"; or more seriously we might answer

2. A unique combination of hardware, software and data which when combined in the proper way creates a realistic illusion of flying/operating the real thing.

The key words here are:
"The real thing" (and I don't mean Coke)
"Illusions"
"Hardware, software and data"

The goal of creating the illusion of the "real thing" leads us to the consideration of the quality of the hardware, software and data. These components are essential, especially the data, to create this illusion. It is the integrated effect of the hardware, software and data which provides the cues for a flight simulator's Cockpit, Instruments, Controls, Sound, Vibrations, Smells, etc. Each of these must progress together. If one is deficient, the quality and acceptability of the simulator suffers.

In the broadest sense we might say that the simulator is acceptable if the analogy holds which says: If it looks like a duck, waddles like a duck, quacks like a duck, swims like a duck, flies like a duck, then it must be a duck. The unfortunate part of this analogy is that it is qualitative and acceptability for today's simulators requires that quantitative criteria be met as well. In fact, if all of the criteria for simulator acceptability could be quantified, it would be possible to build a nearly perfect simulator that is limited only by the physical constraints of the cueing of the hardware. First we must examine the criteria for acceptability.

In the past a simulator was considered acceptable if it met the expectations of the trainer and trainees, whatever that may have been. Generally, it meant that the simulator operated pretty much like the airplane, at least to the point that there were no major distractions that detracted from learning IFR procedures. Over the years the hardware and software have improved to the point of creating some pretty realistic illusions. The expectations of the users have increased as well. "Close" is no longer good enough. If the majority of the training is to be done in the simulator, it must be as exact a replica as possible. Data plays a strong role in achieving this goal. Without good data the best of hardware is just a poor imitation.

SOURCES OF SIMULATOR DATA

Data for yesterday's simulators have come from a number of sources, but primarily it has come from the manufacturer's design data (aerodynamics, propulsion, controls, avionics systems, analytical predictions, weight, center of gravity and inertias, etc.). The problem was integrating the myriad of detailed factors from a multitude of sources into something that makes sense in terms of the flying qualities and the illusions that were desired. The result was not always the best. A classic example of how the detailed derivatives that make up the stability and control of an airplane can be distorted is illustrated in Figure 1 from Reference 1. Here an erroneous value of C_{l_β} was taken as correct at high angle of attack, and the other parameter C_{l_p} was distorted to compensate in such a way as to replicate the maneuver being matched. But the effect on other flying qualities not being matched is devastating and can result either in negative training or simulator compensation on the part of the pilot if he/she recognizes the model as being wrong.

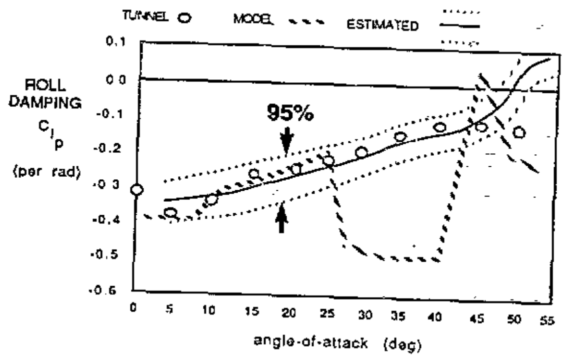

FIGURE 1a. Comparison of roll damping from a variety of sources

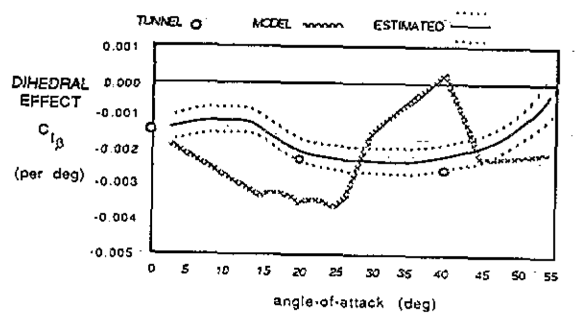

FIGURE 1b. Comparison of Dihedral Effect From a Variety of Sources

An example of the discrepancies that are common when using the data from multiple sources straight away is illustrated in Figure 2a. Here, the values of the various stability derivatives were taken from the manufacturers aerodynamic report for a business jet airplane and compared with the flight test data. The simulator model is being driven by the control surface displacements of the flight tests. When parameter identification techniques are used to determine the stability derivatives for the maneuver, the match is seen in Figure 2b to be identical to the flight test results. These results suggest that simulator data should be derived from flight test but as will be shown this approach is not the whole answer either.

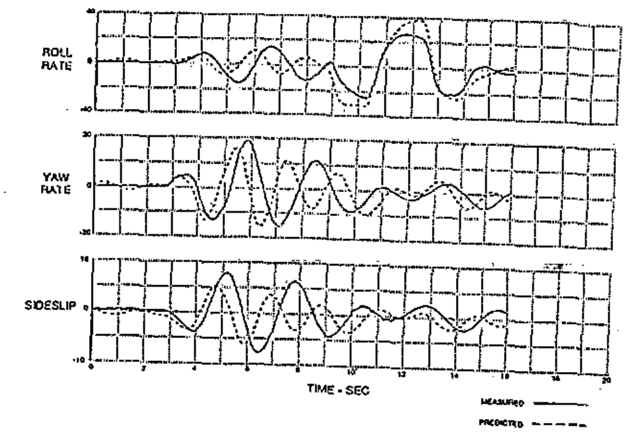

FIGURE 2a. Rudder-Aileron Doublet Predicted Model Response (Mach = .38, Alt. = 10,000 ft.)

