DOT/FAA/AR-98/3

Office of Aviation Research
Washington, D.C. 20591

FAA/NASA Joint University Program for Air Transportation Research 1994-1995

May 1998

Final Report

This document is available to the U.S. public through the National Technical Information Service (NTIS), Springfield, Virginia 22161.

U.S. Department of Transportation
Federal Aviation Administration

National Aeronautics and Space Administration

NOTICE

This document is disseminated under the sponsorship of the U.S. Department of Transportation in the interest of information exchange. The United States Government assumes no liability for the contents or use thereof. The United States Government does not endorse products or manufacturers. Trade or manufacturer's names appear herein solely because they are considered essential to the objective of this report.

Technical Report Documentation Page

1. Report No. DOT/FAA/AR-98/3	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle FAA/NASA Joint University Program for Air Transportation Research 1994-1995		5. Report Date January 1998
		6. Performing Organization Code AAR-201
7. Author(s) James H. Remer, Compiler		8. Performing Organization Report No.
9. Performing Organization Name and Address Department of Transportation, Federal Aviation Administration, William J. Hughes Technical Center, Atlantic City International Airport, NJ 08405		10. Work Unit No. (TRAIS)
		11. Contract or Grant No. NGL-22-009-640 NGR-36-009-017;NGL-31-001-252
12. Sponsoring Agency Name and Address Department of Transportation, Federal Aviation Administration, Office of Aviation Research, Washington, DC	National Aeronautics & Space Administration, Washington, DC 20546-0001	13. Type of Report and Period Covered Technical Report Fiscal Year 1995
		14. Sponsoring Agency Code AAR-201
15. Supplementary Notes		

16. Abstract

The Joint University Program for Air Transportation Research (JUP) is a coordinated set of three grants co-sponsored by the Federal Aviation Administration (FAA) and the National Aeronautics and Space Administration (NASA). Through the year covered by this report (FY1995), NASA technical and financial sponsorship was provided by the NASA Langley Research Center.

The JUP seeks to implement new methods which foster interdisciplinary research and education which would provide researchers and innovative solutions to the large scale system problems which emerge in the continued development of the National Airspace System. Under JUP, three institutions: the Massachusetts Institute of Technology, Princeton, and Ohio Universities receive research grants and collaborate with FAA and NASA in defining and performing civil aeronautics research in a multitude of areas. Some of these disciplines are artificial intelligence, control theory, atmospheric hazards, navigation, avionics, human factors, flight dynamics, air traffic management and electronic communications.

The unique feature of the JUP are the quarterly reviews held on a rotating basis at the universities and FAA and NASA sites. At these reviews FAA and NASA technologists are afforded the opportunity to critique ongoing research, as well as to suggest new projects and avenues of research needed by their respective agencies.

17. Key Words Air Traffic Control; Air Traffic Management; Avionics; Aircraft Guidance and Control; Human Factors	18. Distribution Statement This document is available to the U.S. Public through the: National Technical Information Service Springfield, Virginia 22161

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	119	

CONTENTS

OHIO UNIVERSITY

EXECUTIVE SUMMARY

The FAA/NASA Joint University Program for Air Transportation Research(JUP) is a research partnership of three universities, which conducts scientific and engineering research on problems of a long term nature related to the ultimate improvement and development of the National Airspace System. JUP research covers a broad scope of technical disciplines which contribute to civil aviation , including , but not limited to air traffic control theory, human factors, satellite navigation and communications, aircraft flight dynamics, avionics and meteorological hazards. The universities seek validation and suggestions for improvement of their research, as well as proposed new avenues for investigation via a series of quarterly reviews. These reviews are conducted at the campuses of the participating academic institutions, as well as at the sites of the governmental sponsors, the Federal Aviation Administration(FAA), and the National Aeronautics and Space Administration(NASA). This report represents a compilation and summary of the research conducted by the three academic partners, Massachusetts Institute of Technology, Princeton University and Ohio University. As such, it is intended to serve as an archival reference of this work for the federal fiscal year 1995(or academic year 1994-1995), as well as a means to communicate the notable accomplishments and findings of this sponsored research throughout FAA and NASA.

PURPOSE

The FAA/NASA Joint University Program (JUP) is a long-term cooperative partnership between FAA and NASA to pursue common research goals by promoting research and education in selected aviation technologies. The program, which awards grants annually to 3 American universities to support aeronautical research by faculty and students, is dedicated to the principle that solutions to large-scale systems problems in the National Airspace System (NAS) come only after the technological foundations have been laid through long-term basic and applied research. This report summarizes research conducted during fiscal year 1995 under the sponsorship of the JUP. As such, it draws on, compiles and compresses salient presentations from the four quarterly reviews conducted at the participating universities and one of the funding agencies(NASA).

BACKGROUND

Initiated in 1971, JUP was initially a NASA program, managed by the NASA Langley Research Center. FAA elected to participate in 1979, with FAA sponsorship centered at FAA Headquarters through 1985, and at the William J. Hughes Technical Center thereafter. An administrative framework consisting of two successive five year memoranda of understanding and interagency agreements between FAA and NASA, was employed to govern the conduct of the JUP from 1985 through 1995. Under these agreements FAA provided approximately 1/2 of the program financial support to NASA Langley for the program. By combining its funding with the FAA's, NASA Langley was able to award three grants annually to the participating universities for sponsorship of their research. These grants were NGL-022-009-640 to the Massachusetts Institute of Technology, NGR-36-009-017 to Ohio University, and NGL-31-001-252 to Princeton University. Funding was also included in the foregoing grants to provide for hosting of, participation in a series of four quarterly review conferences at the universities and one of the government sponsors(in 1995 this was NASA Langley).

JUP provides an interdisciplinary team approach to research and education in aviation technologies. By bringing this multi-agency, multi-university approach to bear on large scale NAS management and technical problems, highly original and creative solutions emerge. An additional benefit is the creation of a talented cadre of engineers and scientists who presently form a core of advanced aeronautical expertise in industry, academia, and Government.

This report summarizes the research conducted by the FAA/NASA Joint University Program during the Fiscal Year of 1995.

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

An Investigation of Air Transportation Technology at the Massachusetts Institute of Technology, 1994-1995

Robert W. Simpson, Flight Transportation Laboratory
R. John Hansman, Aeronautical Systems Laboratory
Department of Aeronautics and Astronautics
MIT, Cambridge MA 02139

Summary of Research

The Air Transportation Research Program at MIT focused on the following areas during the past year.

- Generalized Methods for Evaluating Alerting System Performance
- Closely-Spaced Parallel Runway Operations
- Measurement of Situation Awareness Using Testable Responses
- Mode Awareness in Advanced Autoflight Systems

The research supported partially or fully by the FAA/NASA Joint University Program has resulted in a number of reports, papers and journal articles. A complete list of reports and papers is provided in the attached bibliography. Several papers which illustrate the work are included below.

In the area of hazard alerting system performance, a generalized method is developed which allows formal analysis and trade studies to balance the probability of false alarms against the probability of missed detection. This method develops a System Operating Characteristic (SOC) curve for the alerting system which is analogous to a Receiver Operating Characteristic (ROC) curve from signal detection theory. This SOC formally defines the performance of the alerting system and provides the basis for trade studies. The SOC method has been evaluated against existing systems such as TCAS and current Ground Proximity Warning Systems (GPWS) and is found to be consistent with alerting thresholds which have been developed by ad hoc methods.

In the area of closely-spaced parallel runway operations, the hazard alerting methods described above have been applied to the problem of alerting flight crews to threats from aircraft on parallel instrument approaches. During closely-spaced operation the possible warning time is very short and careful consideration to system design and procedure are necessary. A prototype alert logic was developed and was tested in piloted part task simulation studies. Pilots were found to have low conformance to TCAS maneuver instructions during some conditions.

In the area of Situation Awareness Measurement, techniques were developed to evaluate situation awareness by specifically scripting testable responses into flight simulation experiments. If the scenarios are well designed (i.e., the appropriate response is clear if the subject is aware) then the response can be used to probe situation awareness. In contrast with other methods of assessing situation awareness, the testable response is though to be somewhat less invasive.

In the area of Mode Awareness in Advanced Autoflight Systems, accident, incident, and ASRS databases were evaluated to identify key mode awareness problems. In addition, the structure of current autoflight systems were studied based on flight manuals and interviews with pilots. Based on these studies, many of the problems were found to be in the vertical elements (altitude, airspeed, flight path angle, vertical speed, thrust) which are highly coupled. In addition, relatively little feedback is provided in the vertical channel compared with the strong feedback the moving map displays provide in the horizontal channel. In response, a prototype Electronic Vertical Situation Display (EVSD) was developed and preliminary simulator studies were conducted.

Bibliography

Reports:

Amar, M., Hansman, R.J., Hannon, D.J., Vaneck, T.W. and Chaudhry, A.I., "A Human Subject Evaluation of Airport Surface Situational Awareness Using Prototypical Flight Deck Electronic Taxi Chart Displays," MIT Aeronautical Systems Laboratory Report, ASL-94-7, August 1994.

Anagnostakis, I., Khayms, V., and Hansman, R.J., "Experimental Investigations of Physical Factors Which Influence Runoff and Failure Behavior of Aircraft De-Icing Fluids," MIT Aeronautical Systems Laboratory Report, ASL-94-8, August 1994.

Velazquez, M.T., and Hansman, R.J., "Preliminary Implementation of Improved Heat Transfer and Impingement-Direction Ice Growth Modeling in Numerically Simulated Ice Accretion (LEWICE/X)," MIT Aeronautical Systems Laboratory Report, ASL-94-9, August 1994.

Velazquez, M.T., and Hansman, R.J., "Preliminary Implementation of Improved Heat Transfer and Impingement-Direction Ice Growth Modeling in Numerically Simulated Ice Accretion (LEWICE/X)," MIT Aeronautical Systems Laboratory Report, ASL-94-9, August 1994.

Kuchar, J.K., and Hansman, R.J., "A Unified Methodology for the Evaluation of Hazard Alerting Systems," MIT Aeronautical Systems Laboratory Report, ASL-95-1, January 1995.

Johnson, E., and Hansman, R.J., "Multi-Agent Flight Simulation with Robust Situation Generation," MIT Aeronautical Systems Laboratory Report, ASL-95-2, January 1995.

Journal Articles:

Mykityshyn, M., Kuchar, J., and Hansman, R.J., "Experimental Study of Advanced Electronic Cockpit Displays for Instrument Approach Information," *The International Journal of Aviation Psychology*, **4**, No. 2, 141-166, July 1994.

Henry, R.C., Hansman, R.J., and Breuer, K.S., "Heat Transfer Variation on Protuberances and Surface Roughness Elements," *Journal of Thermophysics and Heat Transfer*, **9**, No. 1, 175-180, January 1995.

Bibliography (continued)

Conference Papers:

Velazquez, M.T., and Hansman, R.J., "Implementation of Combined Feather and Surface Normal Ice Growth Models in LEWICE/X," AIAA-95-0753, AIAA 33rd Aerospace Sciences Meeting, January 1995.

Anagnostakis, I., Hansman, R.J., and Khayms, V., "Microphysical Observations of Precipitation Impact and Melting on Non-Newtonian De-Icing Fluids," AIAA-95-0657, AIAA 33rd Aerospace Sciences Meeting, January 1995.

Orr, D.J., Breuer, K., and Hansman, R.J., "Quantitative Analysis of Ice Accretion Roughness Using Spectral and Stochastic Techniques," AIAA-95-0888, AIAA 33rd Aerospace Sciences Meeting, January 1995.

Clarke, J.P., and Hansman, R.J., "Parametric Study of the Noise Impact of Approach and Departure Using Advanced Flight Guidance Techniques," AIAA-95-0835, AIAA 33rd Aerospace Sciences Meeting, January 1995.

Amar, M., Hansman, R.J., Vaneck, T., Chaudhry, A., and Hannon, D., "A Preliminary Evaluation of Electronic Taxi Charts with GPS Derived Position for Airport Surface Situational Awareness," Eighth International Symposium on Aviation Psychology, April 1995.

Pritchett, A., and Hansman, R.J., "Variations on 'Party Line' Information Importance between Pilots of Different Characteristics," Eighth International Symposium on Aviation Psychology, April 1995.

Vakil, S., Hansman, R.J., and Midkiff, A., "Feedback Mechanisms to Improve Mode Awareness in Advanced Autoflight Systems," Eighth International Symposium on Aviation Psychology, April 1995.

Clarke, J.P., Hansman, R.J., and Fricke, J.R., "Determining Aircraft Exclusion Zones for Noise Sensitive Areas Using Acoustic Reciprocity and Linear Superposition," CEAS/AIAA First Joint Aeroacoustics Conference, Munich, Germany, June 1995.

Yang, L., and Hansman, R.J., "Application of the Analytic Hierarchy Process for Making Subjective Comparisons Between Multiple Automation/Display Options," 6th IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design, and Evaluation of Man-Machine Systems, June 1995.

Conference Papers: (continued):

Vakil, S., Hansman, R.J., Midkiff, A., and Vaneck, T., "Mode Awareness in Advanced Autoflight Systems," 6th IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design, and Evaluation of Man-Machine Systems, June 1995.

Kuchar, J., and Hansman, R.J., "A Probabilistic Methodology for the Evaluation of Alerting System Performance," 6th IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design, and Evaluation of Man-Machine Systems, June 1995.

A PROBABILISTIC METHODOLOGY FOR THE EVALUATION OF ALERTING SYSTEM PERFORMANCE

James K. Kuchar
R. John Hansman, Jr.

Department of Aeronautics and Astronautics
Massachusetts Institute of Technology
Cambridge, MA 02139

Abstract: A probabilistic methodology for evaluating hazard alerting systems is described that can be used in vehicle, transportation system, and process control applications. A means of showing the tradeoff between false alarms and missed detections is presented using signal detection theory concepts. The methodology accounts for uncertainties in measurement sources, alerting thresholds and displays, the human operator, and the situation dynamics. An example demonstration of the methodology is provided using the Traffic Alert and Collision Avoidance System (TCAS), an alerting system designed to prevent mid-air collisions between aircraft.

Keywords: Alarm systems, Human factors, Human supervisory control, Probabilistic models, Probabilistic risk assessment, Safety analysis, System analysis, System models

1. INTRODUCTION

A hazard alerting system is one of several safety components typically found in complex human-operated systems such as vehicles, traffic control systems, and process control applications. In these applications, the operator performs a task with the system using operating procedures and informational displays that provide feedback on the system and on the environment in which it operates (Sheridan, 1992; Wickens, 1992). Because unexpected hazards can be encountered, the alerting system serves to warn the operator that additional action may be needed to avoid an undesirable incident.

As the alerting system operates, it either remains silent or issues an alert that indicates that a hazard exists. Typically, this decision is based on whether certain states exceed critical values that form an alerting threshold. However, due to errors in measurements or limitations in design, faulty alerting decisions can occasionally occur. In particular, a system may fail to alert when necessary (termed a missed detection), or may issue an alert when one is not needed (termed a false alarm). Missed detections can result in an accident if the operator does not become aware of the hazard through other informational displays. False alarms have been shown to reduce the operator's confidence in the alerting system (DeCelles, 1991; Mellone and Frank, 1993). Alerting system performance is therefore dependent on the probabilities of missed detection and false alarm. Both types of error are undesirable and generally cannot be eliminated simultaneously. Rather, some tradeoff between false alarms and missed detections must be made.

A number of issues affect the ability of the alerting system to operate without missed detections and false alarms. Although many of these issues are common across applications, current design methods typically follow an *ad hoc*, evolutionary process. As alerting systems become more capable and more complex, however, the issues on which to focus design efforts become less evident. In particular, the effect of the human's response to an alert on the performance of the alerting system can be difficult to determine.

This paper presents a generalized model of alerting systems and a methodology for evaluating alerting system performance. By recasting the alerting decision as a signal detection problem, classical Signal Detection Theory methods can be used. Uncertainties in the human response and in

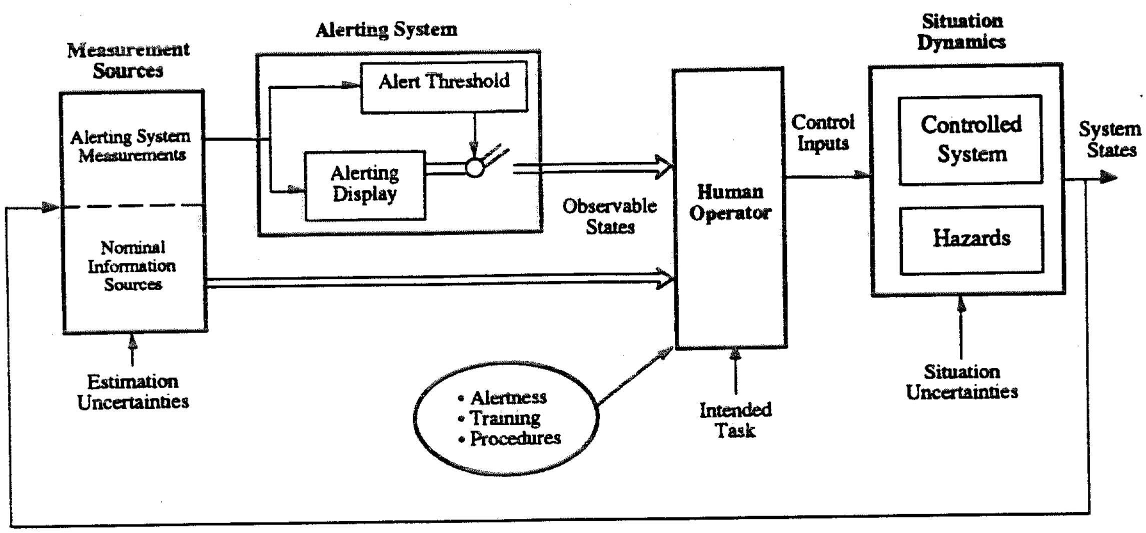

Fig. 1. Generalized Alerting System Model

measurements are considered. The methodology is then demonstrated through an example evaluation of the Traffic Alert and Collision Avoidance System (TCAS), used on civil jet transports.

2. ALERTING SYSTEM MODEL

Figure 1 shows a generalized control-system model of an alerting system that was developed by the authors. This model is based on a state-space representation of hazard situations. The state vector includes internal parameters such as a vehicle's spatial location, velocity, or acceleration, and external parameters describing a hazard's location, size, or severity. In this representation, there are four main elements in the control loop: Measurement Sources, Alerting System, Human Operator, and Situation Dynamics.

The Measurement Sources block represents the components that provide estimates of the system states to the human and to the alerting system. In general, these observable states are only a subset of the complete set of states that describe the situation.

The Alerting System uses a set of alerting thresholds to determine if an alert is warranted. When appropriate, alert information is provided to the Human Operator through aural or visual alerting displays.

The Human Operator uses information from the alerting system and other measurement sources to make necessary control inputs to the Situation Dynamics. The operator's response is also a function of the intended task and factors such as experience, fatigue, and training. These control inputs affect the evolving state vector, which is fed back to the Measurement Sources as the loop repeats.

Figure 2 shows an example situation in which the state vector is located in state-space near a hazard. Ideally, the alerting system should only alert the operator when it is clear that the hazard will be encountered unless action is taken. However, due to errors in the state measurements, the current state values are uncertain. In addition, errors in the extrapolation of the state trajectory (from an uncertain operator response, for example) result in a probabilistic future state trajectory. Thus, whether an alert is truly needed in a given situation is likewise uncertain, and missed detections or false alarms can occur. It is therefore necessary to consider the problem from a probabilistic standpoint to analyze the effects of uncertainties on the ability of the system to operate without missed detections and false alarms.

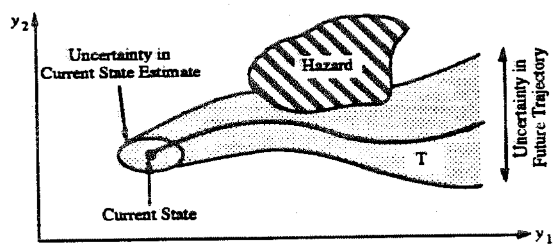

Fig. 2. Example Hazard Encounter Situation

3. PROBABILISTIC ANALYSIS METHOD

To determine if an alert is needed in a given situation, it is first necessary to find the probability that an encounter with a hazard will occur along a given trajectory, termed an *incident*, and denoted I. For example, in Figure 2 an alert may be necessary if the probability of an incident is high. If the probability of an incident is low, then an alert may not be needed.