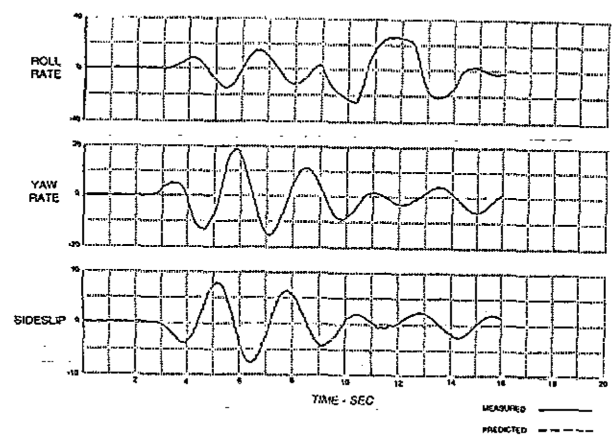

FIGURE 2b. Rudder-Aileron Doublet Final Results After 22 Iterations

FLIGHT TEST GENERATED SIMULATOR DATA

There is a sound rationale for developing simulator models from flight data. The manufacturer can tell you how the airplane is supposed to work. The airplane can tell you how it really does work providing you have the right tools and use them properly.

The benefits of creating a simulator model from flight test results are many compared to the traditional method of using manufacturer's data. First, an aircraft can be chosen that is representative of those operating in the fleet. This fact is especially important as aircraft age and the mechanical components show some slop and wear. Examples of having an airplane representative of the fleet will be given later.

Second, validation data can be collected from the same source as the design data for the simulator. Thus, there is a consistent set of data which, if of high quality and if analyzed properly, can produce very creditable results and a high quality simulator.

Third, the engineering process of parameter identification nulls out the effects of errors, ever present in the physical mass characteristics of the airplane *(weight, center of gravity and inertias)*, because the same values are used in the flight testing as are used in the simulator. It is a happy case of "right or wrong, be consistent."

It must be remembered though that the flight coefficients and derivatives may differ from the wind tunnel values, and that the flight test aerodynamic data and physical data are a matched set and cannot be interchanged with data from other sources (e.g. wind tunnel and computational predictions).

Fourth, aeroelastic effects are an integral part of flight derived coefficients and derivatives. Unlike the wind tunnel model, the airplane is not rigid so it flexes depending on the flight condition, mass distribution and the dynamics of the maneuver that is being analyzed. Thus, aeroelastic effects are buried in the coefficients and derivatives derived from flight tests. Aeroelastic corrections should not be applied when flight test data is incorporated in the simulator model.

Sometimes in using the flight test method we learn things that even the manufacturer didn't discover. This fact is borne out by comparisons of manufacturers design data with flight data.
(Reference 1)

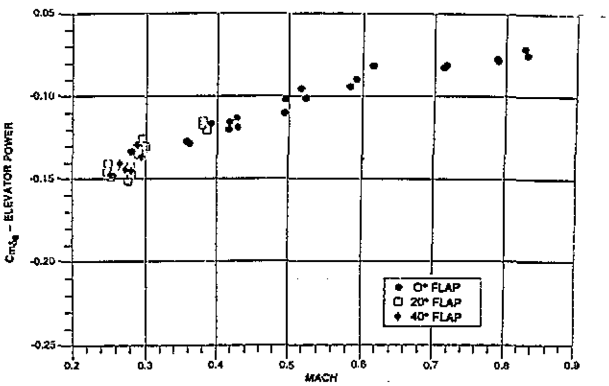

FIGURE 3. Results of Dynamic Maneuvers to Determine Elevator Power by using MMLE Methods

Figure 3 illustrates a set of data points for the elevator power derivative, $C_{m_{\delta_e}}$. Each point is the result of one time history match of a maneuver, like that of figure 2, at a given Mach number. The scatter is typical of that obtained with most flight derived derivatives. Note that this figure includes both flaps down and flaps retracted data. Figure 4 shows a comparison of these flight derived MMLE data, after fairing, with the manufacturers predicted values. The fifty percent error seen here is certainly not good enough.

Another example is shown in Figure 5, which is a comparison of aileron power, $C_{l_{\delta_a}}$, from flight test and the manufacturer's aerodynamic summary. It is clear *that significant differences in magnitude and shape* may exist between MMLE flight test results and predicted values.

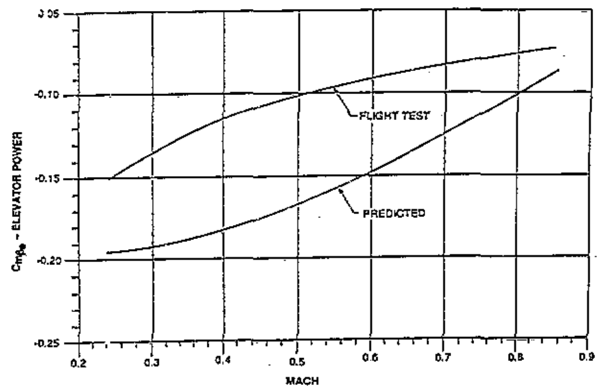

FIGURE 4. Comparison of Flight Test (MMLE) and Predicted Values of $C_{m_{\delta_e}}$ for a Business Jet

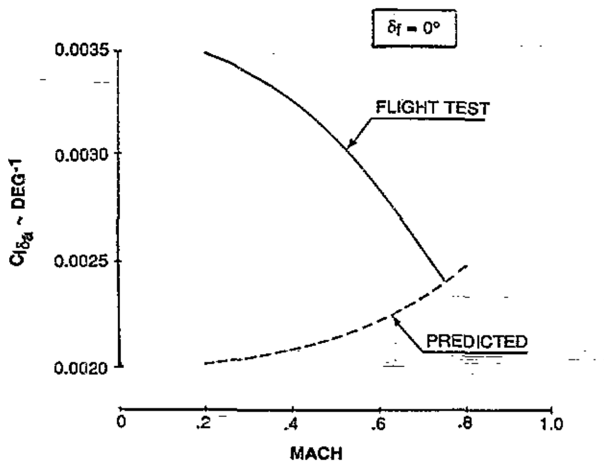

FIGURE 5. Comparison of Flight Test and Predicted Values of Aileron Power, $C_{l_{\delta_a}}$

An area in which flight test parameter ID methods have been particularly useful is in the determination of the effect of quasi-steady aeroelasticity on stability and control derivatives. Figure 6 shows the result of extracting elevator power, $C_{m_{\delta_e}}$, of a popular business jet over a large range of Mach number and dynamic pressure. Elevator deflection was measured near the actuator rod at the root of the elevator. Elevator *twisting at high dynamic pressure causes a significant* attenuation in elevator power, resulting in important departures from the data produced using relatively rigid wind tunnel models.

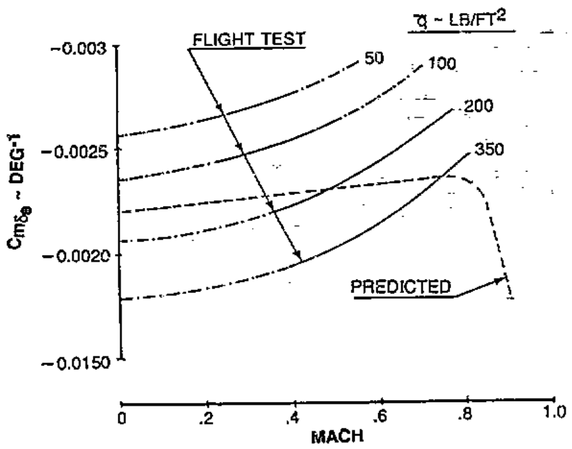

FIGURE 6. Comparison of Flight Test and Predicted Elevator Power with Aeroelastic Effects Included

Another similar example is shown in Figure 7, which shows a data comparison of the dihedral effect, C_{l_β}, of an airplane with winglets. Although the magnitude and shape were well predicted for a rigid airplane, high dynamic pressure causes aeroelastic bending of the winglets, which attenuates the dihedral effect. Such changes must be properly accounted for in simulators.

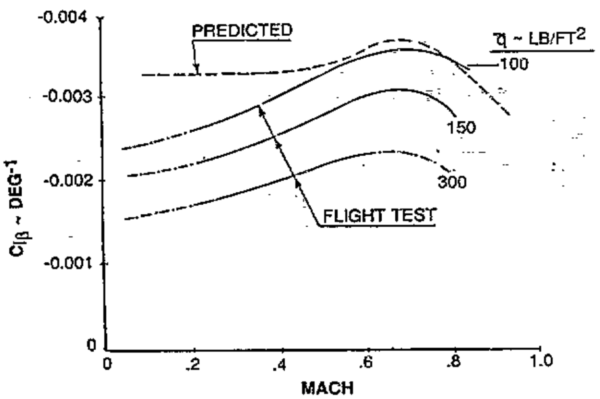

FIGURE 7. Effect of Aeroelasticity on Dihedral Effect, C_{l_β}, of an Airplane with Winglets

Some derivatives like yaw damping due to roll rate, C_{n_p} and rolling moment due to yaw rate, C_{l_r} have been difficult, if not impossible, to extract from flight tests before parameter identification was developed. Now it is relatively straightforward to determine these cross-derivatives, with reasonably good accuracy as shown in Figures 8a, b.

Finally, Figure 9 shows the resolution capability of parameter identification methods. Thus, even when predictive methods are relatively good, effects of secondary parameters, such as angle of attack and flap deflection, can be accurately determined from flight test data.

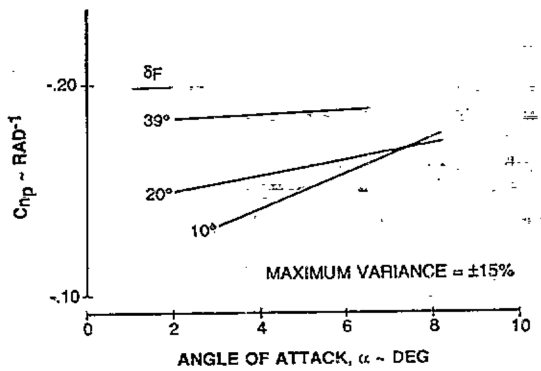

FIGURE 8a. Lateral-Directional Cross Derivatives C_{l_r} and C_{n_p} Determined from MMLE Maneuvers