Given a current state estimate, y, the probability that an incident will occur along a probabilistic trajectory T is written as $P_T(I \mid y)$. A generalized method for calculating $P_T(I \mid y)$ has been developed by Kuchar and Hansman (1995) and is outlined here.

The basic methodology for calculating $P_T(I \mid y)$ is shown in Figure 3. A set of equations that describe the dynamics of the situation and probability density functions (PDFs) describing the uncertainties in the parameters that define the trajectory are needed. For example, measurement uncertainties, the response time delay, and the aggressiveness of an avoidance maneuver can each be described by appropriate PDFs. The PDFs are obtained from hardware specifications or by a statistical analysis of actual or simulated hazard encounter situations.

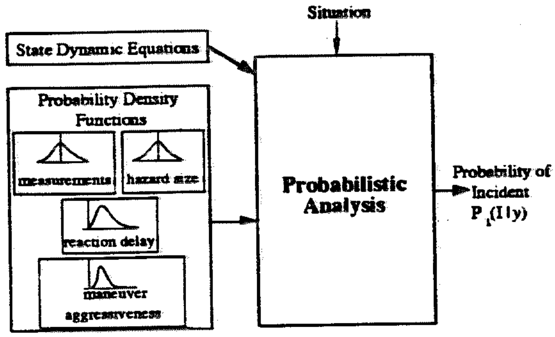

Fig. 3. Probabilistic Analysis Schematic

The procedure takes the equations of dynamics and the PDFs and, through numerical integration or Monte Carlo evaluation, produces the probability of an incident. The flexibility of the procedure lies in its ability to account for varied hazard types and to treat the PDFs as generic modules, allowing for a rapid comparison between the effects of different PDFs on $P_T(I \mid y)$.

3.1 Probabilistic Models

Because the measurements, the human response, and the system dynamics contain uncertainties, a probabilistic model of these components is needed to examine their effects on the overall alerting system. Measurement and dynamic uncertainties can generally be modeled based on hardware specifications. For example, a certain sensor may provide measurements that contain normally-distributed noise with a known variance.

The human's response to an alert is likewise uncertain. This uncertainty can generally be broken down into a probabilistic response latency and a set of random variables that specify the type of avoidance maneuver that is used. In an aeronautical application, for example, an avoidance maneuver to avoid terrain can be modeled as having a certain probabilistic aircraft pitch rate, bank angle, and thrust schedule.

3.2 Alerting Decision Outcomes

The alerting decision must balance the need to alert the operator sufficiently early that an incident can be avoided against the desire to only alert when absolutely necessary. To determine if an alert is warranted in a given situation, it is therefore necessary to examine the hypothetical outcomes of the alerting decision. This decision is analogous to the signal detection problem of determining if a signal is present in background noise (Swets and Pickett, 1982). In particular, if no alert is issued, the state continues along what is termed the projected *Nominal Trajectory*, denoted N. In response to an alert, there is, in general, a discrete change in the actions of the operator and the state follows a different trajectory that may avoid an incident, termed the *Avoidance Trajectory* and denoted A. Both N and A are, in general, probabilistic just as T is in Fig. 2.

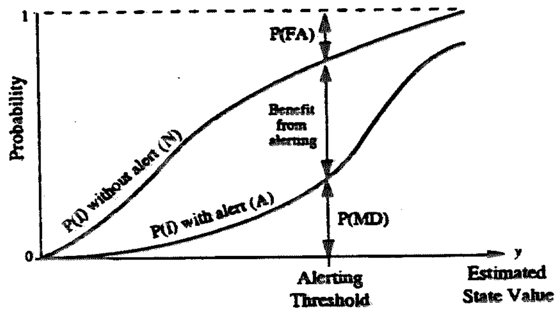

Fig. 4. Example Probability of an Incident Along Nominal and Avoidance Trajectories

An example plot of $P_N(I \mid y)$ and $P_A(I \mid y)$ is shown in Figure 4 as a function of an estimated state value, y. As shown, as the value of y increases, the probability of an incident occurring along N or A likewise increases. The probability that an incident would have been avoided without an alert is given by $1 - P_N(I \mid y)$, and is analogous to the probability of a false alarm, P(FA), in Signal Detection Theory. The probability that an incident will occur even though an alert is issued is given by $P_A(I \mid y)$, and is analogous to the probability of a missed detection, P(MD). One additional metric of system performance is often used: the probability of correct detection, P(CD), defined as P(CD) = 1 - P(MD). Also shown in Figure 4 is the benefit from alerting: the reduction in the probability of an incident that is possible because the alert is issued. As the threshold location is changed, there is a tradeoff between false alarms and missed detections. For example, moving the threshold to the left in Figure 4 increases the probability of false alarm, and therefore reduces the benefit from alerting.

3.3 System Operating Characteristic (SOC) Curves

The tradeoff between false alarms and missed detections can be viewed using a System Operating Characteristic (SOC) curve, similar to the Receiver Operating Characteristic (ROC) used in Signal Detection Theory (Swets & Pickett, 1982; Sheridan & Ferrell, 1974). SOC curves show the tradeoff between P(FA) and P(CD) as a function of the alerting threshold location, as shown in Figure 5. Each choice of an alerting threshold maps onto a single point along the SOC curve (two examples are shown in the figure).

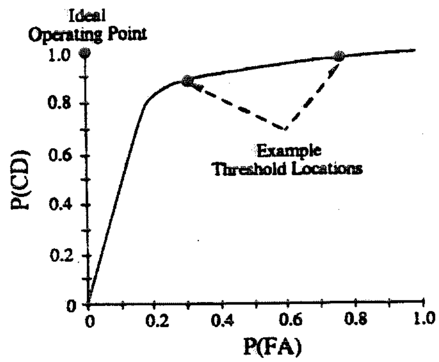

Fig. 5. Example System Operating Characteristic (SOC) Curve

Given a certain system design in terms of sensor accuracy, for example, the tradeoff between P(FA) and P(CD) is constrained to lie on a single SOC curve. P(FA) can then be balanced against P(CD) by changing the alerting threshold location, which changes the operating point on the SOC curve. Increasing sensor accuracy or improving the operator's response results in a shift in the SOC curve toward the upper-left ideal operating point.

Thus, given a definition of the system (which includes a probabilistic description of the human's response), the threshold location can be chosen according to the relative desirability of false alarms and correct detections. However, unless changes are made in the system design or in the operator's response, the system's performance will be constrained to follow a certain SOC curve.

4. EXAMPLE APPLICATION OF THE METHODOLOGY

The Traffic Alert and Collision Avoidance System (TCAS) is being used on transport aircraft in the U.S. to alert flight crews to potential mid-air collisions. This example addresses a situation in which an intruder is flying directly toward an aircraft equipped with TCAS (Figure 6). The intruder is currently above the TCAS aircraft and is projected to descend through the TCAS aircraft's altitude without leveling off. This situation has been known to produce false alarms in actual practice and provides an interesting example with which to apply the methodology.

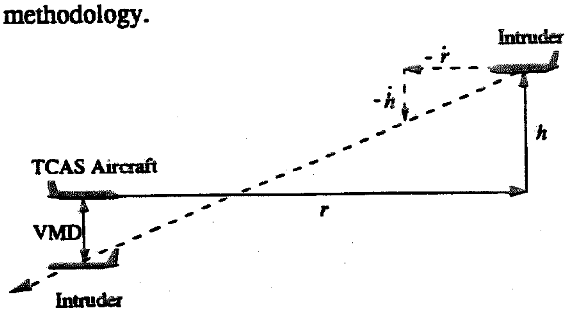

Fig. 6. Example Potential Collision Situation

Assuming straight-line extrapolations of each aircraft's trajectory, the Vertical Miss Distance (VMD) is a function of the range (r), range-rate ($\dot{r}$), relative altitude (h), and relative altitude-rate ($\dot{h}$):

$$\text{VMD} = h - \dot{h}\frac{r}{\dot{r}} \tag{1}$$

TCAS uses a complex set of alerting thresholds based on estimates of the four parameters discussed above (range, range-rate, altitude, altitude-rate). A complete description of the TCAS alerting logic (called Version 6.04A) is provided in a set of operating specifications (RTCA, 1983; MITRE, 1993). In the situation used in this example, TCAS V6.04A issues an alert approximately 22 seconds before the projected time of impact.

When a TCAS alert is issued, the system determines whether a climb or a descent will provide the largest vertical separation between the aircraft. In the situation used here, if the intruder continues its descent, a climbing maneuver by the TCAS aircraft will provide the greatest separation between the aircraft. If, however, the intruder levels off after the alert is issued, a potential collision is induced by the alert as the TCAS aircraft climbs into the now-level intruder (Figure 7).

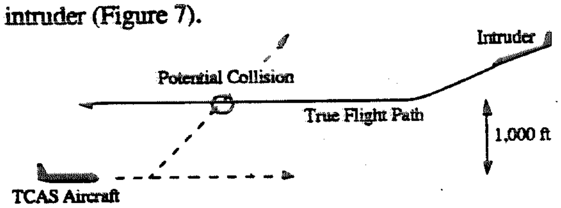

Fig. 7. Potential Induced Collision Due to Climb Alert

Therefore, TCAS must balance a desire to alert early and provide greater separation if the intruder continues to descend, and a desire to postpone alerting until the intruder's intentions are more well known.

4.1 Situation Parameters

The example described here assumes that the two aircraft are flying at an altitude of 4,500 m (15,000 ft) above Mean Sea Level, on a direct collision course at a closure rate of 200 m/s (400 knots). The intruder is descending toward the TCAS aircraft at 13 m/s (2,500 ft/min).

The intruder is assumed to continue its descent with probability 0.25 and to level off 305 m (1,000 ft) above the TCAS aircraft with probability 0.75. If the intruder levels off, it begins a maneuver such that a 0.1 g acceleration pull-up will result in a nominal separation between the aircraft of 305 m (1,000 ft). The actual altitude at which the level-off maneuver begins is modeled as a normal distribution with a standard deviation of 24 m (80 ft) (approximately 2 seconds of flight time).

If an alert is not issued, the TCAS aircraft is assumed to fly at a constant altitude (the Nominal Trajectory). In response to an alert, the TCAS aircraft flies a standard TCAS avoidance maneuver: 5 second delay followed by a 0.25 g acceleration pull-up maneuver to a climb rate of 7.6 m/s (1,500 ft/min) (RTCA, 1983). This climbing maneuver defines the Avoidance Trajectory.

A Monte Carlo simulation of the TCAS parameter estimation logic was performed to obtain the steady-state standard deviations of each state estimate. Approximate standard deviations on the parameters for this situation are shown in Table 1. Normally distributed probability density functions were created with the standard deviations shown in Table 1 to model the uncertainty in state estimates.

Table 1 TCAS Parameter Estimate Accuracies

Parameter	Estimate Standard Deviation
r	5.5 m (18 ft)
$\dot{r}$	1.8 m/s (6 ft/s)
h	25 m (83 ft)
$\dot{h}$	0.17 m/s (0.6 ft/s)

4.2 Human Response Model

The human response is assumed to include a probabilistic response delay, τ, followed by a 0.25 g pull-up maneuver to a 7.6 m/s (1,500 ft/min) climb rate. The response delay is modeled as a Gamma Distribution (Figure 8), which is a smooth, skewed distribution similar to that which might be expected (Hogg & Tanis, 1988). This PDF also provides a non-zero probability of very long response times representative of a flight crew that disregards or misunderstands an alert. A more detailed study of TCAS could use statistics from actual events or flight simulations to build other PDFs of response time delay.

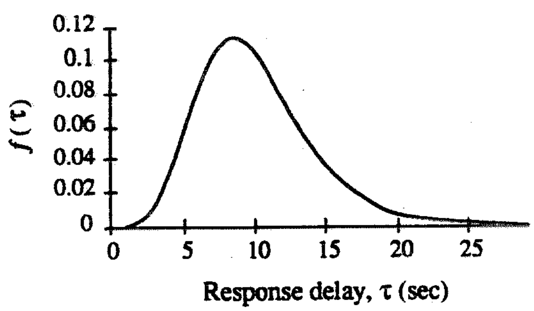

Fig. 8. Example Gamma Distribution $\alpha = 7$, $\theta = 1.4$

The Gamma Distribution is given by:

$$f_\tau(\tau) = \frac{1}{\Gamma(\alpha)\theta^{\alpha}} \tau^{\alpha-1} e^{-\tau/\theta} \qquad (2)$$

where $\Gamma(\alpha) = (\alpha - 1)!$

The mean, $\bar{\tau}$, of the Gamma Distribution is given by:

$$\bar{\tau} = \alpha\theta \qquad (3)$$

The PDFs used here have a parameter value of $\alpha = 7$ and θ is varied to provide a desired mean value.

4.3 Calculation of the Probability of an Incident

The PDFs used in this example, representing uncertainties in range, range-rate, altitude, altitude-rate, and response latency, were used to calculate the probability of an incident along the Nominal and Avoidance Trajectories. These PDFs were numerically integrated to determine the probability that an incident would occur along a given trajectory. This integration can be summarized by Equation (4).

$$P(I) = \int_{All\ PDFs} G(r,\dot{r},h,\dot{h},\tau) f(r) f(\dot{r}) f(h) f(\dot{h}) f(\tau) \qquad (4)$$

where, for example, $f(r)$ is the PDF for r, and G is a function that takes a value of 1 for combinations of parameters that result in a collision, and 0 otherwise. A collision is defined to occur if the Vertical Miss Distance is less than 30.5 m (100 ft).

4.4 Results

An SOC curve for this example is shown in Figure 9 for the case in which the mean response delay is 5 seconds. Note that the SOC curve for this example appears quite different from conventional ROC curves used in Signal Detection Theory, but the SOC curve can still be used in the same manner. As shown, the current alerting threshold location (TCAS V6.04A) is such that a near-minimum of false alarms are expected. Any further changes to reduce P(FA) will result in a rapid increase in P(MD). Due to the large uncertainties in the altitude of the intruder, the false alarm probability remains above approximately 0.4 regardless of the threshold location.

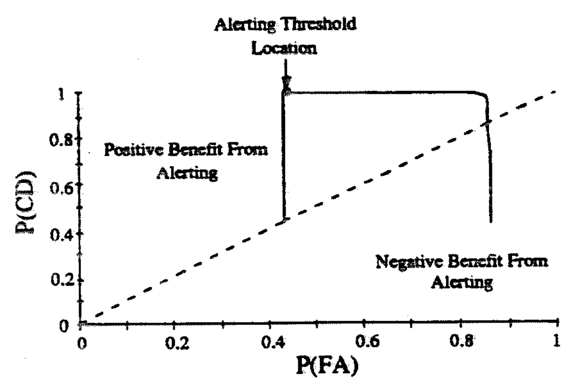

Fig. 9. System Operating Characteristic (SOC) Curve for TCAS Example ($\bar{\tau}$ = 5 sec).

Only by increasing the breadth or accuracy of information available to TCAS could the false alarm probability be reduced. There is also a region on the SOC curve where P(CD) is less than P(FA): it is more dangerous to alert than to not alert. This region corresponds to alerting thresholds that are set such that there is a greater probability of climbing into a leveled intruder than colliding with an intruder that continues its descent.

Figure 10 shows the probability of an incident when an alert is issued, P(MD), at the V6.04A threshold location as a function of the mean response latency, $\bar{\tau}$. A mean response delay of 5 sec results in a value of P(MD) of approximately $2x10^{-5}$. An increase in the mean response time to 7 seconds results in an increase in the value of P(MD) to approximately $5x10^{-4}$. Thus, an additional two second delay over that assumed by TCAS logic reduces the safety of the system by a factor of 25.

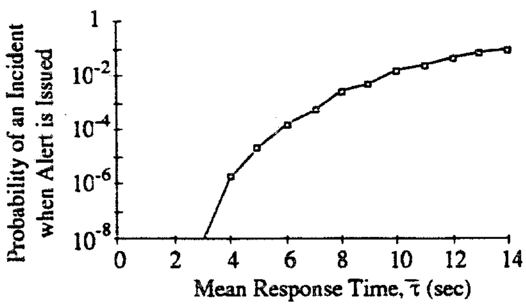

Fig. 10. Effect of Mean Response Time on Probability of an Incident.

5. CONCLUSION

A methodology for modeling and evaluating hazard alerting systems has been developed. This methodology accounts for uncertainties in measurements, the human response to an alert, and situation dynamics. The alerting decision is shown to be analogous to problems in Signal Detection Theory. This connection enables the use of conventional analysis tools to examine the tradeoffs between false alarms and missed detections. Parametric studies are also possible to examine the effect of the human response on alerting system performance.

The methodology is demonstrated in an application using the Traffic Alert and Collision Avoidance System (TCAS) in use on U.S. jet transport aircraft. For the situation covered here, the current alerting threshold is shown to be located to effectively minimize false alarms before a rapid increase in missed detections occurs. The sensitivity of the alerting system's performance is also shown as a function of the mean response latency. This example serves to illustrate the basic concepts of using probabilistic analysis to evaluate alerting systems. A more detailed examination is certainly possible, and will be necessary in an actual design situation.

The methodology is presented here using an aeronautical example. However, the methodology is easily extended to any alerting system application, including other vehicle types, large-scale transportation systems, and process control or medical applications.

ACKNOWLEDGMENT

This work was supported by the Federal Aviation Administration under Grant #92-6-0001, and by the Boeing Commercial Airplane Company.

REFERENCES

DeCelles, J. L. (1991). The delayed GPWS response syndrome. Aviation Research & Education Foundation. Herndon, VA. July.

Hogg, R. V., and E. A. Tanis (1988). *Probability and Statistical Inference*. Macmillan. New York.

Mellone, V. J. and S. M. Frank (1993). Behavioral impact of TCAS II on the national air traffic control system. *Seventh International Symposium on Aviation Psychology*. The Ohio State University. April 27.

The MITRE Corporation (MITRE) (1993). *Final V6.04A Change Pages*. Courtesy of Andrew D. Zeitlin. September 17.

Kuchar, J. K. and R. J. Hansman (1995). A Unified Methodology for the Evaluation of Hazard Alerting Systems. MIT Aeronautical Systems Laboratory Report No. ASL-95-1.

Radio Technical Committee on Aeronautics (RTCA) (1983). *Minimum Performance Specifications for TCAS Airborne Equipment*. Document No. RTCA/DO-185. Washington, D.C.

Sheridan, T. B. and W. R. Ferrell (1974). *Man-Machine Systems: Information, Control, and Decision Models of Human Performance*. MIT Press. Cambridge, MA.

Sheridan, T. B. (1992). *Telerobotics, Automation, and Human Supervisory Control*. MIT Press. Cambridge, MA.

Swets, J. A. and R. M. Pickett (1982). *Evaluation of Diagnostic Systems*. Academic Press. New York.

Issues in Airborne Systems for Closely-Spaced Parallel Runway Operations

A. Pritchett, B. Carpenter, K. Asari, J. Kuchar, and R.J. Hansman
Department of Aeronautics & Astronautics
Massachusetts Institute of Technology
77 Mass. Ave Room 37-447
Cambridge MA 02139
(617) 253-7748
amyruth@mit.edu

ABSTRACT

Efforts to increase airport capacity include studies of aircraft systems that would enable simultaneous approaches to closely spaced parallel runways in Instrument Meteorological Conditions (IMC). The time-critical nature of a parallel approach results in key design issues for current and future collision avoidance systems. These issues are being studied in two ways. First, a part-task flight simulator study has examined the procedural and display issues inherent in such a time-critical task. Second, a prototype collision avoidance logic capable of generating this maneuver guidance has been designed using a recently developed methodology.

INTRODUCTION

To reduce flight delays and increase airport capacity, several methods of enabling closely spaced, independent parallel approaches in Instrument Meteorological Conditions (IMC) are being studied. Without specialized radar, current criteria allow independent parallel approaches to runways spaced 4300 feet or more apart; the use of new technologies to reduce this minimum separation would allow airports to effectively maintain their Visual Meteorological Conditions (VMC) capacity in IMC.