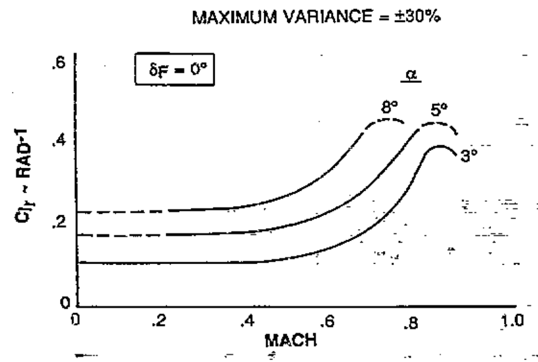

FIGURE 8b. Rolling Moment Due to Yaw Rate Determined from MMLE Maneuvers

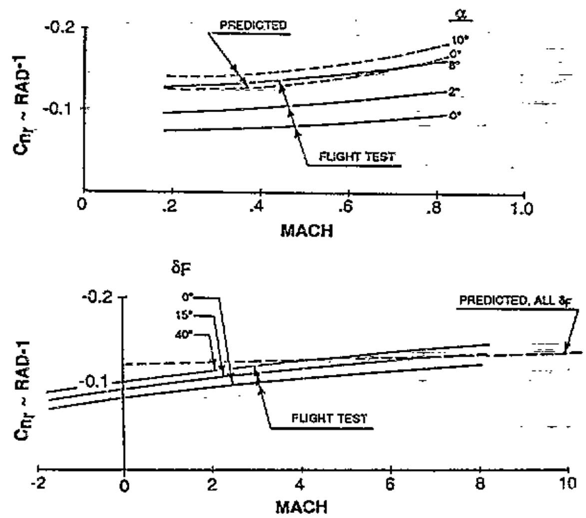

FIGURE 9. Comparison of Flight Test and Predicted Values of Yaw Damping Derivative, C_{n_r}

UNSIMULATED EFFECTS

There is strong evidence of unsimulated effects and derivatives in flight test data. Simulator models include all of the primary derivatives and some of the more commonly encountered secondary derivatives but others are frequently missed especially if the design data model has been derived from maneuver sets that were not designed to reveal them (e.g. a systematically intermeshed combination of Mach number, dynamic pressure and angle of attack).

The implication of the data scatter of Figure 3 is more than simple scatter that is normally found in experimental data. For, if the derivative set for any one test maneuver of Figure 3 is placed in the flight simulator, the match is nearly identical as shown in Figure 10. This figure shows three sets of time history results: the flight test, the match of the flight test using the MMLE model that generated the derivatives, and the flight simulator results using the MMLE derivative set obtained from the flight test data. Both the MMLE and flight simulator models were driven by the flight test control surface deflections for the maneuver. The minor discrepancies between the flight and the model results could be due to small experimental measurement errors in the calibration for the flight control surface deflections of the airplane or to unrecognized and therefore unsimulated terms in the models that produce subtle differences in the results. Note that the match is less exact late in the maneuver during the free response and recovery portion when