The task of ensuring adequate aircraft separation during parallel approach operations is very difficult. The aircraft are closer together than during any other airborne phase of flight, which severely limits the potential warning time should one aircraft blunder into the other's approach path. Studies have shown that, for runways at least 4300 feet apart, the controller using today's radar can ensure aircraft separation. [1] The Parallel Runway Monitor (PRM) system, which uses special radar with a high update rate, has been implemented at Raleigh-Durham and Memphis airports. Using special displays, Air Traffic Controllers can determine if an aircraft will enter a 'No-Transgression Zone' (NTZ) between the two approach paths and can give each aircraft commands to steer away from a potential collision. Recent studies have concluded this system can be used to reduce runway separation to 3400 feet. [2,3]

Further reduction of the runway spacing using PRM has not been recommended. Not only are the required reaction times reduced, but aircraft, given the limits of their localizer tracking performance, may occasionally enter the NTZ while attempting to track the localizer, causing nuisance alerts and aborted approaches. [3]

Overview

This paper details two studies of airborne systems capable of ensuring adequate aircraft spacing during parallel approaches in IMC. First, a baseline flight simulator study examined the pilot responses to a potential collision, both with and without the aid of an alerting system. The current Traffic Alert and Collision Avoidance System (TCAS II) was used as the baseline alerting system. TCAS II was not designed for close parallel approaches, however, and would generate a high number of false alarms for runway separations less than 3000 feet [4]

The second study involved developing a prototype alerting logic specifically for parallel approach conditions. By taking into account the constrained aircraft trajectories during parallel approaches and by using cross-link of information between aircraft, the system can allow a further reduction of aircraft separation. Using a recently developed methodology, a probabilistic analysis of the system's performance has been made. [5] The performance of this system has been compared to both the performance of the pilots in the flight simulation study and the theoretical performance of the TCAS II system.

FLIGHT SIMULATOR STUDY

A preliminary simulator experiment had active airline pilots fly many parallel approaches on the MIT part-task Advanced Cockpit Simulator in order to study the pilot effectiveness in avoiding encroaching traffic, both with and without the aid of an alerting system. During each approach, traffic on a parallel approach (to runways separated by 4300 feet) would blunder towards the subject, and the subject's response was recorded to find the allowable maneuver strength and reaction time. The study also examined several cockpit traffic display enhancements, and the relative merits of flying the approach (before any avoidance maneuver) manually or on autopilot.

The MIT Advanced Cockpit Simulator provides pilots with the relevant controls and displays of a generic glass cockpit aircraft. A Silicon Graphics workstation provides the display of the glass cockpit screens and traffic displays; it also calculates the dynamics of the simulator, which has the performance of a Boeing 737. The pilot can use the Flight Management Computer, Mode Control Panel or sidestick to control the aircraft. An experimenter acts as co-pilot, setting gear, flap and autopilot settings as commanded by the subject.

A second Silicon Graphics workstation steered the 'intruder' aircraft on an approach parallel to the subject's, and then turned the intruder into the subject at a scripted point during the scenario. This Robust Situation Generation system made possible repeatable, scripted near-collisions while allowing flexibility for varied flight paths between pilots. [6]

The 18 subjects were qualified airline flight crew from two major airlines, with a mean of over 15,000 total flight hours. All but one were considered current on glass cockpit aircraft.

Each subject flew a total of 36 approaches. These approaches were flown in 12 blocks of three. Each of the 12 blocks were flown under a different condition, representing all the combinations of four different traffic displays and three different procedures. The test matrix was counter-balanced between pilots to reduce any learning effects.

The study tested four displays: a TCAS traffic display integrated with the Electronic Horizontal Situation Indicator (EHSI); enhancements to the current traffic display on the EHSI, including an indication of the localizer beams for both runways and a split screen; a display of the parallel approach traffic on the pilot's Primary Flight Display (PFD); and a combination of the new displays on the PFD and EHSI.

Three procedures were studied: the subject monitors an autopilot approach, and then takes manual control to follow the alerts and avoidance maneuvers shown by a TCAS II - type system; the subject manually flies the approach, and follows the alerts and avoidance maneuvers shown by a TCAS II - type system; and the subject handflies the approach but is not shown any alerts or avoidance maneuvers.

Within each test block, each subject flew three approaches. Using Robust Situation Generation, these three approaches were scripted to each be one of three types. One represented an intruding aircraft that has overshot its own localizer and is established on a collision course with the subject's aircraft. The next type represents an intruding aircraft that strays from its own approach course to a collision course. The final type represents an intruding aircraft that, while straying enough from its path to generate a TCAS Resolution Advisory (RA), is established on a trajectory that should cause it to pass at least 1000 feet away from the subject. Several different approaches were scripted for each type so that the subjects could not second guess when evasion maneuvers would be needed.

The primary goal is to ensure adequate separation between aircraft on parallel approaches. Therefore, the first measurement of interest is the resulting miss distance between aircraft. Overall, the intruder and subject aircraft came within 500 feet of each other 4% of the time, and within 1000 feet of each other 20% of the time. These percentages were found to be significantly lower when the approach was flown on autopilot and significantly higher when TCAS avoidance maneuvers were not displayed.

The characteristics of the avoidance maneuvers can be described as follows: the mean load factor of the pitch maneuver was .59 'G's and, when the subjects performed a turning avoidance maneuver, they used a mean bank angle of 19 degrees; given the part-task, fixed-based level of simulation, however, these values may not be exactly those which would be used in the real aircraft. The pilots' mean reaction time to a displayed alert was 3.0 seconds (discarding values beyond three standard deviations).

During 16% of the approaches the subjects performed an avoidance maneuver so early that a TCAS alert was never given. These early go-arounds often occurred long before the intruder aircraft was straying from its proper approach path, and may indicate the pilots' concern over the unusual proximity of these aircraft. Significantly fewer early go-arounds were commanded during the approaches flown on autopilot.

With the presentation of a TCAS generated maneuver comes the assumption that the pilot will follow it, both by reacting within five seconds, and then by matching or exceeding the TCAS pitch command. However, examination of the trajectories after-the-fact has shown that the actual maneuvers flown by the pilots, when the TCAS maneuvers were shown, met the vertical maneuver commanded by the TCAS only 40% of the time.

No single causal factor of the low conformance rate can be isolated. Pilot reaction time alone does not show a strong effect. 66% of the pilots reacted within the five second allowance assumed by the TCAS system, and of these only 61% matched the displayed TCAS maneuver. Of the pilots who acted shortly before the alert or after the five second allowance (13% and 20% respectively), a significant number of pilots still matched what the TCAS guidance commanded (71% and 33% respectively).

Conformance to the (vertical) TCAS maneuver may be affected by the turning maneuvers that the pilots often performed at the same time. Overall, pilots did not turn in 32% of the approaches (ie. the maximum bank angle after the alert was less than five degrees); 34% of the time the pilots turned away from the intruder, 11% of the time pilots turned toward the intruder, and 23% of the time pilots turned one way and then another. Pilots who did not follow the TCAS maneuver turned away from the intruder significantly more often than pilots who followed the TCAS maneuver; this may suggest that the pilots, by executing a turn, felt a vertical maneuver was no longer required.

Pilots, given the enhanced traffic displays tested in this experiment, conformed significantly less often than when they were given the current TCAS II type traffic display. This may also suggest that pilots, given a more explicit traffic picture, may have felt a vertical maneuver was not longer required; this perception may have been erroneous, however, as more near-misses happened with these new displays.

Other possible factors for the low conformance rate have also been investigated. Examining the aircraft trajectories for the approaches where the pilots were not shown any TCAS alerts or maneuver guidance, the pilots' reactions only satisfied what the TCAS would have commanded in 25% of the approaches, suggesting that the TCAS maneuver is not what the pilot would do instinctively. As well, the conformance rate varies widely between pilots, from a high of 68% to a low of 25%.

The mere presentation of the TCAS alerts caused a significant improvement in aircraft miss distance, whether or not the TCAS maneuvers were followed exactly by the pilots. As shown in Figure 1, more incidents were caused when the pilots were not shown alerts, regardless of whether their maneuver happened to match what TCAS would have commanded; when the pilots were shown alerts and an avoidance maneuver, far fewer incidents occurred. Significantly fewer incidents happened when the pilots, shown an avoidance maneuver, conformed to it.

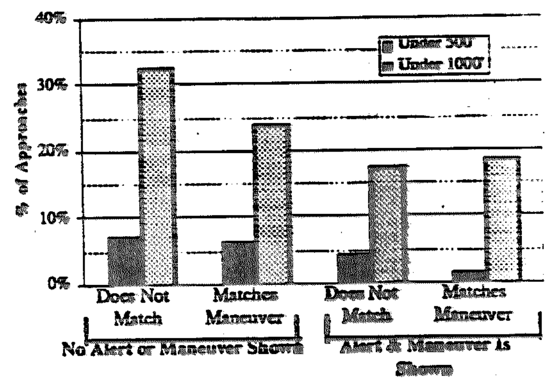

Figure 1. Frequency of Incidents With and Without Presentation of Alerts & Avoidance Maneuvers, and Pilot Conformance to Them

Pilot ratings during the debriefing showed an overwhelming preference for parallel approaches in IMC to be flown on autopilot; most pilots cited both decreased workload, which would allow them to monitor the parallel traffic better, and the reliability of the autopilot system.

Pilot rankings of the preferred role of the collision avoidance system showed a strong preference for a TCAS II - type alerting system, ie. one that provides avoidance maneuvers for the pilot to follow manually. A lack of any alerting system and a completely automatic avoidance system were ranked the lowest.

PROTOTYPE ALERTING LOGIC

A prototype alerting system logic was developed as the second study discussed in this paper. TCAS II uses the projected time to impact to determine whether an alert should be issued, and it has been successful in resolving several conflicts. However, when using TCAS for closely spaced parallel approaches, nuisance alarms can occur as the parallel traffic oscillates along its approach path. Furthermore, when parallel traffic does blunder, TCAS may not provide enough warning time to avoid an accident. It was felt that a specialized alerting system designed specifically for parallel approach could improve safety while producing fewer nuisance alerts than TCAS.

The prototype alerting system uses estimates of intruder position, heading, and bank angle to determine whether the intruder is in a position to potentially cause a collision. Differential Global Positioning System position accuracy is assumed, and heading and bank angle errors are assumed to be normally distributed with standard deviations of 2.5° and 5° respectively. Based on these errors, a probabilistic analysis of the parallel approach situation was performed to determine the probability of a collision. This approach was based on a previously-developed methodology to evaluate alerting system performance and to illustrate the tradeoffs between false alarms and accidents when alerting thresholds are designed. [5]

Given a particular intruder relative position, heading, bank angle, and velocity, the probability that a collision will occur was estimated using Monte Carlo simulations. These simulations assumed that the intruder aircraft flew a constant-rate turn and remained at the same altitude throughout the event.

To examine whether an alert based on a particular intruder state is appropriate, several potential future trajectories for the own aircraft were examined over a range of measured intruder positions, headings, bank angles, and velocities. Once an alert is issued the own aircraft performs an avoidance maneuver following a response delay. Several avoidance maneuvers were examined, including a 0.25g pitch up until a 2,000 ft/min climb rate was achieved and a 10°/sec rolling maneuver to a 30° bank angle, held until a 30° heading change away from the intruder was achieved. A combined climbing and turning maneuver was also examined. For comparison, the miss distance achieved without an avoidance maneuver was also determined.

By examining the curves of collision probabilities for the two potential future trajectories of the own aircraft (non-maneuvering and maneuvering), it is possible to determine whether an alert is appropriate. For example, an alert should be issued when there is a high probability of a collision if the own aircraft does not maneuver. The probability of a collision, even when the own aircraft does maneuver, provides a measure of the timeliness of the alert: if this probability is also high, then the alert may be too late to prevent an accident.

Figure 2 shows a representative plot of two regions within which the probability of collision is greater than 0.001, both for a non-maneuvering and maneuvering own aircraft. The choice of probability levels of 0.001 is based on PRM safety levels.

In the situation shown in Figure 2, the own aircraft is shown at the origin. The intruder's measured heading is parallel to the own aircraft but the measured bank angle is 15° toward the own aircraft. If the intruder is located at position 1, then the probability of a collision, regardless of what the own aircraft does, is below 0.001 and an alert could be considered unnecessary. In effect, an aircraft at position 1 is unable to collide with the own aircraft without a great increase in speed. An intruder at position 2 is projected to collide with the own aircraft unless a climbing turn avoidance maneuver is performed, warranting an alert. An intruder at position 3 is projected to collide with the own aircraft if no avoidance maneuver is performed, but an alert could be delayed because time is yet available to collect information about the intruding aircraft's flight path.

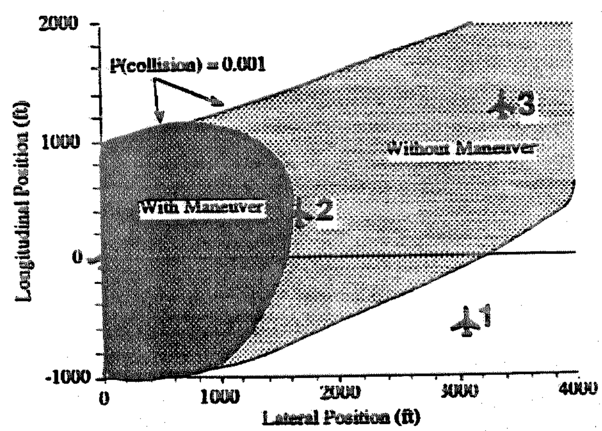

Figure 2.
Probability Contours and Alert Threshold

Thus, a prototype alerting threshold was constructed along the thick solid line in Figure 2. When the intruder crosses this threshold, the probability of a collision rises above 0.001 and an alert is issued. The own aircraft should then perform a climbing turn avoidance maneuver. A different threshold must be constructed for each combination of intruder heading and bank angle.

EVALUATION OF ALERTING METHODS

The prototype alerting method was compared to both the theoretical performance of TCAS II-type thresholds and to the achieved performance of pilots using the aircraft track data from the part-task simulation studies. When an alert was issued from the prototype system, the own aircraft was assumed to perform a climbing-turn maneuver following a 5 second delay. A TCAS-like system was also examined using slightly modified TCAS II alerting logic. When a TCAS alert was issued, the aircraft performed a vertical avoidance maneuver following a 5 second delay. Third, the achieved performance of pilots was also compared using the observed results from the study, for the cases where the pilots were and were not shown TCAS II-type maneuvers.

A collision was defined to occur if the distance between aircraft was below 500 ft at any time during the run. The behavior of the alerting system was classified into one of six categories, as summarized in Table 1.

Table 2 shows the following: the achieved performance by pilots during the flight simulation study, both when they were and were not shown alerts, and the theoretical performance of the exact TCAS II type system and of the prototype system.

The pilots' reactions were very conservative. The most false alarms were generated by the pilots when TCAS avoidance maneuvers were presented to them. However, without the presentation of TCAS avoidance maneuvers, some potential collisions were never spotted, resulting in collisions. In addition, the pilots' reactions, when required to avoid a collision, were usually strong enough to successfully evade the intruding aircraft; however, some pilot reactions caused an accident that otherwise would not have occurred.

The TCAS II-type maneuvers, flown automatically after a five second delay, were also very conservative, correctly detecting all possible intrusions but generating many false alarms. Like the pilots' maneuvers, the TCAS system caused some accidents; unlike the pilots' maneuvers, some TCAS maneuvers were of insufficient strength to successfully evade the intruder.

The prototype alerting system performed the best, with more correct detections and significantly fewer false alarms. However, the single type of avoidance maneuver flown after an alert was insufficient in some cases and caused collisions in others.

Category	Correct Rejection	Correct Detection	False Alarm	Missed Detection	Insufficient or Late Alert	Induced Collision
Alert Issued?	No	Yes	Yes	No	Yes	Yes
Collision Without Maneuver	No	Yes	No	Yes	Yes	No
Collision With Maneuver	N/A	No	No	N/A	Yes	Yes

Table 1. Situation Outcome Categories

Category	Correct Rejection	Correct Detection	False Alarm	Missed Detection	Insufficient or Late Alert	Induced Collision
Pilot, Without TCAS II System	12%	3%	80%	1%	5%	4%
Pilot, With TCAS II System	8%	2%	86%	0%	2%	4%
Theoretical TCAS II Type System	13%	2%	79%	0%	4%	3%
Prototype Alerting System	54%	4%	39%	0%	2%	2%

Table 2. Comparison of Achieved and Theoretical Outcomes

Further improvements to the prototype alerting system may help eliminate these remaining problems. This logic was designed assuming the intruder would perform a constant-rate turn at constant altitude and at constant airspeed. A significant portion of the Insufficient Alerts (64%) occurred in scenarios in which the intruder greatly increased airspeed during the interval between the alert and the closest point of approach. The development of alerting thresholds can easily be modified to accommodate such maneuvers resulting in system which can guarantee the same level of safety while minimizing false alarms.

Similarly, the alerting logic currently dismisses any information it receives about the intruder's climb rate. The logic senses a potentially dangerous situation and suggests a climbing turn avoidance maneuver, without regard to the climb rate of the intruder. A preliminary study of assigning avoidance maneuvers based on the intruder's climb rate has shown a significant decrease in the number of those false alarms which will cause accidents.

The False Alarm rate, while high, may also reflect that unusual proximity of the aircraft during close parallel approaches. An alert was termed a False Alarm if the own and intruding aircraft would have missed each other by more than 500 feet; this threshold may be less conservative than pilots are currently accustomed to. An investigation of the false alarm trajectories finds that the aircraft often would have missed by a distance less than 1000 feet, suggesting an alert was still valuable in these cases.

CONCLUSIONS

Both pilot comments and flight simulator results show the benefits of displaying to the pilot an avoidance maneuver for them to fly. Although the low conformance of pilots to their displayed TCAS maneuvers is surprising, it is difficult to determine whether it is good or bad. Numerical studies show that both the maneuvers flown by pilots and the maneuvers commanded exactly by the TCAS systems have their flaws. However, these results do show that pilot responses are variable and must be taken into account when designing a pilot-in-the-loop alerting system.

Numerical simulations provide several insights. The presence of an alerting system eliminates missed detections. Pilots and the TCAS II - type system are conservative and generate many false alarms. The prototype system is successful in improving the correct detection rate while reducing the amount of false alarms.

Some improvements may be advisable for the prototype alerting system. Measurement of the intruder's relative velocity should be integrated into the design of the alerting thresholds. Also, further work should investigate updating the suggested avoidance maneuver based on real-time intruder state measurements.

ACKNOWLEDGMENTS

The flight simulation study was supported by the National Aeronautics and Space Administration / Ames Research Center under grant NAG 2-716. The prototype alerting system development was supported by the National Aeronautics and Space Administration / Langley Research Center. The authors would also like to thank the pilots who participated in the study.

REFERENCES

[1] "PRM (Precision Runway Monitor System)" FAA Document produced by MSI Services Inc. November 1989

[2] A.T. Lind "Two Simulation Studies of Precision Runway Monitoring of Independent Approaches to ely Spaced Parallel Runways" DOT/FAA/NR-92/9 March 1993

[3] M.R. Owen "The Memphis Precision Runway Monitor Program Instrument Landing System Final Approach Study" DOT/FAA/NR-92/11 May 1993

[4] V.A. Folmar, X. P Szebrat, & N.E. Toma "An Extension to the Analysis of Traffic Alert and Collision Avoidance System (TCAS) Advisories During Simultaneous Instrument Approaches to Closely Spaced Parallel Runways" MITRE Paper MTR 94W0000056 July 1994

[5] J. K. Kuchar & R.J. Hansman, "A Unified Methodology for the Evaluation of Hazard Alerting Systems" ASL Paper 95-1, January 1995

[6] E.N. Johnson & R.J. Hansman, "Multi-Agent Flight Simulation with Robust Situation Generation" ASL Paper 95-2, January 1995

Preprint, International Conference on Experimental Analysis and Measurement of Situation Awareness
Nov. 1-3, 1995, Daytona Beach, Florida

USE OF TESTABLE RESPONSES FOR PERFORMANCE-BASED MEASUREMENT OF SITUATION AWARENESS

A. R. Pritchett, R.J. Hansman & E.N. Johnson
MIT Aeronautical Systems Laboratory
MIT Room 37-447
77 Mass. Ave
Cambridge MA 02139
email: amyruth@mit.edu

INTRODUCTION

The use of testable responses as a performance based measurement of situation awareness is a valuable measurement technique for testing of a wide-range of systems. Unlike measurement techniques that attempt to ascertain the subject's mental model of the situation at different times throughout an experiment, performance based testing focuses solely on the subject's outputs. This quality makes it ideal for comparing the desired and achieved performance of a human-machine system, and for ascertaining weak points of the subject's situation awareness.