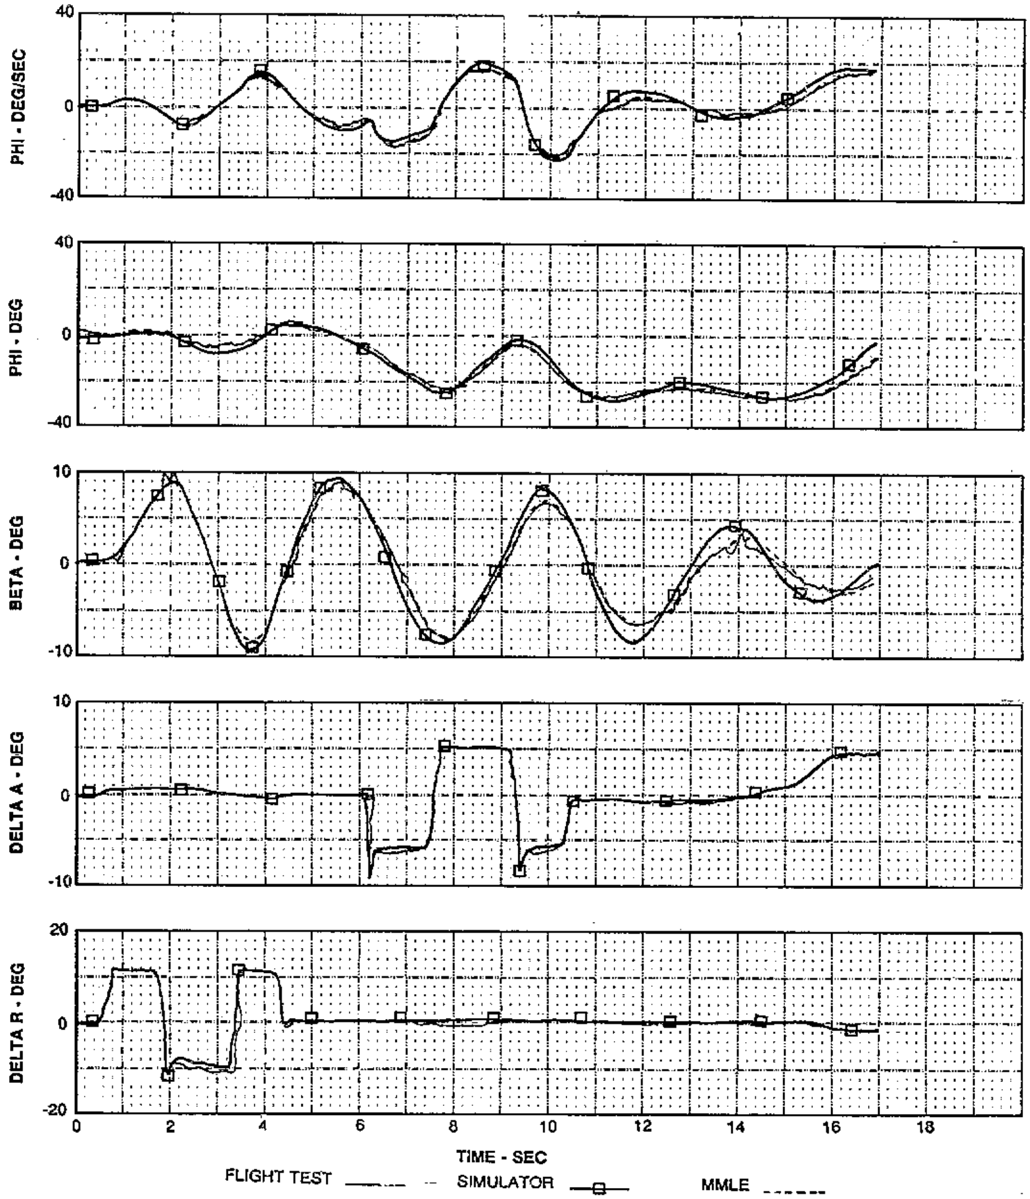

FIGURE 10. Comparison of MMLE Model, Simulator Model and Flight Test using MMLE Derivative Without Tuning the Simulator Model

the magnitude of the disturbances is small. Frequently these discrepancies are attributed to minor errors in the derivative set and the drifts they cause as the integration duration becomes large. However, they could also be caused by the relative magnitudes of the simulated variables compared to the unknown and therefore unsimulated parameters that should have been included in the models.

Thus, the scatter of the data points of Figure 3 may not be scatter at all, but rather, minor differences in modeling where unsimulated and perhaps non–linear parameters are concerned. It should be noted that the data points of Figure 3 that are at the same Mach numbers were done within one minute of each other so the center of gravity and moments of inertia were essentially identical. Yet the derivatives that were derived were slightly different.

These differences as indicated by the scatter become very large for secondary derivatives and for primary derivatives when that mode of motion has not been excited. For example, one would not expect to extract good lateral directional derivatives from a longitudinal maneuver where the lateral–directional motions simply have not been excited.

An example of the effects of variables that are not normally present, and thus are not modeled for most simulators, are shown in Figure 11 for a C–141 airplane. Its configuration, its age and perhaps its aeroelasticity create some anomalous behaviors. Here, a simple rudder doublet has excited both the lateral–directional and the longitudinal modes. The pitching oscillation is twice the frequency of the lateral–directional. A derivative for pitching moment due to sideslip (C_{m_β}) had to be added to the model to properly reproduce the observed motion. Also present in this maneuver is a non–linear floating aileron trim tab oscillation condition and a non–linear asymmetric out of phase thrust variation that contributes to the resulting motion of the airplane. The trim tab slop of almost one half degree is probably a minor effect but the sawtooth and out of phase thrust variation due to sideslip indicated by the EPR1 (left outboard engine) and the EPR3 (right inboard engine) variations are very large, 800 pounds for engine one and 600 pounds for engine three. The total moment produced for all engines is of the order of 25,000 foot pounds which is not of secondary magnitude.

The interplay of these and perhaps other secondary effects, cause the lateral directional oscillatory modes to be unstable at higher altitudes and quite stable at low altitude. A detailed analysis and discussion of these characteristics are beyond the scope of this paper and will have to wait for another paper on another day. But the fact remains that there are other derivatives and effects that are not represented in today's simulator models.

It is the author's belief that the primary and secondary contributors to the motions of an airplane can best be identified through the flight test approach along with some very perceptive analysis of the data by some very skilled interpreters. Some very high quality instrumentation and advanced flight test procedures are required to reveal what might otherwise be missed by other traditional flight test technicians.