This paper will focus on the use of situations with testable responses during simulations. During the simulation runs, the subjects are presented with situations. The situations are designed such that, if the subject has sufficient situation awareness, an action is required. This provides an unambiguous accounting of the types of tasks for which the pilots had sufficient situation awareness.

First, this method of assessing situation awareness will be briefly compared with other methods. The use of situations with testable responses in a representative flight simulator study will be detailed. Then, because the subject's responses depend heavily on the precision with which the situations are generated, techniques for robust generation of pre-determined situations will be discussed, and the performance of a current implementation will be discussed.

A COMPARISON OF PERFORMANCE BASED MEASUREMENT WITH OTHER METHODS OF SITUATION AWARENESS ASSESSMENT

Performance-Based Measurement of Situation Awareness has taken several forms. Some techniques measure the overall final performance of the human-in-the-loop system in any or all of its tasks (Endsley, 1995). This paper focuses on the use of Testable Responses for evaluating situation awareness, where the subjects are presented with realistic situations during the simulation runs which, if they have sufficient situation awareness, require decisive and identifiable actions.

Several other methods of testing situation awareness have been documented (Endsley, 1995; Adams, Tenney & Pew, 1995). Several complex techniques exist which attempt to determine or model the subject's knowledge of the situation at different times throughout the simulation runs. For example, the Situation Awareness Global Assessment Technique (SAGAT) freezes the simulator screens at random times during the runs, and queries the subjects about their knowledge of the environment. This knowledge can be at several levels of cognition, from the most basic of facts to complicated predictions of future states.

Several causal factors affect the actions of the subject, as shown in Figure 1. Comparing knowledge-based and performance-based techniques of evaluating situation awareness, we find they take measurements at different points in the process of user cognition. This illustrates the different purposes for these two measurement techniques.

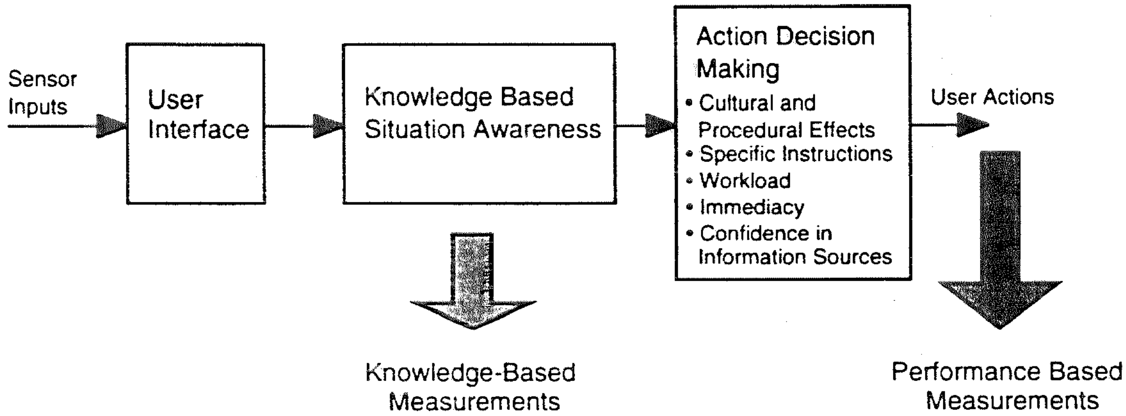

Figure 1. A Comparison of Measurement Points Between Knowledge-Based and Performance-Based Measurement Techniques

For providing a detailed, theoretical assessment of the subject's situation awareness, the knowledge based techniques are more accurate, as they measure these variables directly. Performance-based measurement can only make inferences based upon the particular information the subject acted upon, and how it was interpreted.

However, performance-based measurements can satisfy several goals that knowledge-based techniques can not. The most apparent is its ability to ascertain the timing and substance of a user's reaction to realistic situations. For testing of systems, final decisions must be based on whether the user will be provided with sufficient situation awareness to perform the correct actions, which performance-based techniques measure directly. Knowledge-based measurement techniques, on the other hand, can only make reasonable guesses about the likely user's actions given their knowledge state.

In addition, performance-based measurement provides measures of situation awareness that are not otherwise easily achievable. It can identify constraints on a user, arising from their training and standard procedures, that would not be anticipated by a strict knowledge-based model of situation awareness. For example, in a flight simulator study by Midkiff and Hansman, ATC neglected to turn the subject towards the landing runway although the subjects could overhear the aircraft before and after them being giving the proper instructions; although the subjects' actions indicated they were aware of the situation, they did not take a strong reaction because of their reticence to assume the Air Traffic controller had made an error (Midkiff & Hansman, 1993). A knowledge-based measurement of the pilots situation awareness also would have provided a measurement, in this case, of the pilot's awareness of the problem; only performance-based measurement, however, could ascertain how the pilots would act upon this information within an established set of Air Traffic Control procedures.

Performance-based measurement is also able to determine perceived reliability of the knowledge users gather from any of a multitude of sources. For example, the same simulation study by Midkiff and Hansman found pilots were often unwilling to act upon information only overheard on ATC voice frequencies because they did not have confidence in the mental model it provided (Midkiff & Hansman, 1993). The study was therefore able to ascertain whether pilots had sufficient confidence in their mental model to take action. A knowledge-based measurement, in the same study, might have concluded that the pilots had correct knowledge, but might not realize the pilots would refuse to act upon it in the same manner as if they had verifiable, correct knowledge.

Finally, performance-based measurement works well in time-critical situations to find the real-time response, rather than planned or thought-through response. Subtle variations in situation awareness or current conditions may be causal factors in different actions by the user, as shown in autopilot mode-awareness simulation, where the pilot's actual, real-time reactions often varied significantly from those they named as 'what they would do' during non-time critical questioning afterwards (Johnson & Pritchett, 1995).

In summary, performance-based measurement is complementary to knowledge-based measurement in the development of a human-in-the-loop system. Each is useful at different times, and for different purposes, throughout the design process. For final testing of a system, performance-based measurement is very useful because of its ability to ascertain the resulting performance of the entire system, and to point to areas of situation awareness that are deficient. Although performance-based measurement does not provide as pure a measurement of a user's knowledge base as other techniques, it is able to illustrate the inter-relationship between the user's knowledge and the manner in which they use it.

USE OF SITUATIONS WITH TESTABLE RESPONSES IN A REPRESENTATIVE FLIGHT SIMULATION STUDY

This section shall use a recent flight simulator study to demonstrate the use of testable responses in measuring situation awareness and overall system performance. Both the development and performance of the measurement techniques shall be discussed.

The flight simulator study by Midkiff & Hansman was conducted to evaluate pilot utilization of the Party Line Information they can overhear on shared Air Traffic Control frequencies (Midkiff & Hansman, 1993). Two-pilot air transport flight crews, using the NASA Ames Man-Vehicle System Research Facility (MVSRF), flew a 3 leg flight, during which they were exposed to nine different situations.

The design and scripting of the situations is the most crucial aspect of the experiment design. The situations must be designed to have several traits. Most importantly, the situations must be designed such that, should the user have sufficient situation awareness, a clear and unambiguous response is mandated. As illustrated in Figure 2, the task of the experimenter is to expose the user to situations which force a measurable action, without attempting to examine the specifics of the 'inner' workings of the subject, such a their knowledge state.

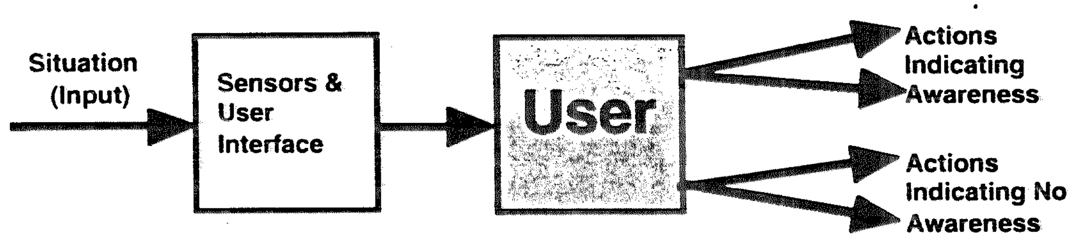

Figure 2. Use of Testable Responses to Situations

When expert-users, such as airline pilots, are used as subjects, situations can be chosen for which standard operational criteria demand a certain response. For example, one situation in the Midkiff and Hansman simulator allowed pilots to overhear communications which suggested that another aircraft had not departed the runway the subjects were very close to landing on. In this case, action was required to avert a collision; a lack of action by the pilots could be considered to represent a lack of pilot situation awareness.

In addition, the situations should be chosen to cover the domain of important situations in which the system is expected to perform. For example, in the Midkiff and Hansman simulator study, the nine situations tested were the testable situations which had received the highest importance ratings in a pilot survey of Party Line Information importance. Testing of a final prototype system may include situations which test all conditions given in the system design specifications.

Finally, the situations must represent believable and recognizable occurrences to which the subject can be expected to react as they would in the real, non-simulated environment. For example, in the Midkiff and Hansman study, the subjects were flying an air transport simulator and believed they were over-hearing other air transport aircraft. Therefore, the 'Potential Collision' situations were staged to happen at a rate which was physically reasonable and were carefully scripted to portray to the subject a believable scenario of pilot confusion and/or mechanical failure on the part of the intruding aircraft.

The testable responses should be capable of examining the range of all probable actions and in-actions by the subject throughout the experiment. Care must be taken to look for actions which are different, less severe or incorrect in addition to just looking for the expected or desired result. For example, the response to the situation "Aircraft on Landing Runway" might be expected to be an immediate go-around. However, the subject's actions were often less severe, with pilots instead attempting to query ATC or each other to verify the knowledge they had gained from Party Line Information.

The strong reactions can be considered an indication of good situation awareness; correspondingly, the lack of any indication of awareness can be considered an indication of insufficient situation awareness. As discussed earlier in this paper, the uncertain or weak responses are also valuable measurements. They may illustrate problem areas such as lack of pilot confidence in information, feelings by the subjects that the expected reaction would defy accepted procedures, or other such unexpected impediments to action.

Performance-based measurement does not preclude other concurrent methods of assessing situation awareness. For example, Midkiff and Hansman also debriefed their subjects in an attempt to get pilot opinions on their situation awareness during the experiment.

GENERATING REPEATABLE SITUATIONS

When the purpose of an experiment is to test subjects' responses to specific situations involving multiple agents, there is a need to repeatably generate these situations across multiple trials. This is often complicated since subjects may not act consistently or as expected before the desired situation. As an illustration, consider the creation of an aircraft collision hazard. If the subject does not fly at exactly the speeds that were expected, the resulting situation can be completely different than that desired, or, as in this example as depicted in Figure 3, not occur at all.

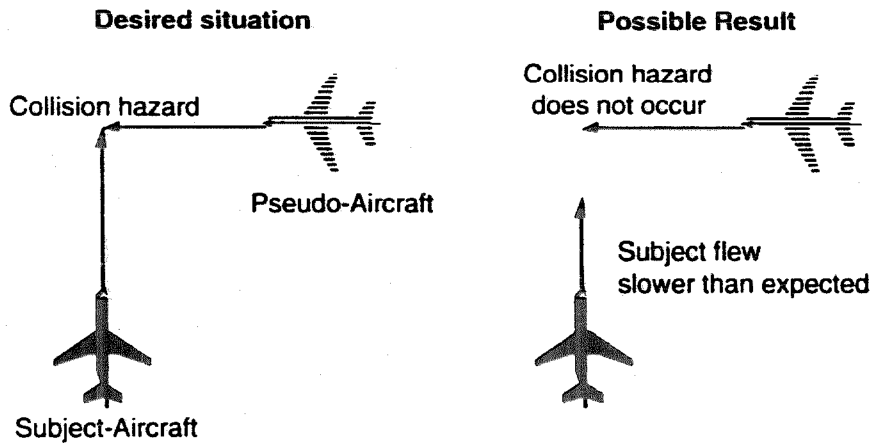

Figure 3. Situations are Dependent on Subject Actions

In order to make situations repeatable, some form of feedback of system state must be used to control the pseudo-agents (agents other than the subject), constantly controlling their actions to create the desired situations. Traditionally, this has been achieved by using experimenters to control pseudo-agents, in real-time, during the simulation run. A Robust Situation Generation architecture has been developed (Johnson & Hansman, 1995) whereby system state information is used to automatically generate scripted situations for a human subject, shown in Figure 4.

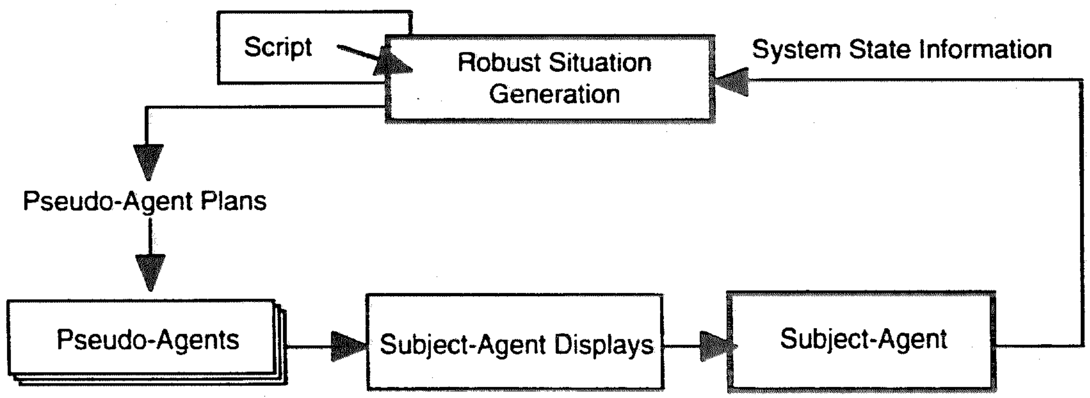

Figure 4. Overview of Robust Situation Generation

Pseudo-agents have plans that consist of a desired trajectory specified by waypoints and a discrete action plan. System state is utilized in three fundamental ways: pseudo-agent waypoints specified as relative to the subject, discrete actions of pseudo agents triggered by a cue, and cued amendments to pseudo-agent flight plans. Instances of these features are specified in a pre-determined script.

A Robust Situation Generation system has been implemented as part of the MIT Aeronautical Systems Laboratory (ASL) Advanced Cockpit Simulator (ACS), illustrated in Figure 5. A single workstation is used to simulate the pseudo-agents, consisting primarily of aircraft and controllers, and is referred to as the experimenter's station. Pseudo-aircraft state and digitally pre-recorded radio transmissions are presented to a subject operating the cockpit simulator. The scripts can be designed interactively in preliminary simulation runs using the experimenter's station, and are then stored and used as often as required.

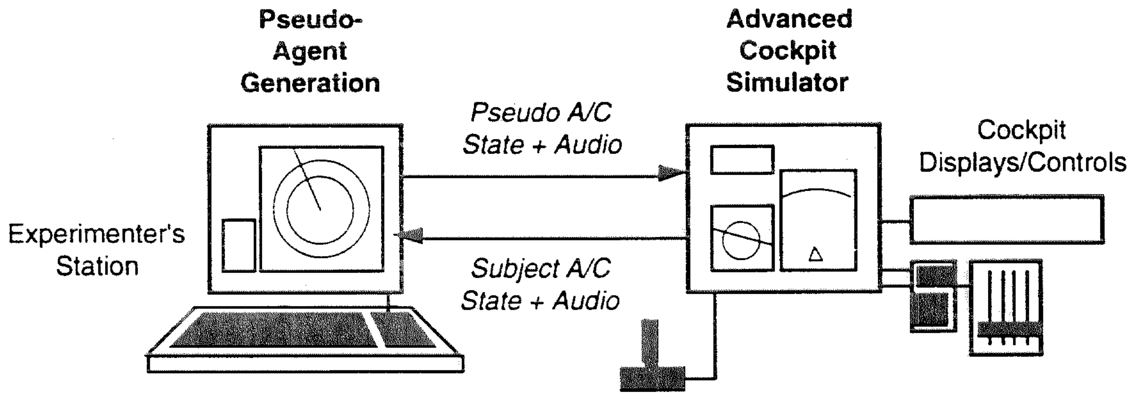

Figure 5. Implementation of Robust Situation Generation

The achieved robustness of the system, i.e. the maximum subject variation that can occur while still producing scripted situations, has been tested by varying subject-aircraft speed and position, as well as testing blunders by the subject, such as missing a turn. Unless the subject operates at an extreme limit of performance, situations were demonstrated to occur repeatably. The level of robustness depends on the level of fore thought and detail in the script, which can be made to react an arbitrary amount of subject variation, as required by the simulation.

CONCLUSIONS

Performance based measurement of situation awareness is a powerful tool for measuring the performance of a human-in-the-loop system and for identifying areas of inadequate situation awareness. The use of situations with testable responses can provide valuable insight into the user's situation awareness and how the user will act upon it.

The development of automatic robust situation generation has created a reliable mechanism for repeatable, consistent situations, making performance based measurement more reliable and easy to implement. Although the current implementation has been designed specifically for flight simulator experiments, Robust Situation Generation can also be implemented for any simulation involving multiple controllable agents.

REFERENCES

Adams, M.J. et al. (1995), "Situation Awareness and the Cognitive Management of Complex Systems," *Human Factors*, 37(1): 85-104.

Endsley, M.R. (1995), "Measurement of Situation Awareness in Dynamic Systems," *Human Factors*, 37(1): 65-84.

Johnson, E.N. & Pritchett, A.R. (1995), "Experimental Study of Vertical Flight Path Mode Awareness," Sixth IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design and Evaluation of Man-Machine Systems, Cambridge MA.

Johnson, E.N. & Hansman, R.J. (1995), "Multi-Agent Flight Simulation with Robust Situation Generation," MIT Aeronautical Systems Laboratory Report ASL-95-2.

Midkiff, A.H. & Hansman, R.J. (1993), "Identification of Important 'Party Line' Information Elements and Implications for Situational Awareness in the Datalink Environment," *Air Traffic Control Quarterly*, 1(1): 5-30.

Preprint, 14th AIAA/IEEE Digital Avionics Systems Conference, Nov. 6-9, 1995, Cambridge, MA

IMPACT OF VERTICAL SITUATION INFORMATION ON VERTICAL MODE AWARENESS IN ADVANCED AUTOFLIGHT SYSTEMS

Sanjay S. Vakil, R. John Hansman, Jr., Alan H. Midkiff

Aeronautical Systems Laboratory
Massachusetts Institute of Technology
Cambridge, MA 02139

ABSTRACT

An examination of autoflight systems in modern aircraft was made, with emphasis on the complex mode structure which is suspect in several recent accidents. Aviation Safety Reporting System reports and Flight Mode Annunciator conventions were examined. Principal results identified the lack of a consistent global model of the Autoflight System architecture and identified the vertical channel as requiring enhanced feedback. Functional requirements for an Electronic Vertical Situation Display (EVSD) were created based on established conventions and identified mode awareness problems. A preliminary version of this display was prototyped and an evaluation methdology was proposed. A set of experimental scenarios based on various types of mode awareness problems was established and discussed.

1.0 INTRODUCTION

Current advanced commercial transport aircraft rely on Autoflight Systems (AFS) for flight management, navigation and inner loop control. In the simplest form, these Autoflight Systems switch between different modes or states of aircraft automation. Each mode has a particular mechanism to control the state of the aircraft. These modes are strung together to create complicated flight trajectories. As the AFS becomes more complicated, it is speculated that the proliferation of modes and the variety of specific characteristics of each mode become more difficult to model mentally, leading to problems in mode awareness.

In flight mode awareness problems, there can be confusion between the pilots' expectations of the AFS and what it is actually doing. In this research, the working definition of a mode awareness problem is one in which the aircraft AFS executes an action, or fails to execute an action in the manner that was anticipated or expected by one or more of the pilots.

Mode awareness problems are suspected in a number of aircraft incidents, specifically vertical path control confusion on an Airbus A320 at Strausbourg (Sparaco, 1994) and pitch control in go-around mode on an A300 at Nagoya (Mecham, 1994). Additional incidents include events such as overspeed during takeoff or go-around on Boeing 757/767s and inconsistent envelope protection (loss of the alpha protection floor) on the A320.

2.0 PREVIOUS WORK

Several mechanisms have been used to explore the issues surrounding mode awareness problems. These included a review of Aviation Safety Reporting System (ASRS) reports, an investigation of the structure of current AFSs, including the Flight Mode Annunciator, as well as flight observations and focussed interviews with pilots and check airmen.

Two main observations emerged from the investigation. First the results indicated that vertical modes appear to be more prone to serious mode related problems. A second problem was a lack of underlying structure to the automation, making it difficult for pilots to develop consistent mental models.

The three hundred ASRS reports (from 1990-94) returned by the keyword search were reduced to 184 relevent cases (Vakil, et al, 1995). A total of 62.7% of the categorized reports were vertical/speed related. In particular, the Mode Transition Problems causal category is dominated by vertical/speed problems.

The FMA investigation highlighted that more recently designed FMAs display Target State Values, for example, the commanded vertical speed in Vertical Speed Mode. The Fokker 100 and the MD-11 FMAs also displayed the *control allocation* in each mode. This information explicitly identifies which actuator is controlling which output state. This is particularly relevent in the vertical domain since pitch and thrust can be used interchangeably to control the vertical path and speed. Note that the availability of control allocation information in these FMAs implies that a set of parallel SISO (Single Input, Single Output) controllers is being used. This SISO model was found to be consistent with pilot mental models.

In contrast, two distinct types of AFS input-output relationships can exist. The first is simply a quasi-steady-state model where each output variable is controlled by a single input (SISO). This relationship appears to be functionally adequate for simple (base-)modes. More complex (macro-)modes and certain mode transitions require Multi-Input Multi-Output (MIMO) controllers, where each output variable can be controlled by more than one input. These modes are typically of short duration and do not appear to be modelled in detail by the crew.

Based on interviews and observations, pilots appear to model the AFS modes as independent SISO control loops. They do not appear to model MIMO transitions in detail, relying instead on an understanding of the final target criteria and some smoothness criteria to monitor the AFS performance.

3.0 INVESTIGATION OF FEEDBACK MECHANISMS

In order to improve vertical mode awareness, a set of crew information requirements was developed to help mitigate some identified Mode Awareness Problems. An Electronic Vertical Situation display was designed based on these requirements. The display is currently being prototyped on the Aeronautical Systems Laboratory's Advanced Part Task Simulator.

3.1 Information Requirements for Enhanced Vertical Mode Feedback

Information requirements were based on operator requirements and on known mode awareness problems. As shown in Table 1, the information requirements have four components: information regarding the current mode, the anticipated modes, the mode transitions and the consequences based on no further inputs.

For the *current mode*, the mode identification and the specific attributes of the mode, namely target state values and control allocation, are displayed in text. Where possible, graphical depiction is given of the current state and targets.

Anticipated modes have the same set of information as the current mode when this information can be determined. While this determination is straightforward for many preprogrammed macro-modes, it may be inaccurate for uncommanded mode changes. For example, the prediction of entering an envelope protection mode downstream would be difficult.

Mode transition alerting requires anticipation of mode transitions. The depiction of envelope protection limits alerts the pilot when an envelope protection mode change is imminent. In certain situations, it may not be possible to predict transitions. Retroactive cueing may be necessary to show the mode that was previously engaged. This is important in situations where an uncommanded mode change creates a confusing set of transitions.

Table 1. Crew Information Requirements

Current Mode
- Mode Identification
- Target States
- Control Allocation

Anticipated Mode*
- Mode Identification
- Target States
- Control Allocation

Mode Transitions*
- Envelope Limits
- Retroactive Cueing

Consequences*

*if predictable

The *consequences* of a mode change are based on an extrapolation of the current state of the aircraft automation. A predictive profile based on the current automation state is shown to display the locations of anticipated mode transitions and their consequences on the flight path.

3.2 Prototype Electronic Vertical Situation Display

It is hypothesized that an Electronic Vertical Situation Display (EVSD) which incorporates the information elements discussed earlier will provide enhanced vertical mode feedback. The use of an EVSD has been suggested by other researchers for differing reasons: to improve mode awareness by Hutchins (1995) and Palmer (1995) and to accurately control in the vertical domain by Fadden, Braune, and Wiedemann (1988). A simple EVSD is currently flying on the GulfStream Corporation's G4 cockpit.

An EVSD is envisioned to provide an analog to the Electronic Horizontal Situation Indicator (EHSI), or "map" display currently available in glass cockpits. The display shows the programmed vertical path of the aircraft and the associated modes referenced to that path. The functional requirements for a prototypical EVSD were based on a functional model of the AFS and were to be consistent with pilots' mental models and the control allocation displays on newer aircraft. Basing the requirements on a SISO model maintains this consistency.

A preliminary prototype of the EVSD has been implemented on the Aeronautical Systems Laboratory's Part Task simulator. An example of this display is shown in Figure 1. This prototype has four major areas. At the top of the display is the mode display window, showing the current and anticipated modes, control allocations and target states. At the left is a scaleable altitude tape. The bottom window can either display the path distance (if in a lateral navigation mode), or the range directly ahead of the aircraft. Finally, the main window shows the aircraft in vertical relation to the upcoming waypoints and mode transition points. It should be noted that because of the prototypical nature of this EVSD, certain informational requirements (retroactive cueing, some graphical elements) have yet to be implemented.

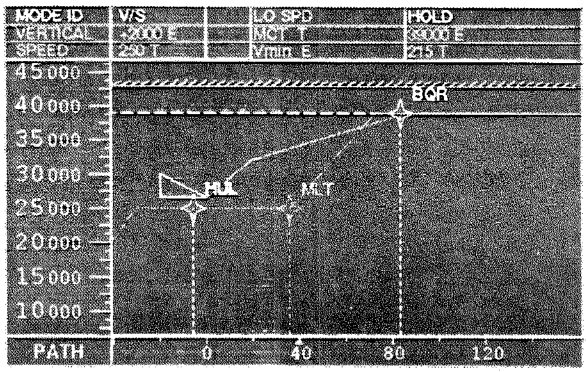

Figure 1. Electronic Vertical Situation Display during Vertical Speed Climb, reaching an envelope protection limit

This display is implemented in color conventions consistent with the Boeing 747-400 Electronic Horizontal Situation Indicator (EHSI) as shown in Table 2. The *current mode* is identified in this prototype on the top window, in white directly above the aircraft symbol. Underneath are the control allocations for the mode and the target states. In this example, the aircraft is in Vertical Speed (VS) Mode with the vertical path controlled by elevator (E) and the speed controlled by throttle (T). The *anticipated mode* is shown in the top window above the point where it is predicted to be engaged. The anticipated target state and control allocations are depicted in a manner similar to the current mode. In this particular case, the system is predicting a envelope protection violation and a mode transition to the LO SPD protect mode. In this mode, the vertical path is controlled by the throttles (at Maximum Climb Thrust), and the speed is controlled by elevators to Vmin. Note that both the target states and the control allocation changes when the new mode is engaged. *Mode transition alerting* currently consists of the anticipated mode translating across the top window and into the current mode slot when engaged. In this case, an elbow in the path also highlights where the transition is calculated to occur. An example of the *target state* is the dashed magenta line at 39000ft, which is the altitude currently dialed into the MCP. The target state is also shown in the mode identification window, as a vertical path at 39000ft at the transition to Altitude Hold (HOLD) mode. The Aircraft Path Line shows the path that the aircraft will travel based on the current state of the automation.

Table 2. Color Conventions

Information Element	Graphical Representation
MCP Input	Dashed WhiteLine
FMS Path	Magenta Line
Programmed Waypoint	White Symbol
Current Waypoint	Magenta Symbol
Target Altitude Intercept Location	Green Triangle
Top of Descent	Green Circle
Aircraft Path Line	Green Line

The solid magenta line connecting waypoint crossing restrictions is a graphical display of the current vertical path programmed into the FMS. In the main section of the screen are the lateral waypoints programmed into the FMS. Altitude crossings are shown by the waypoint altitudes. Those without restrictions are placed on the ground. Dashed lines descend to the path scale on the bottom of the display. The path distance between each waypoint defines the horizontal scale.

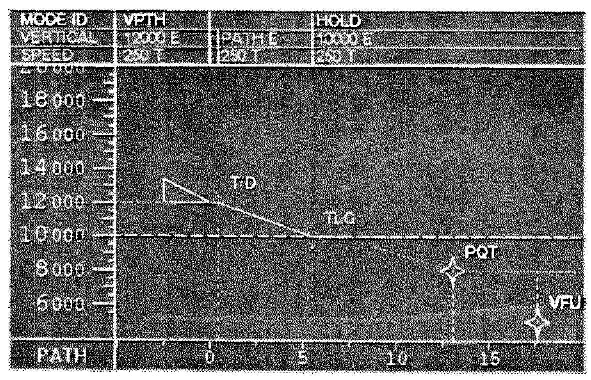

Figure 2. Electronic Vertical Situation Display during VNAV descent

The example shown in Figure 2 is a Vertical Navigation (VNAV) macro-mode descent. While still in VNAV macro-mode, note that the aircraft changes the target state of the vertical path at the Top of Descent (T/D) point: a predetermined transition criteria. In the mode identification window is the current target state, an altitude of 12000 ft. At the anticipated Top of Descent point the target state becomes a path-based descent. In addition, the Aircraft Path Line highlights the mode change that will occur at 10000ft as the aircraft intersects the preprogrammed MCP altitude and goes into an Altitude Hold mode. Along the bottom of the display is the altitude above sea level of the terrain under the aircraft flight path. Also, the instantaneous value of the terrain height is shown on the bottom of the altitude tape. The terrain information is drawn from a database.

4.0 EVALUATION METHODOLOGY

In order to evaluate the effect of the prototype EVSD on mode awareness, a series of mini-scenarios has been developed based on a sampling of the common mode awareness problems that were found in the ASRS review.

The evaluation methodology is very similar to the SAGAT techniques (Endsley, 1995) but has some important differences. In contrast with SAGAT evaluations, where the subjects are interacting with the simulator, the scenarios are completely passive. The only interaction is for the pilots to signal the computer when they have observed a specified event. The other major difference is that the questions that are asked are not selectively drawn from a larger pool of available questions, rather they are static, generalized questions applicable to each scenario.

Subjects, who are current airline pilots, will act as Pilot Not Flying (PNF), passively observing the operation of the Autoflight System. A compatriot researcher will act as the Pilot Flying (PF) and will respond to the Air Traffic Control cues that will occur at various times during the scenarios. The PF will handle all of the interaction with the Mode Control Panel in addition to controlling the scaling of the display to ensure that the relevant information is always available on the displays.

The dependent variable in the experiment will be the presence of the EVSD, so each scenario will be run at least twice: with and without the EVSD. To "keep pilots honest", some scenarios which do not have any uncommanded mode changes will also be run.

When the subject feels that an error that an uncommanded mode transition has occurred, or that an error has been made in interfacing with the FMS, or an unsafe or nonprocedural operation has taken place, they will signal the computer. The simulation will be paused and the screen will be blanked. The subject will be queried with a series of questions designed to determine the level of mode and situational awareness. Questions address the cause of the mode transition, the consequences of the new mode state and the cues used by the subject to determine that an incorrect automation state was reached:

1. Why did you stop the scenario?
2. What caused the event? Why did it happen?
3. What cues did you use to determine that the event occurred?
4. What would have eventually happened if you had not stopped the scenario?

The quantitative measurements include the elapsed time from the uncommanded mode transition or other incident (PF error, aircraft put in danger) and, when relevant, any deviation from the nominal flight path.

5.0 SCENARIO OVERVIEWS

The following is a brief description of each of the candidate scenarios that will to be tested in the part task facility.

Loss of glide slope: During a standard ILS approach, loss of signal due to hardware glitch causes a mode reversion to the instantaneous vertical speed. The aircraft transitions into Vertical Speed mode and the flight path trajectory intersects the ground short of the runway threshold. EVSD-specific cues include a change in the mode header and the green path prediction line indicating the dangerous flight path.

High Speed Protect: During a steep descent in vertical speed mode, the envelope protection limits of the aircraft are surpassed. The aircraft transitions to High Speed Protect mode; AFS control allocation shifts from speed-on-thrust to speed-on-pitch. EVSD-specific cues include mode header change and green path prediction line indicating altered flight trajectory. Note that the EVSD predicts this mode change before it occurs, allowing the pilot to anticipate the change.

Loss of Altitude Capture: During a Flight Level Change Descent, the aircraft fails to level-off at the armed/pre-selected MCP altitude. Instead, the aircraft continues an open descent at some small interim descent value in Vertical Speed Mode. Since the MCP altitude is now above the aircraft, the vehicle is in an open descent. EVSD-specific cues include mode header change and green path prediction line indicating the mild descent. This glitch has been documented on some early MD-80 aircraft.

FMS STAR: During an FMS STAR procedure, unanticipated wind conditions result in the inability of the AFS to accommodate both the FMS path and speed targets, causing a VNAV PTH to VNAV SPD mode transition. The consequences include non-compliance with STAR altitude and speed restrictions. EVSD-specific cues include mode header change and green path prediction line indicating change of descent trajectory and the point of MCP altitude intercept, downstream of the next waypoint.

Non-Precision Approach: During a non-precision approach with multiple steps, an intermediate altitude target between the Final Approach Fix and the Missed Approach Point is inadvertently omitted by the Pilot Flying. EVSD-specific cues include inappropriate altitude target in mode header and green path prediction line indicating neglected intermediate step.

Climb Crossing Restriction: A climb crossing restriction is issued by Air Traffic Control. A few minutes later an amendment is issued. This amendment results in a target that exceeds the climb performance capability of the aircraft, resulting in a transition to LO SPD Protect mode. EVSD-specific cues include mode header change and green path prediction line indicating that the aircraft will not make the restriction.

Terrain Depiction: The EVSD terrain depiction is demonstrated in an emulation of Controlled Flight Into Terrain (CFIT) situation in which an incorrect flight mode caused an incorrect Vertical Speed to be selected. EVSD-specific cues include an inappropriate Vertical Speed target in mode header and green path prediction line indicating (CFIT) flight path trajectory.

6.0 DESCRIPTION OF SELECTED SCENARIOS

Shown below are screen snapshots from two of the cases to demonstrate current scenario implementation of the EVSD.

6.1 High Speed Protect Scenario

The scenario begins with the aircraft in Altitude Hold mode, travelling towards MLT at an altitude of 35000ft as shown in Figure 3. Air Traffic Control then issues a command to descend and maintain 20000ft by MLT. This is possible by dialing in a rather steep Vertical Speed descent of 6000fpm. Some time later, the AutoFlight System realizes that the current aircraft configuration is going to lead to an envelope protection limit violation, and anticipates where this overspeed will occur.

When the HI SPD mode is engaged, the throttles latch to idle and the speed is controlled by the elevators in a manner designed not to exceed Vmax. At this point, both the standard ADI mode windows and the MCP will show that a mode transition has occurred, and the pilot should be prompted to register the transition by signalling the computer.

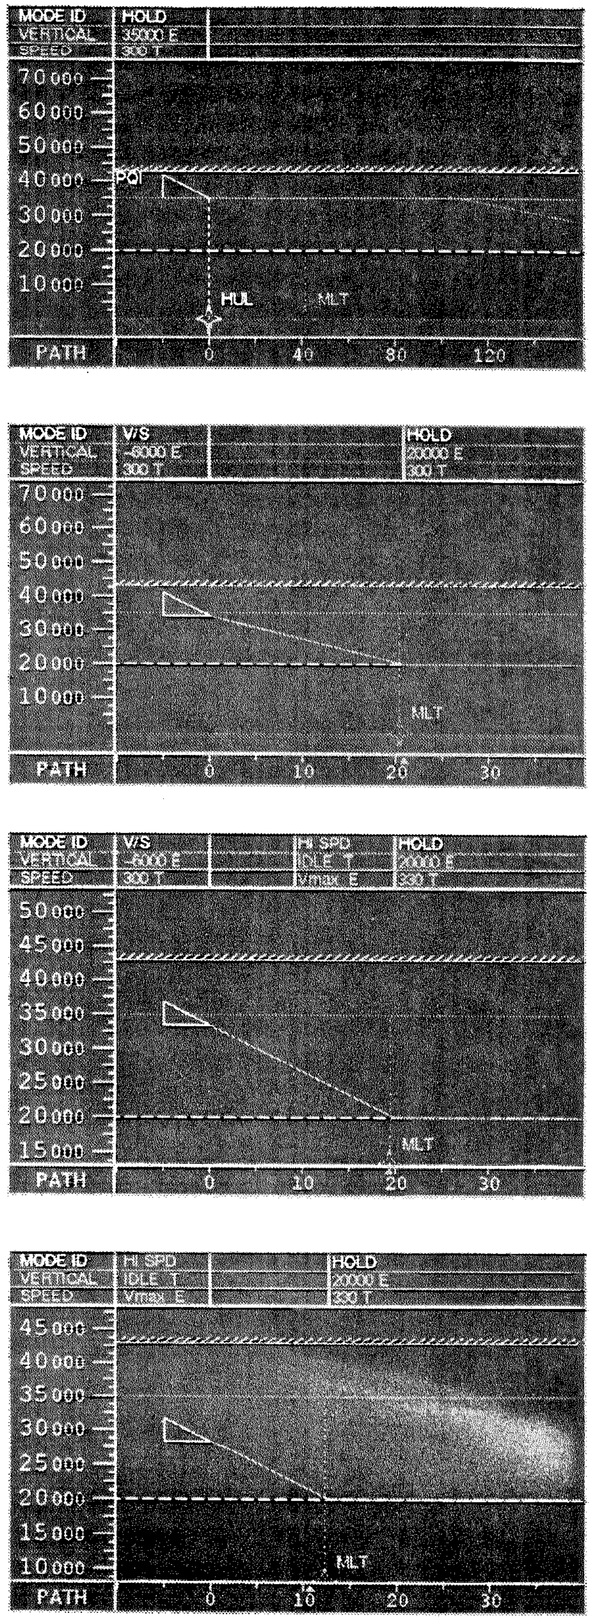

Figure 3. High Speed Protect Scenario

6.2 Loss of Glide Slope Signal Scenario

As shown in Figure 4, the aircraft starts out in an armed mode, with the localizer and glide slope capture anticipated by the AFS. A few minutes later, both the localizer and the glide slope have engaged and the aircraft begins a descent down the 3° glide path.

When the loss of glide slope signal occurs, the AFS in most current aircraft respond by setting the mode to the instantaneous Vertical Speed that was being used to bracket the glide slope signal. Since this loss of signal occurred early in the descent, the controller logic was oscillating about the correct vertical speed and called for an overly steep descent, in this case, -900fpm. If this mode error remains undetected, the aircraft will impact short of the runway.