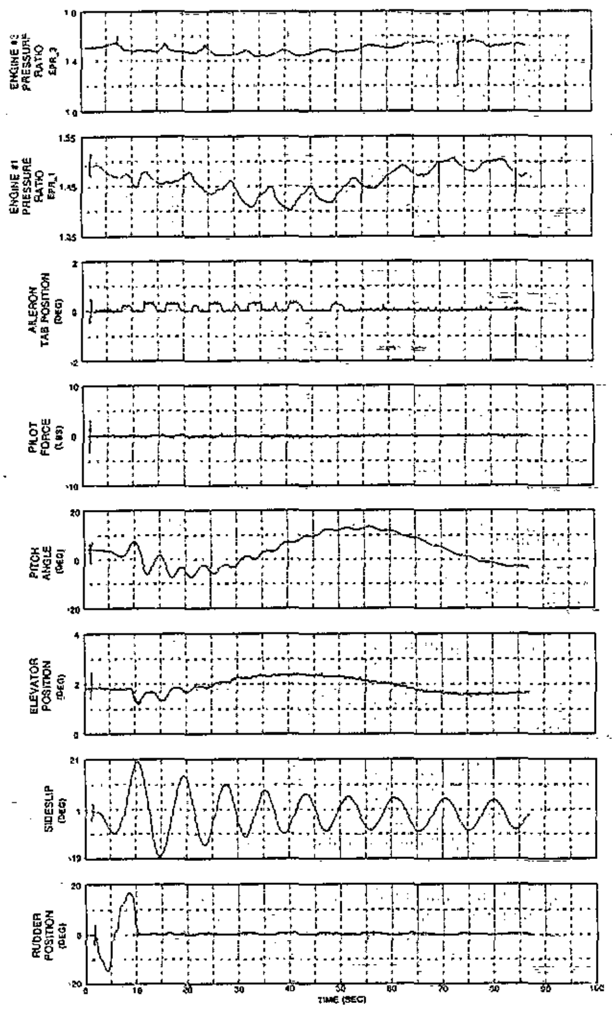

FIGURE 11. Lateral/Directional Controls Free Oscillation of a C-141 Airplane, 29,000 ft.

FLIGHT TEST DATA

Traditional flight test data that is done for development, test and evaluation (DT&E) in the military sector and for FAA certification in the civil sector is none too good for the validation of flight training simulators. Large development and certification programs frequently are done using several test articles each having different purposes (e.g. performance, stability and control, propulsion, avionics, loads and structural dynamics). Each has different types and numbers of sensors to accommodate the specific test areas and objectives assigned to it. Rigorous consistency of type and quality of measurements between the various test articles is not a high priority.

In some cases the various test articles from which the simulator validation data comes are not of the same lineage. For example data may come from a prototype, a preproduction airplane and a production airplane which are similar but different in ways that do not affect the development process but do seriously impact the simulator acceptability. Data from the power plant and flight control systems are often significantly different between prototype and production airplanes. Yet the data set provided to the simulator manufacturer frequently consists of an unholy mix of data from a number of different aircraft having similar but different characteristics of the various components and systems. Even when all of the different test articles are of the same lineage (i.e. all production airplanes) there are differences in data acquisition systems and the quality of a particularly key measurement from one airplane to the other. The differences are far more important for simulator validation than for developmental testing. Since developmental testing requirements drive the quality of the data gathered during the early stages of testing and even later, much of the simulator data gathered today is inadequate either because of the way the test was flown or the quality of the data taken.

A classic example of the pitfalls of using "best available" DT&E data taken from similar but different test articles is described on reference 3 for the AH–1W helicopter. Data were available from a prototype (DT–11F) and a production (DT–111) aircraft. The prototype data were to be used as the primary set and holes (tests not available) were to be filled with data from the production aircraft. Almost every possible ill was present in these data: noise, data scatter, data shifts, uncertain parameter scaling, missing parameters, uncertain fight conditions, data trend differences, misunderstood pilot techniques, wide test tolerances, etc. Many of these ills are seen in Figure 12a which shows large scatter and poor trend information coming from the two test articles DT–11F and DT–111. The simulator, though not yet finely tuned showed a trend similar to the DT–11F results. A consensus was reached with the test and simulator procuring organizations that the curve fit along with the wide tolerance shown in figure 12b should be the standard for simulator acceptance. For a full discussion of the sort of problems that are encountered when forced to use data that was not specifically intended for validating simulators, the reader should review reference 3 for further details and horror stories. The term "Best Available Data" is frequently used in the simulator validation business. The acronym for this term aptly describes this type of data, BAD. Using best available data is like finding something in the garbage dump, it usually stinks.

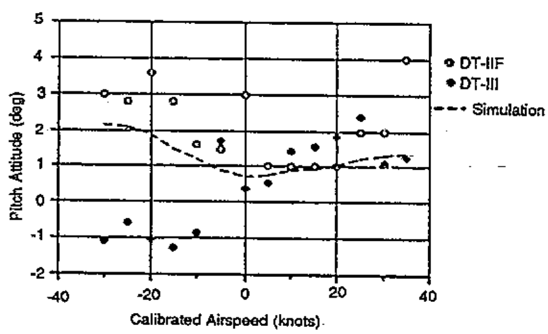

FIGURE 12a. Pitch Attitude Trend Disparity

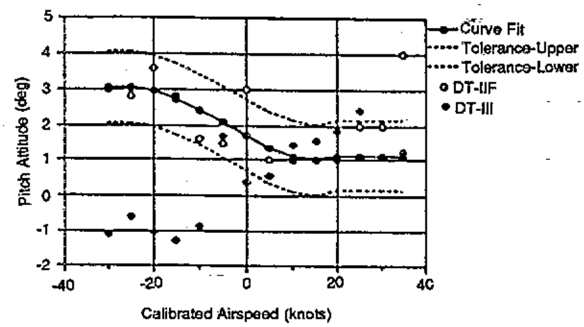

FIGURE 12b. Pitch Attitude Final Test Criteria Definition

Simulator data requirements have redefined the term "Quality Data" with standards that exceed the best available flight research data. Quality data has no noise (< .1%), high precision (accuracy < .05%, resolution < .025%) and has inherent stability (repeatability of .05% and no zero shifts). Great care is required to see that slight zero shifts, calibration changes, flight test techniques and analysis methods are "the best possible", because "the best possible" are sometimes still not good enough. Special flight tests should be run to obtain design and validation data for high quality flight simulators that fly like the airplane.