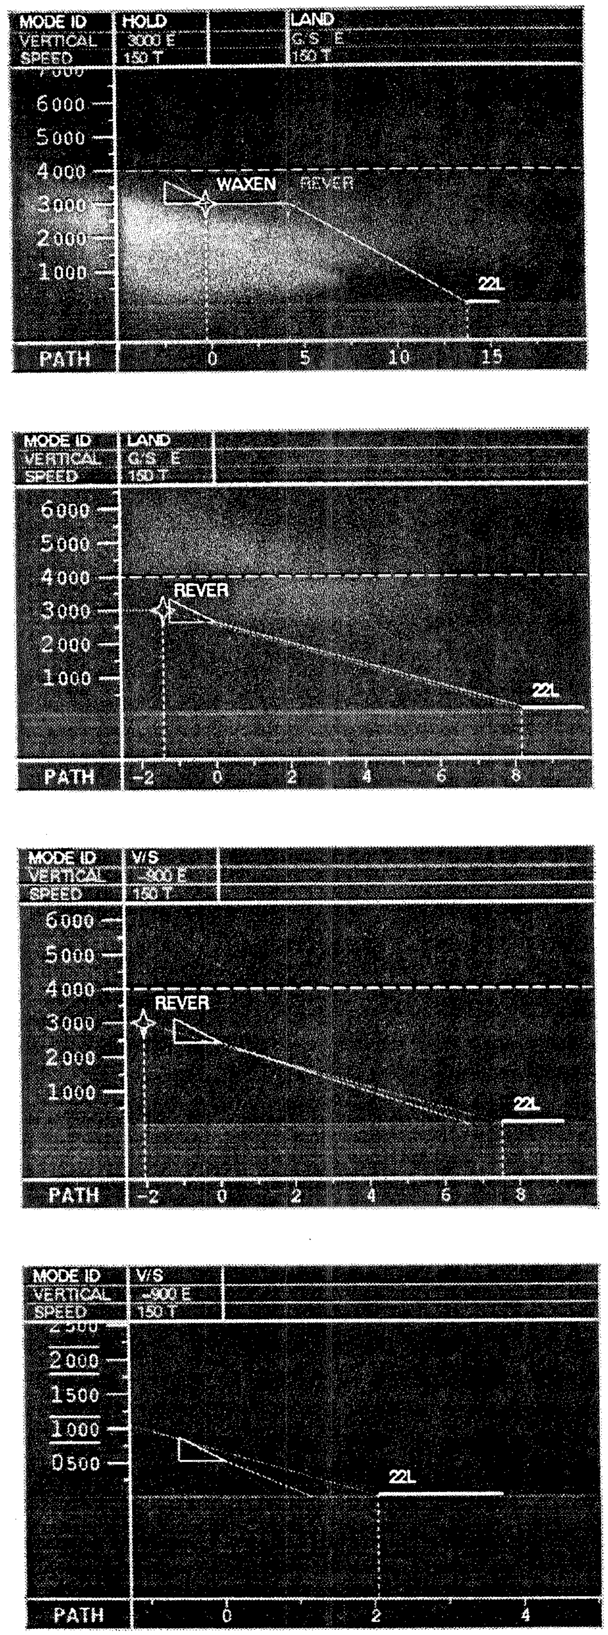

Figure 4. Loss of Glide Slope

ACKNOWLEDGEMENTS

This work was supported by the National Aeronautics and Space Administration under grant NAG1-1581. The authors would like to thank the following individuals for their suggestions and contributions: William Corwin, Honeywell; Peter Polson, Jim Irving, Sharon Irving, University of Colorado; Michael Palmer, Kathy Abbot, Terrence Abbot, Everett Palmer, NASA.

REFERENCES

Mecham, M. (1994). Autopilot Go-Around Key to CAL Crash. *Aviation Week & Space Technology*. May 9, 1994.

Sparaco, P. (1994). Human Factors Cited in French A320 Crash. *Aviation Week & Space Technology*. January 3, 1994.

Endsley, M. R. (1995). Measurement of Situation Awareness in Dynamic Systems. *Human Factors*. 37(1).

Fadden, Braune, and Wiedemann, 1988. *Spatial Displays as a Means to Increase Pilot Situational Awareness*, Boeing Commercial Aircraft Company, Seattle, Washington.

Hughes, D. (1995), Incidents Reveal Mode Confusion. *Aviation Week & Space Technology*. January 30, 1995.

Vakil, Hansman, Midkiff, Vaneck, (1995). Mode Awareness in Advanced AutoFlight Systems. *Proceedings of the 6th IFAC Symposium*. June 27-29, 1995

Preprint. 6th IFAC/IFIP/IFORS/IEA Symposium, June 27-29, 1995, Cambridge, MA

MODE AWARENESS IN ADVANCED AUTOFLIGHT SYSTEMS

Sanjay S. Vakil
R. John Hansman, Jr.
Alan H. Midkiff
Thomas Vaneck

Aeronautical Systems Laboratory,
Massachusetts Institute of Technology
Cambridge, MA 02139

Abstract: An examination of autoflight systems in modern aircraft was made, with empasis on the complex mode structure which is suspect in several recent accidents. Aviation Safety Reporting System reports from 1990-94 were examined. Flight Mode Annunciator conventions were inventoried. Focussed interviews with pilots and check airmen were conducted. Principal results identified the lack of a consistent global model of the AutoFlight System architecture and identified the vertical channel as requiring enhanced feedback. An Electronic Vertical Situation Display to help mitigate the identified problems was prototyped.

Keywords: Aircraft Contol, Mode Structure, Complex Systems, Displays, Human Supervisory Control, Human Factors, Human-Machine Interface, Mode Analysis

1.0 INTRODUCTION

Current advanced commercial transport aircraft rely on AutoFlight Systems (AFS) for flight management, navigation and inner loop control. These systems have evolved into multiple computers which are capable of sophisticated tasks including automatic flight planning, navigation, and automated landings. These systems also provide envelope protection, which prevents pilots from committing obvious mistakes, such as stalling or lowering flaps at a high speed.

In the simplest form, these AutoFlight Systems switch between different modes (states of aircraft automation), stringing together modes to create complicated flight trajectories. As the AFS becomes more complicated, it is speculated that the proliferation of modes and the specific characteristics of each mode become more difficult to model mentally, leading to problems in mode awareness.

1.1 AutoFlight System Overview

The AFS in a modern aircraft typically separates the guidance into uncoupled horizontal and vertical components. Examples of horizontal guidance modes include flying a preprogrammed trajectory using a Lateral Navigation mode (LNAV), or flying on a selected heading. For horizontal flight, the path is controlled through roll. Typical vertical guidance modes include flying level (Altitude Hold), or maintaining a selected Vertical Speed (V/S). The vertical flight path coupled with the speed of the aircraft is controlled by a combination of thrust and elevator.

Unfortunately, the increasing complexity of AutoFlight Systems has caused an increase in problems associated with the management of the system. In flight mode awareness problems, there can be confusion between the pilots' expectations of the AFS and what it is actually doing. In this research, the working definition of a mode awareness problem is one in which the aircraft AFS executes an action, or fails to execute an action that is anticipated or expected by one or more of the pilots.

1.2 Motivation

Mode awareness problems are suspected in a number of aircraft incidents, specifically vertical path control confusion on an Airbus A320 at Strausbourg (Sparaco, 1994) and pitch control in go-around mode on an A300 at Nagoya (Mecham, 1994). Additional incidents include events such as overspeed during takeoff or go-around on Boeing 757/767s and inconsistent envelope protection (loss of the alpha protection floor) on the A320.

1.3 Approach

The approach involved two parts: an investigation to identify the problems and issues involved in mode awareness, and an investigation of potential mechanisms to mitigate mode awareness problems, including enhanced mode feedback

2.0 ISSUE IDENTIFICATION METHODOLOGY

Several mechanisms were used to explore the issues surrounding mode awareness problems. These included a review of Aviation Safety Reporting System (ASRS) reports, an investigation of the structure of current AFSs, including the Flight Mode Annunciator, as well as flight observations and focussed interviews.

2.1 Aviation Safety Reporting System

The Aviation Safety Reporting System allows pilots to detail safety problems or incidents with a degree of amnesty. A search was performed on the ASRS database over the years 1990-94 with a set of keywords designed to elicit problems related to mode awareness. The keywords consisted of the following: annunciation, annunciator, FMC, flight management computer, FMS, flight management system, CDU, mode, capture, arm, automatic flight system, vertical, horizontal, and program.

2.2 Inventory of Current AutoFlight Systems

The structure of the AutoFlight Systems in current aircraft was examined, including an inventory of flight mode annunciation schemes, via training manuals, and a review of the open literature. The aircraft examined included the Boeing 757/767, the MD-11 and the Airbus 300-600R

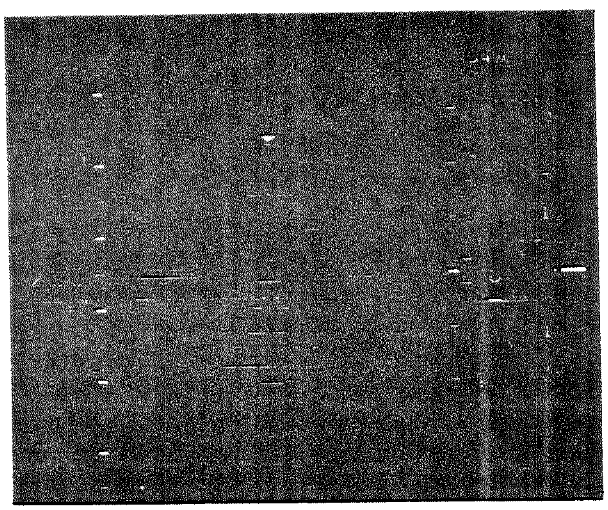

Figure 1. Primary Flight Display on 747-400

The Flight Mode Annunciator (FMA) on the Primary Flight Display (PFD) of modern aircraft is normally the primary location of mode status information. FMAs typically display the current mode configuration of the aircraft in text. A PFD from the Boeing 747-400 is shown in Figure 1. The FMA is located in the middle of the display directly above the artificial horizon. The FMA shows the aircraft's AutoThrottle is engaged, that it is in LNAV Mode, and in Altitude Hold Mode. The display conventions used in the FMA were compared across a set of aircraft, including the Boeing 737-500/600, 757, and 767, MD-11, MD-80, Fokker 100 and Airbus 300-600R.

2.3 Focussed Interviews

Information was gathered from a variety of informal sources, including direct flight deck observations, discussions with flight crews, and discussions with simulator check airmen about their observations of crews during recurrent training.

3.0 PRELIMINARY OBSERVATIONS

Two main observations emerged from the investigation. First the results indicated that vertical modes appear to be more prone to serious mode related problems. Second was a lack of underlying structure to the automation, making it difficult for pilots to develop consistent mental models. Tertiary observations included a lack of commonality between Flight Mode Annunciators and a reduction in vestibular cueing.

3.1 Vertical Flight Path

Three hundred ASRS reports (from 1990-94) were returned by the keyword search and analyzed. Of these 184 were categorized as mode awareness problems using the working definition detailed earlier.

As shown in Figure 2, these reports were then categorized by the perceived cause of the problem and by the flight path (vertical/speed, horizontal or both) that was impacted. Since the vertical flight path and the speed are implicitly coupled, problems with either were grouped together. In instances where the problems spanned multiple causal categories, the reports were counted in each relevent causal category.

In Figure 2, it can be seen that vertical/speed problems dominate many of the categories. A total of 62.7% of the categorized reports were vertical/speed related. In particular, the Mode Transition Problems causal category is dominated by vertical/speed

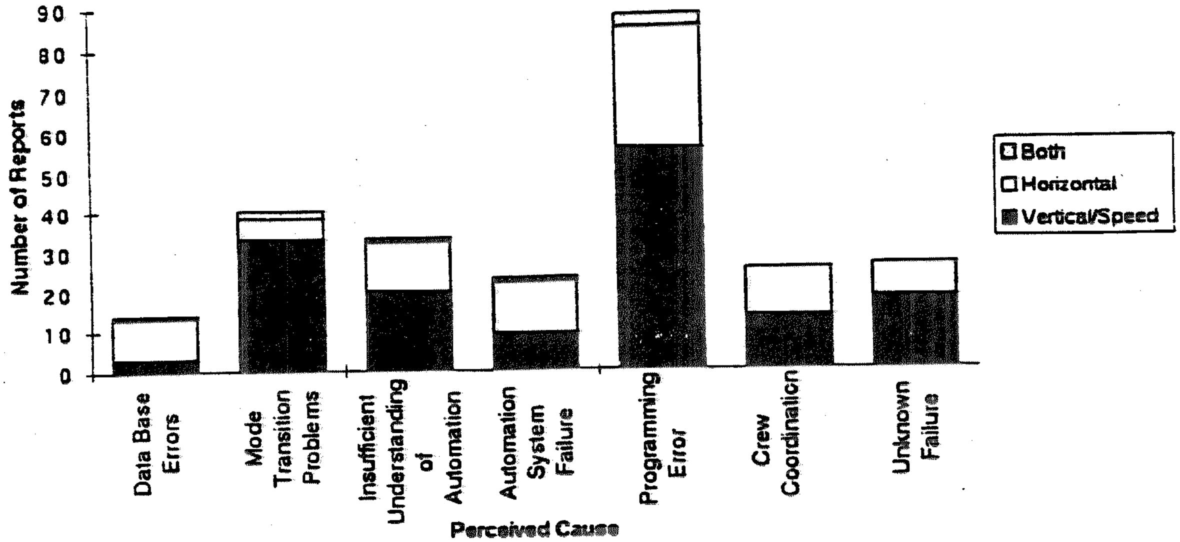

Figure 2. Breakdown of ASRS reports into Perceived Causes and Flight Domain.

problems. The data in the Insufficient Understanding of Automation causal category also suggests a deficiency in knowledge of the vertical domain automation.

It should be noted that there exists a potential for over-reporting vertical deviations. Since Air Traffic Control radar has very precise surveillance of altitude versus position, there may be more cause to report vertical/speed incidents due to the amnesty clause in the ASRS.

3.2 *Structure of Current Autoflight Systems*

AutoFlight Systems have evolved from more basic autopilots as manufacturers incorporated additional functionality, including fuel optimal descents, flight path angle controlled vertical maneuvers and speed based vertical maneuvers. The evolutionary development of these systems and their increased capabilities have resulted in a complex structure of operating modes. Since each generation of AutoFlight System incorporates much of the functionality and modes of the previous generation, without reducing the number of features, there has been an entropic growth of complexity.

The other effect of this entropic growth has been to create a system that appears to lack a simple, consistent, global model. This lack requires pilots to create their own ad-hoc models.

Mode Definitions. A mode can be defined as a specific state of the aircraft automation. A mode includes the set of targets (heading, speed, vertical speed, pitch etc.) to which the aircraft is to be controlled as well as the actuators that are to be used (elevators, thrust, ailerons). Current AutoFlight Systems switch between two basic types of operating modes which are termed base-modes and macro-modes. Base-modes are used in quasi-steady-state conditions and have an invariant set of targets. A base-mode example would be a Vertical Speed Mode where the aircraft attempts to maintain a specific vertical speed target by controlling pitch (via the elevators) and airspeed with thrust.

Macro-modes consist of a linked sequence of base-modes. Each base-mode in the macro-mode has its own set of targets, implying a set of targets which vary over the course of the macro-mode. Transitions between the base-modes are made based on specific state criteria, such as altitude or indicated air speed. An example of a macro-mode is the Autoland sequence, which transistions (in the vertical channel) between Flight Level Change, Glideslope Capture, Flare, and Rollout with a different set of targets in each base-mode.

Mode Transitions. There are three types of transitions between modes. A *commanded transition* is active as soon as selection is made. An example is pressing the vertical speed button. An *uncommanded transition* is one that is not directly activated by the pilot. These transitions are usually some type of envelope protection. An example is a transition caused by overspeed protection. Finally, *automatic/conditional transitions* occur when, after arming, a mode engagement occurs at a target state. An example is the use of Glide Slope Capture to transition to a descent mode after the aircraft intersects with the glide slope trajectory.

AFS Input-Output Relationships. Two basic types of AFS input-output relationships can exist. The simpler is a quasi-steady-state model where each output variable is controlled by a single input (SISO). A typical Vertical Speed mode engages two independent SISO controllers: the aircraft's pitch controls the vertical speed target and the thrust controls the air speed target. SISO models appear to be functionally adequate for simple modes.

Some mode transitions appear to utilize Multi-Input Multi-Output (MIMO) controllers, where each output variable can be controlled by more than one input. These transitions are typically of short duration and they do not appear to be modelled in detail by the crew. An example of a complex mode transition is the 0.05g capture used in an Altitude Capture transition of the MD-11: when the aircraft is approaching a selected flight altitude, the intercept maneuver limits the normal acceleration to 0.05g.

3.3 Flight Mode Annunciators

Preliminary analysis of Current Flight Mode Annunciators showed a limited commonality between displays. There are differences in the number, type and location of windows, in the usage of color, in the layout of text, and in the conventions used to indicate arming. More recently designed FMAs include Target State Values. An example of a target state value would be the commanded vertical speed in Vertical Speed Mode.

The Fokker 100 and the MD-11 FMAs also displayed the *control allocation* in each mode. This information explicitly identifies which actuator is controlling a particular output. Knowledge of control allocation is particularly important in the vertical domain, since pitch and thrust can be used interchangeably to control the vertical path and speed. Note that the availability of control allocation in these FMAs implies a set of parallel SISO (Single Input, Single Output) controllers being used.

3.4 Pilot models of AutoFlight Systems

Based on the analysis of the structure of current AutoFlight Systems, there does not appear to be a simple, consistent, global model of current AFS. Such a model is not available in flight manuals, which focus on crew interface and procedures. In the absence of a simple consistent model, pilots appear develop their own ad hoc models of the AFS.

The observation that pilots are constructing their own ad hoc models in the absence being presented with one has serious implications. Any empirically derived model of the AFS is going to be largely based on nominal operation. Without explicit knowledge of the system structure during non-nominal situations, pilots may find their models inadequate when they are most critical. The implication for the training of new pilots is that a consistent model of the automation would provide a solid basis upon which to understand aircraft automation.

Based on interviews and observations, pilot's appear to model the AFS modes as independent SISO control loops. They do not appear to model MIMO transtions in detail, relying instead on an understanding of the final target criteria and some smoothness criteria to monitor the AFS performance.

Pilot models were observed to be of a SISO variety. As such, they may not accurately represent AFS operation, especially in transitions between modes. Transitions to and from envelope protection modes appear to be particularly troublesome.

Since these models are constructed empirically, individual pilot models may vary significantly. In future aircraft, this empirical model construction may cause additional difficulties, especially as control designers migrate to more common MIMO control designs.

3.5 Vestibular Cueing Concerns

An additional factor in mode awareness is the reduction of vestibular cues of mode transitions in some aircraft. For the purpose of improving ride quality, some AutoFlight Systems impose nominal vertical acceleration limits close to the human vestibular detection threshold. This provides a smooth ride for passengers, but may reduce mode awareness.

The vestibular thresholds of blindfolded subject in the Z body-axis is between 0.1Hz to 1Hz measured under laboratory conditions is between 0.001G and 0.015G (Gundry, 1978). The nominal goal of the control system on the Airbus A320 and the McDonald-Douglas MD-11 during nominal maneuvering is thought to be 0.05G.

Under operational conditions, 0.05G may be below the detection threshold of individual pilots. More work on human detection capabilities in operating environments is required to address this issue.

4.0 INVESTIGATION OF FEEDBACK MECHANISMS

In order to improve vertical mode awareness, a set of crew information requirements was developed to help mitigate some identified Mode Awareness Problems. An Electronic Vertical Situation display was designed based on these requirements. The display is currently being prototyped on the Aeronautical Systems Laboratory's Advanced Part Task Simulator.

4.1 Information Requirements for Enhanced Vertical Mode Feedback

Information requirements were based on operator requirements and on known mode awareness problems.

As shown in Table 1, the information requirements have four components: information regarding the current mode, the anticipated modes, the mode transitions and the consequences based on no further inputs.

Table 1. Crew Information Requirements

Current Mode
Mode Identification
Target States
Control Allocation
Anticipated Mode*
Mode Identification
Target States
Control Allocation
Mode Transitions*
Envelope Limits
Retroactive Cueing
Consequences*

*if predictable

For the *current mode*, the mode identification and the specific attributes of the mode, namely target state values and control allocation, are displayed in text. Where possible, graphical depiction is given of the current state and targets.

Anticipated modes have the same set of information as the current mode when this information can be determined. While this determination is straightforward for many preprogrammed macro-modes, it may be inaccurate for uncommanded mode changes. For example, the prediction of entering an envelope protection mode far in the future would be difficult.

Mode transition alerting requires anticipation of mode transitions. The depiction of envelope protection limits alerts the pilot when envelope protection is not available. In certain situations, it may not be possible to predict transitions. Retroactive cueing may be necessary to show the mode that was previously engaged. This is important in situations where an uncommanded mode change creates a confusing set of transitions.

The *consequences* of a mode change are based on an extrapolation of the current state of the aircraft automation. A predictive profile based on the current automation state is shown to display the locations of anticipated mode transitions and their consequences on the flight path.

4.2 Prototype Electronic Vertical Situation Display

It is hypothesized that an Electronic Vertical Situation Display (EVSD) which incoporates the information elements discussed earlier will provide enhanced vertical mode feedback. This conclusion has been drawn by other researchers both to improve mode awareness by Hutchins and Palmer and to accurately control in the vertical domain by Fadden, Braune, and Wiedemann. A simple EVSD is currently flying on the GulfStream Corporation's G4 cockpit.

An EVSD is envisioned to provide an analog to the Electronic Horizontal Situation Indicator (EHSI), or "map" display currently available in glass cockpits. The display shows the programmed vertical path of the aircraft and the associated modes referenced to that path. The functional requirements for a prototypical EVSD were based on a functional model of the AFS and were to be consistent with pilots' mental models and the control allocation displays on newer aircraft. Basing the requirements on a SISO model maintains this consistency.

A preliminary prototype of the EVSD has been implemented on the Aeronautical Systems Laboratory's Part Task simulator. An example of this display is shown in Figure 3. This prototype has four major areas. At the top of the display is the mode display window, showing the current and anticipated modes, control allocations and target states. At the left is a scaleable altitude tape. The bottom window can either display the path distance (if in a lateral navigation mode), or the range directly ahead of the aircraft. Finally, the main window shows the aircraft in vertical relation to the upcoming waypoints and mode transition points. It should be noted that because of the prototypical nature of this EVSD, certain informational requirements (retroactive cueing, predictive profile, some graphical elements) have yet to be implemented.

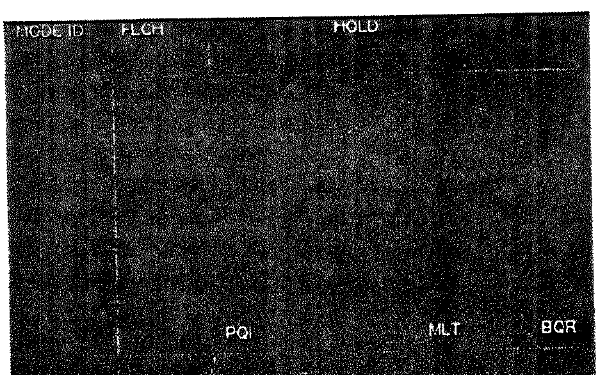

Figure 3. Electronic Vertical Situation Display during Flight Level Change Climb

This display is implemented in color conventions consistent with the 747-400 Electronic Horizontal Situation Indicator (EHSI) as shown in Table 2.

Table 2. Color Conventions

Information Element	Graphical Representation
MCP) Input	Dashed Magenta Line
FMS Path	Magenta Line
Programmed Waypoint	White Symbol
Current Waypoint	Magenta Symbol
Calculated Altitude Intercept Location	Green Arc
Top of Descent	Green Circle
Extrapolation of Aircraft State	White Line

The *current mode* is identified in this prototype on the top window, in white directly above the aircraft symbol. Underneath are the control allocations for the mode, and the target states. In this example, the aircraft is in Flight Level Change (FLCH) Mode with the vertical path controlled by throttle (T) and the speed controlled by the elevator (E).

The *anticipated mode* is shown in the top window above the point where it is calculated to be engaged. The anticipated target state and control allocations are depicted in a manner similar to the current mode. Note that both the target states and the control allocation changes when the new altitude is captured.

Mode transition alerting currently consists of the anticipated mode sliding across the top window and into the current mode slot when engaged. In this case, an altitude capture arc highlights where the transition is calculated to occur.

An example of the *target state* is the dashed magenta line at 24000ft, which is the altitude currently dialed into the MCP. The target state is also shown in the mode identification window, as a vertical path at 24000ft.

The solid magenta line connecting waypoint crossing restrictions is a graphical display of the current vertical path programmed into the FMS. On the bottom of the screen are the lateral waypoints programmed into the FMS. Altitude crossings are shown by the height of the dashed lines which intersect the relevent waypoint. The path distance between each waypoint defines the horizontal scale.

The example shown in Figure 4, is a Vertical Navigation (VNAV) macro-mode descent. While still in VNAV macro-mode, note that the aircraft changes the target state of the vertical path at the Top of Descent point, a predetermined transition criteria.

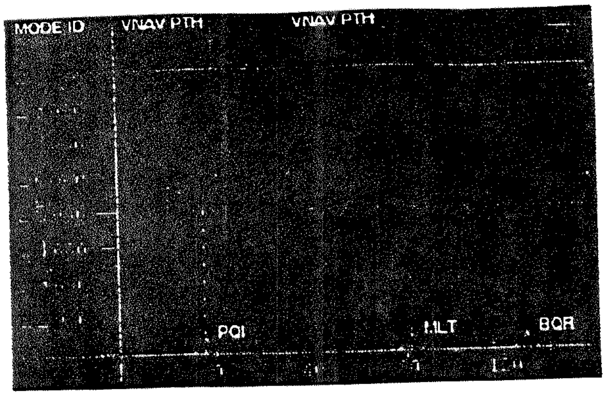

Figure 4. EVSD during VNAV descent

In the mode identification window is the current target state, an altitude of 25000 ft. At the anticipated Top of Descent point, the target state becomes a path based descent.

ACKNOWLEDGEMENTS

This work was supported by the National Aeronautics and Space Administration under grant NAG1-1581. The authors would like to thank the following individuals for their suggestions and contributions: William Corwin, Honeywell; Peter Polson, Jim Irving, Sharon Irving, University of Colorado; Michael Palmer, Kathy Abbot, Terrence Abbot, Everett Palmer, NASA.

REFERENCES

Mecham, M. (1994), Autopilot Go-Around Key to CAL Crash. *Aviation Week & Space Technology*. May 9, 1994.

Sparaco, P. (1994), Human Factors Cited in French A320 Crash, *Aviation Week & Space Technology*. January 3, 1994.

Ministry of Planning, Housing, Transportation and Space (1992). Preliminary Report Relative to the Accident which Occurred on Jan 20, 1992 near Mon Sainte-Odile (BAS-Rhine) to the Airbus A320.

Fadden, Braune, and Wiedemann, 1988. *Spatial Displays as a Means to Increase Pilot Situational Awareness*, Boeing Commercial Aircraft Company, Seattle, Washington.

Hughes, D. (1995), Incidents Reveal Mode Confusion. *Aviation Week & Space Technology*. January 30, 1995.

PRINCETON UNIVERSITY

INVESTIGATION OF AIR TRANSPORTATION TECHNOLOGY AT PRINCETON UNIVERSITY, 1994-1995

Professor Robert F. Stengel
Department of Mechanical and Aerospace Engineering
Princeton University
Princeton, New Jersey 08544

SUMMARY OF RESEARCH

The Air Transportation Research Program at Princeton University proceeded along five avenues during the past year:

- Microburst, Wind-Rotor, and Wake-Vortex Hazards to Aircraft
- Flight Control System Robustness
- Intelligent Aircraft/Airspace Systems
- Failure-Tolerant Flight Control Systems
- System Identification Using Neural Networks

This research has resulted in a number of archival papers and conference papers. An annotated bibliography of publications that appeared between June 1994 and June 1995 appears at the end of this report. The research that these papers describe was supported in whole or in part by the Joint University Program, including work that was completed prior to the reporting period.

Severe downdrafts and resulting high velocity outflows caused by microbursts present a significant hazard to aircraft on takeoff and final approach. *Microbursts*, which are often associated with thunderstorm activity, also can occur in the vicinity of dissipating convective clouds that produce no rainfall at ground level. Microburst encounter is a rare but extremely dangerous phenomenon that accounts for one or two air carrier accidents and numerous general aviation accidents each year (on average). Conditions are such that an aircraft's performance envelope may be inadequate for safe penetration unless optimal control strategies are applied. An expert system for wind shear avoidance that extends the FAA Microburst Windshear Guidelines to account for temporal and spatial variations in wind shear was developed in prior research [1]. Real-time guidance for the case in which wind shear encounter has been encountered also was investigated. Our emphasis shifted from strategies for abort and recovery based on nonlinear-inverse-dynamic controllers to the optimal nonlinear state estimation required to use this logic with realistic sensor suites [2].

The dynamics and control of a twin-jet transport encountering an intense wind "rotor" were studied [3]. It was found that a physically realizable rotor could roll the aircraft to inverted attitude if left unopposed by lateral control. Similarly, unopposed full rudder deflection could invert the aircraft in its landing configuration. Conventional linear-quadratic flight control laws can maintain wings level through such encounters.

A new effort has begun to characterize the dynamics and control of following-aircraft response to leading-aircraft wake vortices. An aerodynamic model for a subsonic jet aircraft exposed to vortical flows will be developed, concentrating on the resultant forces and moments arising from rotating wind velocity distributions. Methods for designing feedback control logic that use available control power to minimize the disturbance to the following aircraft's flight path will be derived. This logic will combine features of optimization and nonlinear-inverse-dynamic control theory to synthesize practical digital control structures.

Control system robustness is defined as the ability to maintain satisfactory stability or performance characteristics in the presence of all conceivable system parameter variations. While assured robustness may be viewed as an alternative to gain adaptation or scheduling to accommodate known parameter variations, more often it is seen as protection against uncertainties in plant specification. Consequently, a statistical description of control system robustness is consistent with what may be known about the structure and parameters of the plant's dynamic model. Rarely will there be a single "most robust" controller, as design tradeoffs must inevitably be considered. For example, stability, settling time, and control usage all may be of concern; controllers that favor one criterion over the other two have dramatically different characteristics.

Our initial research focused on probabilistic analysis of the stability and performance robustness of given controllers, while more recent research has shifted to designing robust controllers [4]. Numerical search using a genetic algorithm produces robust controllers based on proportional-filter linear-quadratic regulators with implicit model-following [5]. These controllers compare favorably to others designed by competing methods (e.g., those that minimize H_∞ cost functions). This research is proceeding with applications to control system design for a hypersonic aircraft [6] and with the incorporation of parallel processors to speed computations [7].

Advanced concepts for air traffic management are being developed by modeling aircraft and air traffic centers as intelligent agents that engage in principled negotiation [8]. Each agent is characterized as a dynamic system that carries out declarative, procedural, and reflexive functions. Principled negotiation entails the proposal of alternative flight plans, evaluation of costs and constraints according to separate and shared interests, and conflict resolution. We are setting the groundwork for an *Intelligent Aircraft/Airspace System* (*IAAS*). The goal is to identify means by which ground-based and airborne flight management systems can cooperate to produce a net gain in the efficiency and robustness of air transportation. To facilitate more general investigations of air traffic management and control, we conducted an assessment of the technology that is likely to be available during the next 30 years [9].

The annotated charts that follow present work in progress on control against wake vortex, air traffic management, failure-tolerant flight control, and aeronautical applications of computational neural networks.

ANNOTATED BIBLIOGRAPHY OF 1994-1995 PUBLICATIONS

1. D. A. Stratton and R. F. Stengel, "Real-Time Decision Aiding: An Application to Wind Shear Avoidance," *IEEE Trans. Aerospace and Electronic Systems*, Vol. 31, No. 1, Jan. 1995, pp. 117-125.

Modern control theory and artificial intelligence are applied to the Wind Shear Safety Advisor, a conceptual airborne advisory system to help flight crews avoid or survive encounter with hazardous low-altitude wind shear. Numerical and symbolic processes fuse diverse, time-varying data from ground-based and airborne measurements. Simulated wind-shear-encounter scenarios illustrate the need to consider a variety of factors for optimal decision reliability. Simulations show the potential for effectively integrating available information, highlighting the benefits of the computational techniques employed.

2. S. S. Mulgund and R. F. Stengel, "Optimal Nonlinear Estimation for Aircraft Flight Control in Wind Shear," *Proc. the 19th Congress of the International Council of the Aeronautical Sciences*, Anaheim, Sept. 1994, pp. 1747-1755.

An Extended Kalman Filter (EKF) is developed to estimate the state of a jet transport aircraft. The EKF is based on the nonlinear longitudinal aircraft equations of motion, and it is designed to provide estimates of horizontal and vertical atmospheric wind inputs. The optimal state and disturbance estimates are incorporated in feedback control laws based on the aircraft's nonlinear inverse dynamics. The EKF produces accurate estimates, and the resultant flight trajectories are very similar to those obtained with perfect state feedback. The EKF is sensitive to uncertainty in the dynamic model, but much of the lost performance can be restored by treating the uncertainty as a random disturbance input.

3. D. R. Spilman and R. F. Stengel, "Jet Transport Response to a Horizontal Wind Vortex," *J. Aircraft*, Vol. 32, No. 3, May-June 1995, pp. 480-485.

The dynamic response of a twin-jet transport aircraft encountering a single-axis wind vortex on final approach to landing is investigated. Severe performance degradation and possible ground impact may result from a low-altitude encounter with a wind rotor, which is formed by strong winds that flow over a mountain range and roll up on the leeward side of the mountain.

The simulation makes use of the similarities between flow induced over the aircraft surfaces by angular rates and the flow induced by a wind gradient. A single-axis vortex model approximates the wind velocity field. Dynamic simulations illustrate the effects of vortex strength, vortex length, lateral entry position, vertical entry position, and encounter incidence angle on the aircraft roll response parameters. Results show that maximum roll rate and roll angle increase proportionally with vortex strength and vortex length, until a "saturation length" is reached. Roll response is highly dependent on entry location: changes in lateral entry position affect maximum roll angle while changes in vertical entry position affect maximum roll rate. Peak roll rate and roll angle obtain their largest values at near-zero incidence angles. The response is highly dependent on the initial conditions of the encounter -- even small variations cause significant changes in aircraft roll response.

4. C. I. Marrison and R. F. Stengel, "Stochastic Robustness Synthesis Applied to a Benchmark Control Problem," *Int'l. J. Robust and Nonlinear Control*, Vol. 5, No. 1, Jan. 1995, pp. 13-31.

Stochastic Robustness Analysis guides the synthesis of robust LQG regulators for a Benchmark Control Problem. Probabilities of exceeding allowable design limits, including stability, settling time, and control usage, are estimated by Monte Carlo evaluation. Robust, low-gain compensators that fulfill objectives are designed by numerically minimizing quadratic functions of these probabilities. The method is straightforward and makes use of uncomplicated design principles.

5. C. I. Marrison and R. F. Stengel, "The Use of Random Search and Genetic Algorithms to Optimize Stochastic Robustness Functions," *Proc. 1994 American Control Conference*, Baltimore, June 1994, pp. 1484-1489.

Stochastic robustness synthesis is a framework for designing practical control systems. It uses Monte Carlo simulation to evaluate the quality of candidate designs, and it searches a parameter space to find the best one. The global minimum of a probabilistic criterion function must be found, ideally with a minimum number of evaluations. This paper examines two approaches to minimizing the probabilistic function: random search and a genetic algorithm. The genetic algorithm is similar to previously published algorithms but has several modifications to improve its performance, most notably a clustering analysis at the beginning of each generation. Statistical tools are incorporated in the search algorithms, allowing intelligent decisions to be based on the "noisy" Monte Carlo estimates. Performance of the two methods is demonstrated by application to a 24-dimensional test function. The genetic algorithm is shown to be significantly better than the random search for this application. The genetic algorithms is then used to design compensators for a benchmark problem, producing control laws with excellent levels of stability and performance robustness.

6. C. I. Marrison and R. F. Stengel, "Synthesis of Robust Control Systems for a Hypersonic Aircraft," *Proc. 33rd IEEE Conference on Decision and Control*, Orlando, Dec. 1994, pp. 3324-3329.

Stochastic Robustness Analysis is a flexible framework for defining the robustness of control systems. It defines the robustness of a compensator by the *probability* that parameter variations will cause the closed-loop system to have unacceptable behavior. Here, robust linear-quadratic regulators are synthesized to control the nonlinear longitudinal dynamics of a hypersonic aircraft with uncertainties in 28 parameters. The compensators are designed using a genetic algorithm to search a design parameter space and Monte Carlo evaluation to estimate the robustness of the compensator at each search point. The method is shown to produce control structures that satisfy nominal stability and performance goals, while minimizing robustness cost functions.

7. W. M. Schubert and R. F. Stengel, "Parallel Stochastic Robustness Synthesis for Control System Design," *Proc. 1995 American Control Conference*, Seattle, June 1995, pp. 4429-4434.

Stochastic Robustness Synthesis is used to evaluate compensator robustness numerically and to automate the design of stochastic optimal controllers. Monte Carlo Simulation (MCS) is applied to quantify robustness, and a Genetic Algorithm (GA) searches for the optimal controller. The overall algorithm is computationally expensive, and parallel computing is utilized to reduce execution times. Parallel Stochastic Robustness Analysis and Design (PSRAD) is introduced as a viable solution for real-time controller design. A Dynamic Scheduler is proposed to alleviate stochastic load imbalances. Results are presented for a shared-virtual-memory computer.

8. J. P. Wangermann and R. F. Stengel, "Principled Negotiation Between Intelligent Agents: A Model for Air Traffic Management," *Proc. 19th Congress of the International Council of the Aeronautical Sciences*, Anaheim, Sept. 1994, pp. 2197-2207.

The worldwide aircraft/airspace system (AAS) is faced with a large increase in air traffic in the coming decades, yet many flights already experience delays. The AAS is comprised of many different *agents*, such as aircraft, airlines, and traffic control units. Technology development will make all the agents in the AAS more intelligent; hence, there will be an increasing overlap of the *declarative functions* of the agents. This paper describes the basis for an *Intelligent Aircraft/Airspace System* (IAAS) that provides improved system performance, redundancy, and safety by utilizing the overlapping capabilities of the agents. *Principled Negotiation* between agents allows all the agents in the system to benefit from multiple independent declarative analyses of the same situation. *Multi-attribute utility theory* and *decision trees* are used as the basis for analyzing the behavior of different types of agents. Intelligent agents are modeled as *rule-based expert systems* whose *side-effects* are the procedural and reflexive functions of the agent. Principled negotiation also is a side-effect of the expert system's declarative functions. A hierarchical organization of agents in the IAAS is proposed to facilitate negotiation and to maintain clear lines of authority.

9. J. P. Wangermann and R. F. Stengel, *Technology Assessment and Baseline Concepts for Intelligent Aircraft/Airspace Systems*, Princeton University Report MAE-2016, Feb. 1995.

This report describes the technology and operating concept for an *Intelligent Aircraft/Airspace System* (IAAS) that will meet the demand for air travel in the year 2025 while improving safety standards and efficiency for all users. Technology development trends are examined, and artificial intelligence and decision-making techniques that can exploit these changes are described. The operation of the IAAS is illustrated by describing a typical flight of an aircraft through such a system.

Today's global *aircraft/airspace system* (AAS) consists of a hierarchy of *agents*, each making decisions and taking actions that affect air transport operations in the AAS. The functions of each agent can be classified as *declarative, procedural,* or *reflexive.* Declarative functions are high-level, involving conscious, strategic decision-making. Reflexive functions are low-level, requiring little or no conscious "thought". Procedural functions are intermediate; often these are the actions taken to convert a strategic plan into a sequence of reflexive actions.

The report surveys the possible short- and long-term technology developments in traffic sensors, weather sensors, head-up cockpit displays, navigation equipment, approach guidance equipment, communications, computation, aircraft design, and airport design. The report concludes that these developments will greatly increase the declarative decision-making capabilities of all agents. The overlap between agents in declarative capabilities will dramatically increase. For example, airliners will have traffic situation awareness that is nearly as good as that of a traffic management unit. There will be far less improvement in procedural and reflexive function performance. New technology will also provide much expanded communications bandwidth between agents, an "Internet in the sky". The combination of dramatically increased, overlapping declarative capabilities and communications bandwidth could significantly improve the safety, robustness, and performance of the AAS.

In order to meet the demands of the year 2025, the AAS must exploit these technological developments. The structure of today's AAS arose when there was little overlap in the declarative functions of the agents in the system. This report proposes a structure for an *Intelligent Aircraft/Airspace System* (IAAS). By using *principled negotiation* as the basis for all interactions between agents, the IAAS harnesses the enhanced and overlapping capabilities of all agents and the increased communications bandwidth to improve system performance. Principled negotiation encourages all agents to search for *options for mutual gain.* Each agent wants to carry out a set of actions that maximizes their own utility function, but they are best able to achieve this if they propose actions that benefit other agents in the system too. Hence, by aiming to improve their own utility, each agent improves system performance for all agents.