SIMULATOR VALIDATION PROCESS

As stated at the outset of this paper, the acceptability of simulators a decade or two ago was by and large qualitative, especially in the flying qualities area. The objectives were training in instrument flying procedures and there were no credits for visual takeoff and landings. These procedures were done in the airplane.

The demand for full flight simulation and training has changed this situation. Acceptability has taken on the meaning of truly representing the airplane. The validation process has changed from one of "being close" to one of being as exact a replica of the real thing as possible. Quantitative criteria have replaced much of what was once a qualitative process. These quantitative criteria can be specified in terms of the frequency, damping and time constants of dynamic maneuvers or the more sophisticated frequency domain parameters can be specified. The leadership in this area has come from the civil sector (FAA, IATA, etc.) which embraces the concept of using flight test measured control inputs to drive the simulator controls, then measuring the outputs of the simulator and comparing them with the flight test responses. The differences between the two are evaluated against established tolerances within which they must fall. This provides an end–to–end check of a variety of different types of maneuvers over the flight envelope of the airplane. This process then serves as the basis for requalifying the simulator at regular intervals (once a year for FAA AC 120–40B). This process is described in reference 4.

The validity of this approach assumes that good flight test data and well flown maneuvers are available.

Matching of time history maneuvers is not a random "tweaking" process. Trading of the values of one derivative in order to compensate and match the time histories is what caused the disparities of stability derivatives illustrated in figure 1. Changing C_{l_p} to accommodate an erroneous C_{l_β} created a reasonable match, but two wrongs don't make a right. The simulation had to be lacking in another area where these derivatives are important (e.g. roll characteristics and steady sideslips). Erroneous adjustments then cascade into other parameters that must then be adjusted. Those who successfully balance the entire set of coefficients and derivatives are real artists

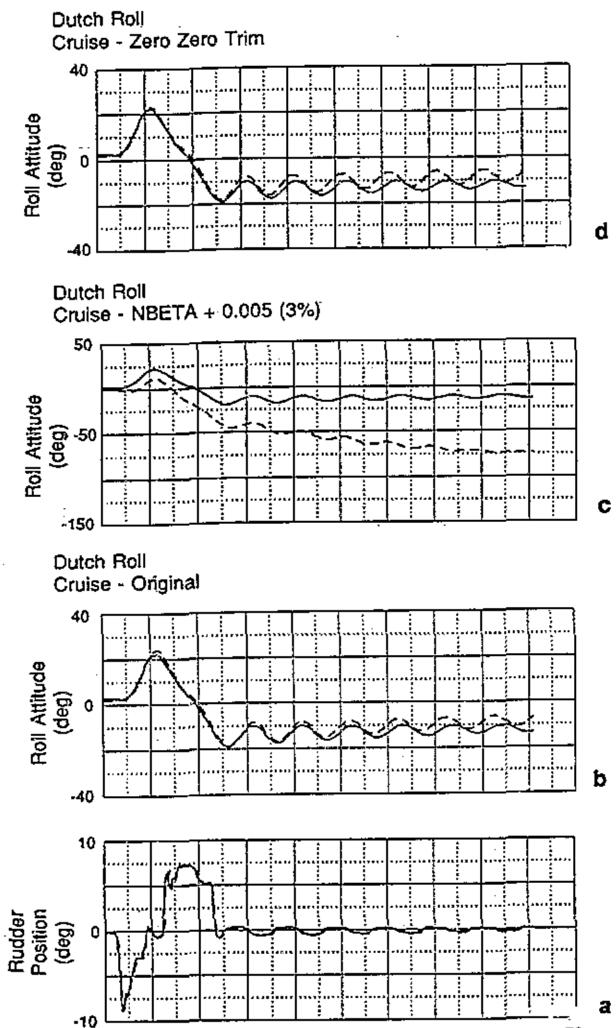

FIGURE 13. Effects of Rate Mistrims and Derivative Error

An example of the sensitivities to subtle and minor changes in trim and stability derivatives are illustrated in figure 13 which show the response of the aircraft bank angle of a C–141 airplane to a rudder doublet (Figure 13a) that excites the dutch roll mode of the airplane. With the simulator trimmed so as to account for the small angular rates due to mistrims of the test data (less than .1 degrees per second) the first cut using the flight test derived derivative set for this flight condition gives a reasonably good match as seen in 13b. But if the angular rates are taken to be zero at the trim condition, the simulated bank angle is seen to diverge as seen in Figure 13c, and exceeds fifty (50) degrees error by the end of the maneuver. It is very important to begin matching maneuvers from proper trim conditions.

The effect of simulator derivative errors in C_{l_β} of 3 percent, illustrated in Figure 13d, is less dramatic than errors in trim. However, there are noticeable effects on the frequency and damping of the oscillation for this very small error in the value of just one derivative. These kinds of errors can be expected in any of the coefficients and derivatives that affect a maneuver. The result could be very dramatic if the errors interacted adversely in the simulator model.

SIMULATOR ACCEPTABILITY

In spite of all of the good hardware, software and data, simulators still will not fly like the airplane. This fact is true because certain flight conditions and types of maneuvers can never be properly represented by the motion or visual systems. The reason is that we don't understand all we know about the system we are analyzing.

Parameter identification is not the whole answer to simulator model development. It does a wonderful job of developing the derivatives for large disturbance maneuvers but those modes, and associated derivatives that are not or cannot be excited are not well identified.

Pilot acceptability of a simulation is greatly influenced by the very small vibrations and responses about the trim condition, either level flight or while maneuvering. These responses are impossible for the manufacturer to quantify in the design and development stage and they are very difficult to evaluate in flight but they are important to pilot opinion. First, the forces are very small and the control movements are almost imperceptible. Being small they are in a region that is most likely non linear so that the traditional equations of motion do not apply. Analysis of this particular area deserves closer examination both in flight and in the simulator.

Even though there has been much progress, we still need to focus on the common complaints of experienced and knowledgeable pilots that still remain even for the best and most sophisticated of simulators. The more common ones are outlined in Table I but surely there are others. The ones enumerated are clearly related to the data issues that are the subject of this paper.

TABLE I

COMMON COMPLAINTS FOR THE BEST OF SIMS

General Comments:

- Simulators Can Pass All Objective Match Tests and Still Not Fully Represent the Aircraft – Small Disturbance Effects Are the Key Reasons.
- Simulators lack the crisp response and feel of the real aircraft.
- Large disturbance motions are generally satisfactory but small inputs and response are where the pilot is most sensitive and most critical

Specific Comments:

- Response of Visual System and Instruments to Control Input lags the input (Latency)
- Subtle vibrations and sounds are not represented properly
- Ground Handling Motion Peculiarities are not correct
- Visual Peripheral Cues are limited
- Low Daylight brightness especially internal to cockpit
- Highly Maneuverable aircraft can't simulate the steady state g–loads
- Ground Effect of the simulator is usually lower where lift and, to a lesser extent, pitching Moment is concerned.

"Full flight simulators such as those that meet FAA levels C and D are very good at creating an illusion of flying the 'real thing' and can make the pilot sweat. But the illusion is shut down and it becomes just a trainer when minor distractions occur" (e.g. an atypical response, the flickering of a light, the extraneous noise of a fire bell in the simulator next door, a comment or even worse a freeze of the simulator action by the instructor, etc.)

Atypical responses are an engineering problem but many of the other distractions are a training procedures problem. Perhaps it would be better to temporarily save a particular training exercise and return to it later at the critical point to critique or re–initiate it.

CONCLUSIONS

No matter how good the hardware and the software, the final simulator will still be a poor simulator if the type and quality of the data is not good.

The traditional approach to generating simulator design data from manufacturers design data is lacking in that it leaves many holes in the data set where parameters must then be adjusted and filled by highly skilled "artists" to match the specific flight condition and also to make smooth transitions from one flight condition to another in a seamless manner. This procedure is time consuming and introduces unreal characteristics into the data model.

The flight test approach provides a more consistent and coherent data model that represents more of the variables as we understand them today. Flight test generated models tend to be insensitive to errors in weight, c.g.'s and inertias because the same model is used to quantify the model in flight as is used in the simulator on the ground. The model derivatives obtained in flight contain all of the flexibility characteristic of the airplane without having to overtly know what modes and stiffnesses of the airplane really exist. Furthermore, if the validation data is obtained on the same airplane as the design model data, the two sets of data are more coherent because the variables that might otherwise be introduced by discrepancies between two different data acquisition systems are not present.

While flight test is the preferred way to obtain both design and validation data it is not the whole answer, especially for small disturbances where nonlinear effects of control friction and small disturbance aerodynamics are present. A considerable amount of study is needed to investigate, quantify and model these effects. Until more knowledge is gained in this area, final adjustment to the simulation model will have to be based on qualitative comments of pilots and the skill of simulator engineering/artists.

Finally, experienced flight test people grossly underestimate the quality and care that must go into producing flight test data for simulator validation let alone the gathering of design data. In large flight test organizations it can be very difficult, if not impossible, to get everyone sensitized and religiously committed to attending to every detail in order to squeeze every last bit of performance out of the data system.

The demand for higher quality training devices and greater training credits will continue to push the simulation industry for better hardware and software but we must remember that the simulator can be no better than the data and the data model mechanization.

In the final analysis, simulator acceptability is totally determined by the pilot. But objective standards have gone a long way toward providing a product that is very close to being acceptable. Increased knowledge and understanding of the secondary aerodynamic derivatives and nonlinear characteristics of the airplane and control system could bring the objective and subjective criteria for acceptability in much closer agreement.

REFERENCES:

1. Hess, Robert A; Hildreth, Bruce L.; Numerical Identification and Estimation of an Efficient Method for Improving Simulation Fidelity I/ITSC Proceedings Nov. 1990 pg 123.

2. Kohlman, David L.; A Comparison of Flight Test and Predictive Results from a Series of Simulator Data Generation Programs, AIAA 88-2131

3. Humphrey, Tom; Winters, Mark; Definition and Validation of the Flying Qualities and Performance Test Criteria for the Modern Operational Flight Trainer; I/ITSC Proceeding, Nov. 1990, pg 130.

4. Tafts, William D.; Kraft, James N.; *Boeing Simulator Approval Process*, AIAA-85-1736.

5. Jacobs, Robert S.; Roscoe, Stanley N.; Simulator Cockpit Motion and the Transfer of Flight Training, ISBN 0-8138-1925-3