The actions and negotiations that occur during a typical flight in an IAAS are outlined. All parts of a flight trajectory, from push-back to nose-in, are negotiable. Airlines and aircraft constantly search for options for mutual gain. For example, aircraft, using their enhanced traffic situation awareness, can propose trajectory changes that not only save fuel, but also improve aircraft separations. Proposed trajectories must be approved by Traffic Management Agents before they are implemented. Provided they are conflict-free, proposals will generally be approved. TrMAs need to resolve conflicts only if the aircraft or operators fail to find suitable solutions in a timely manner. The IAAS provides aircraft and operators with much more freedom to optimize their flights, while improving levels of safety and system performance through their pro-active involvement.

The report concludes by describing a transition strategy to achieve an IAAS by 2025. Principled negotiation can be applied to agent interactions immediately. Doing so would ensure that maximal benefit was gained from any new technology introduced by any agent.

Joint University Program

Nonlinear Inverse Dynamic Control of Aircraft in Wake Vortex Wind Shear

Research Summary, 1994-1995

Gregory R. Wold
Dept. of Mechanical and Aerospace Engineering
Princeton University

The world's airports are today already near capacity, and the projected increase in future air traffic combined with little new airport construction does not bode well. Some other means of expanding capacity must be found. One possibility is to reduce the aircraft spacing requirements.

Following	Generating Aircraft		
Aircraft	Small	Medium	Large
Small	3	4	6
Medium	3	3	5
Large	3	3	3
	(Distance in nmi)		

The wake vortex hazard is the dominant safety issue determining these spacing requirements. There have already been a number of accidents attributed to encounters with wake vortices, some resulting in deaths. As air traffic increases, the number of dangerous encounters is likely to as well. So a means of reducing the danger must found. This is the purpose of this research: to use nonlinear inverse dynamic (NID) control methods in the design of an aircraft control system which can improve the safety margin in a wake vortex encounter.

Nonlinear Inverse Dynamics

Given a nonlinear system of the form

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{G}(\mathbf{x})\mathbf{u}$$
$$\mathbf{y} = \mathbf{H}\mathbf{x}$$

sufficient differentiation of each of the elements of y makes a component of u appear. The results can be assembled as

$$\mathbf{y}^{(d)} = \mathbf{v} = \mathbf{f}^*(\mathbf{x}) + \mathbf{G}^*(\mathbf{x})\mathbf{u}$$

For non-singular $\mathbf{G}^*$, the inverse system takes the form

$$\mathbf{u} = \left[\mathbf{G}^*(\mathbf{x})\right]^{-1}\left[\mathbf{v} - \mathbf{f}^*(\mathbf{x})\right]$$

and applying the NID control law yields the full system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{G}(\mathbf{x})\left[\mathbf{G}^*(\mathbf{x})\right]^{-1}\left[\mathbf{v} - \mathbf{f}^*(\mathbf{x})\right]$$

Conventional flight control systems are developed assuming the aircraft dynamics to be linear and time invariant about some nominal flight condition. The performance of such linear control systems breaks down in extreme flight conditions, as the higher-order terms of linearization become non-negligible. A wake vortex encounter is just such an extreme condition.

Nonlinear inverse dynamics is one means of better accounting for these higher-order effects, providing better control over a greater portion of the flight envelope. An output vector y dependent on x is defined, and is differentiated with respect to time, element by element, until the control effect appears. The assembled output derivative vector **v** becomes the input vector for the inverse system. It is specified as a function of the desired outputs y, accounting for suitable dynamics.

NID controllers must be used with care. They are characterized by a significantly higher computational cost than are linear controllers. Singularity conditions must lie outside the allowed flight envelope, to ensure the inevitability of G*. Also, the design is often based on a reduced plant model, as small effects, if fully accounted for, can result in system instability.

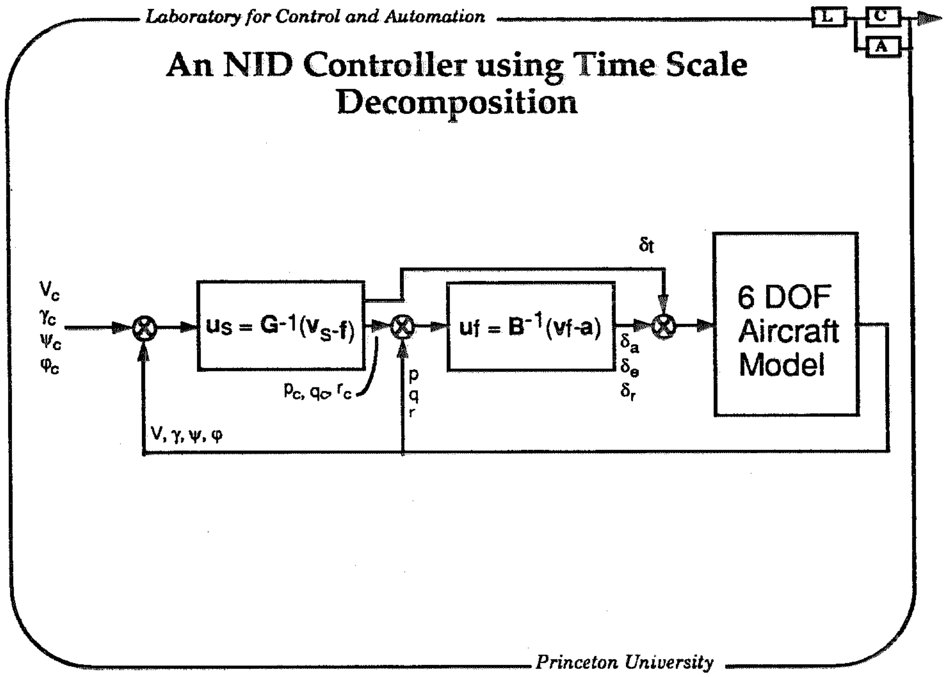

The design of an NID controller can be simplified somewhat. If sufficient time-scale separation exists between elements of the system state vector, the full system can be broken into subsystems in a process known as Time Scale Decomposition. The subsystems generally require lower orders of differentiation than would the full system, decreasing the mathematical complexity. One example of an NID controller using this method is shown above. The flight path variables evolve more slowly than the body-axis angular rates , so they are used to specify "desired" p, q, and r though the slow-time scale inverse system. Then the necessary control surface deflections are generated with the fast-time scale system built around the angular rates.

Another useful feature of this method is the implicit elimination of the small control effects mentioned earlier. The lift effect of rudder deflection is ignored in the design of the slow-time scale inverse system, for example. Using it would result in probable instability, through the nonlinear equivalent of non-minimum phase zeros.

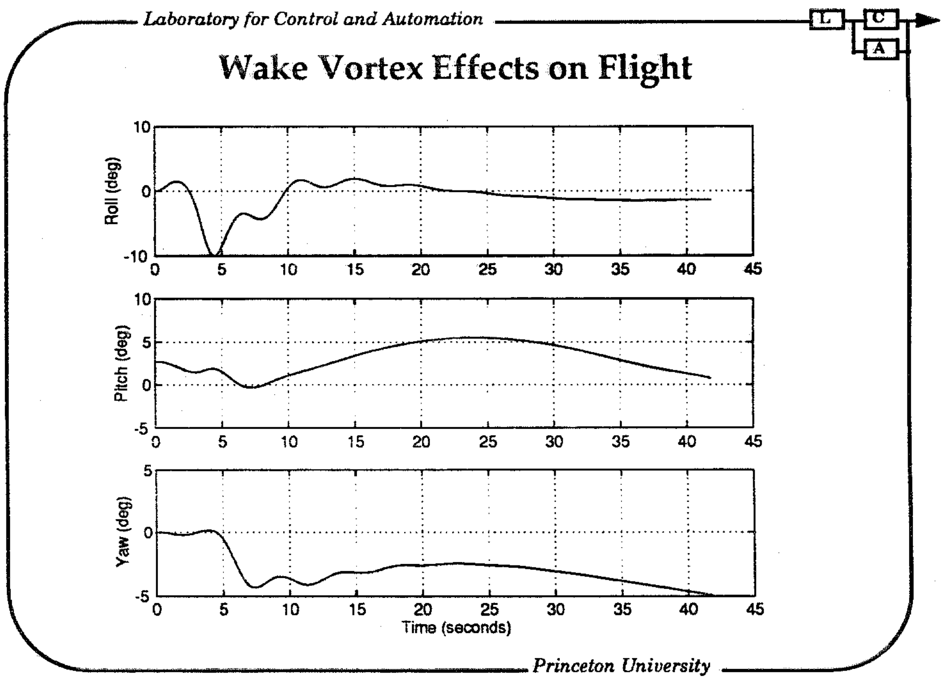

The wake vortex system is composed of dual, counter-rotating vortices spaced just under the generating craft's wingspan apart. Approaching from outside of and level with the system, an aircraft would experience an initial upwash on the near wing, resulting in a roll away from the vortices. Further penetration into the core of the near vortex would result in a snap roll into the center of the system, where there is a strong downwash field. This is the region of greatest hazard, as the resulting roll angle combined with the loss on altitude due to the downwash could result in a crash.

The cores of the vortices have been shown to be much smaller than the span of an airplane, on the order of a meter. The shear field can be modeled as a solid-body rotation within the core, surrounded by a field that falls off inversely with respect to the radius from the core. As such, there is tremendous variation in the induced aerodynamic angles across the span of an aircraft in a wake vortex field. There are several methods for simulating the force and moment effects the field will cause, including vortex-lattice methods. Strip theory is another; it is relatively simple and conservative, and generates the proper general effects, as can be seen in the plots above for a plane passing through a wake vortex system.

Wind Term Effects on Controller Design

Kinematics

Position

$$\begin{bmatrix} \dot{x}_E \\ \dot{y}_E \\ \dot{z}_E \end{bmatrix} = \mathbf{L}_{EB} \begin{bmatrix} u \\ v \\ w \end{bmatrix} + \mathbf{W}_E$$

Orientation

$$\begin{bmatrix} \dot{\phi} \\ \dot{\theta} \\ \dot{\psi} \end{bmatrix} = \begin{bmatrix} p + (q\sin\varphi + r\cos\varphi)\tan\theta \\ q\cos\varphi - r\sin\varphi \\ (q\sin\varphi + r\cos\varphi)\sec\theta \end{bmatrix}$$

Dynamics

Linear

$$\begin{aligned} \dot{u} &= \frac{X}{m} - g\sin\theta - q(w + W_{z_B}) + r(v + W_{y_B}) - \dot{W}_{x_B} \\ \dot{v} &= \frac{Y}{m} + g\sin\varphi\cos\theta - r(u + W_{x_B}) + p(w + W_{z_B}) - \dot{W}_{y_B} \\ \dot{w} &= \frac{Z}{m} + g\cos\varphi\cos\theta - p(v + W_{y_B}) + q(u + W_{x_B}) - \dot{W}_{z_B} \end{aligned}$$

Angular

$$\dot{q} = \frac{1}{I_{yy}}\left[M_m + I_{xz}(r^2 - p^2) + (I_{zz} - I_{xx})rp\right]$$

$$\begin{bmatrix} \dot{p} \\ \dot{r} \end{bmatrix} = \begin{bmatrix} I_{xx} & -I_{xz} \\ I_{xz} & I_{zz} \end{bmatrix}^{-1} \begin{bmatrix} L_m + I_{xz}pq + (I_{yy} - I_{zz})qr \\ N_m + I_{xz}qr + (I_{xx} - I_{yy})pq \end{bmatrix}$$

The wind terms enter both explicitly (WE and WB) and implicitly (multiplying the force and moment coefficients).

One means of improving the performance of a controller is to account for known or estimated disturbance effects in the controller design. However, the introduction of disturbance terms into the aircraft equations of motion complicates the controller design. In the case of wind shear, the effects on the equations are both implicit, through the aerodynamic forces and moments (which are calculated relative to the surrounding air mass), and explicit, in modifying the inertial velocity of the craft. Using a mixed inertial- and body-axis formulation, as is shown above, the added complexity is fairly innocuous. A design which uses wind-axis definitions, however, would be computationally too expensive to use with wind terms added. This is due to new couplings between state vector elements introduced only with the disturbance terms; in particular, derivatives of the wind-axis angular rates must be computed, and these require computation of the wind-axis moments and products of inertia - something certainly to be avoided! Just such a design was presented here, until this difficulty was discovered. Now, a new controller is being designed, based on command variables defined with the equations above.

Further Work

- Design and code new NID controller
- Verify strip theory code
- Evaluate NID controller effectiveness
  - No disturbance terms in controller
  - Perfect wind "measurements"
- Possibly test other NID controllers

Current and future work will complete the design and testing of a new NID controller, based on the knowledge that a wind-axis formulation is not adequate for accounting for disturbances. The strip theory code will be further verified, ensuring incorporation of proper stall effects and by comparing the magnitude of induced roll angles with those of past studies. Then the two code segments will be integrated into an existing simulation of a twin-jet commercial airliner, and the effectiveness of the controller, both with and without wind effects, will be tested. The wind effects will be introduced assuming a perfect estimator; introduction of a more realistic estimator can be done in follow-on studies. Finally, NID controllers utilizing other command vector elements may be tested as well.

Distributed Search Coordinated by Principled Negotiation for Advanced Aircraft/Airspace Systems

John Paul Wangermann

Princeton University

1995

Laboratory for Control and Automation

Principled Negotiation Between Agents

- **Each agent searches for *Options for Mutual Gain***
  - **Increase utility actions to the agent**
  - **Increase utility to other agents**
- **Agent proposes OMG to affected agents**
- **If agreed upon, option is implemented**

- **Such a system provides**
  - **Allows iterative improvements to all flight plans**
  - **Allows all agents to use their declarative capabilities to analyze the traffic situation**
  - **Provides aircraft operators with a payoff for improving the equipment fit on aircraft**

Princeton University 2

The aircraft/airspace system (AAS) is a multi-agent system. The different agents (e.g., airlines, aircraft, traffic management agents) have different interests. With high-bandwidth communications making data readily available to all agents, there is an increasing overlap in their ability to analyze the traffic situation. Each agent may have a different view as to what plan of actions is preferable. A method of negotiation is required that allows the agents to consider all the alternatives and reach a good outcome quickly.

In *principled negotiation*, agents examine the interests of other agents and try to generate options that provide mutual gain: i.e., provide increased utility to other agents as well as themselves.

Allowing agents to propose and negotiate over alternate solutions allows dynamic optimization of AAS operations and should greatly improve system performance.

Searching for Options for Individual Gain

General

- **Set of agents** $N=\{i\} \quad i=1,\ldots,n$ (1)
- **Action profile** $a=(a_1,a_2,\ldots,a_n)$ (2)
- **Outcome function** $c=g(a,\omega,v)$ (3)
- **Utility function and preference**
  - **Maximizer** $a \succ b$ iff $E_i[u_i(g(a))] > E_i[u_i(g(b))]$ (4)
  - **Satisficer** OK if $E_S[Sat(g(a))] \geq S_{min}$ (5)
- **Proposal acceptable to all agents if**

$$E_j[u_j(g(a))] \geq E_j[u_j(g(b))] \quad \forall j \in N \qquad (6)$$

To formulate the problem, we consider a system with n agents (1). Each agent is following a particular action plan (2). The collection of all the action plans is the action profile for the system. The action profile a will result in a particular set of outcomes for the system (3). The outcome depends on the current system state ω and uncertain disturbances **v**. Some agents will be *maximizers,* others *satisficers.* The preference functions for each type of agent depend on the agent's expectation of the outcome (4) (5). A proposal is acceptable to all agents if the effect of the proposal is at worst neutral (6).

Individual Gain

- **Agent *i* searches for options by varying own action plan only**
  - **Existing plan** $a=(a_1,a_2,..\boxed{a_i,}..,a_n)$ (7)
  - **Option** $b=(a_1,a_2,..\boxed{b_i,}..,a_n)$ (8)
- **In general multi-agent system, there is no guarantee that**

$$E_j\left[u_j(g(b))\right] \geq E_j\left[u_j(g(a))\right] \quad \forall j \neq i, \in N \tag{9}$$

In the simplest form of principled negotiation, each agent searches for options that provide individual gain but do not affect other agents. This is done by looking at options (new action profiles) that change only the agent's own action plan. In general, there is no guarantee that because options satisfy equations (7) and (8) that (9) also holds.

Individual Gain in an AAS

- **AAS General form of utility function:**

$$u_i = \mathbf{w}_i^{\mathbf{T}}(\omega)\mathbf{y}_i(\mathrm{g}[a,\omega,v])$$

$$\mathbf{y}_i = \left[y_i^f, y_i^t, y_i^{sa}, y_i^{se}, y_i^o\right]^{\mathbf{T}}$$

y_i^f = fuel usage measure

y_i^t = timeliness measure

y_i^{sa} = separation from other aircraft measure

y_i^{se} = separation from environmental hazards measure

y_i^o = measure of other attributes

- **Cost associated with separation is the only factor affected by other aircraft**

However, in an AAS a common form of the utility function for an aircraft would be (10). None of the components of the utility function are affected by the flights of other aircraft apart from the aircraft separation measure.

Required form of Separation Cost Function

- **If each type of agent agrees on minimum safe separation, then utility functions decouple for safely separated trajectories**
  - **change in a_i affects u_i only**

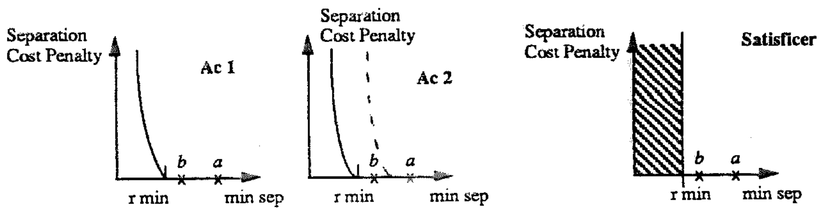

If each agent agrees on the minimum safe separation, then an agent can guarantee that any proposed option that satisfies the separation criteria will be accepted by the other agents. Rather than negotiating with each other aircraft in the system, the agent can negotiate instead with a coordinating agent acting as a satisficer. The satisficing function must be the same as the agreed safe minimum separation.

Agreement on objective criteria for assessment of options is a tenet of principled negotiation.

IAAS Implementation

- **Action Plan**
  - **The action plan of an aircraft is a trajectory to be flown**
  - **Trajectory defined by maneuver sequences**
    - ***(manType stTime target rate ....)***
    - **e.g., (T 260.0 840.0 10.0 ......)**
    - **Compact description of any trajectory**
    - **Compatible with many types of search**
- **Role of Traffic Management Agent**
  - **Defines the separation requirements**
  - **Supplies information to agent (e.g., presence of other aircraft)**
  - **Approves/rejects proposed changes in flights**

A computer model of an Intelligent Aircraft/Airspace System based on principled negotiation is being developed. (Screen shots at end of notes.) The action plan of each aircraft is represented as a *maneuver sequence*. A heuristic rule base is used to search for options for individual gain. Each aircraft uses a bank of hybrid extended Kalman filters to estimate future motion of other aircraft to ensure that separation standards are met. Any changes in an aircraft's action plan must be approved by a traffic management agent (TrMA).

Screen shots from the computer simulation are included at the end.

Searching for Options for Mutual Gain

- **Individual Gain**
  - **Agent i changes a_i only**

$$E_i\left[u_i(g(b))\right] \geq E_i\left[u_i(g(a))\right] \quad (13)$$

  - **If no conflicts**

$$E_j\left[u_j(g(b))\right] = E_j\left[u_j(g(a))\right] \quad \forall j \neq i, \in N \quad (14)$$

  - **If conflict being resolved**

$$E_j\left[u_j(g(b))\right] \geq E_j\left[u_j(g(a))\right] \quad \forall j \neq i, \in N \quad (15)$$

- **Mutual Gain**
  - **Agent i proposes changes to some a_j also**

$$E_i\left[u_j(g(b))\right] \geq E_i\left[u_j(g(a))\right] \quad \forall j \neq i, \in N \quad (16)$$

  - **To be accepted**

$$E_j\left[u_j(g(b))\right] \geq E_j\left[u_j(g(a))\right] \quad \forall j \neq i, \in N \quad (17)$$

When searching for mutual gain, the option does not have to obey (8). In general, it can propose a new action profile involving changes in the action plans of other agents if it expects that this will provide mutual gain (16) (17).

Slot Negotiation Scenario

Effect of Original EDCT

AL and #	Ac Type	Dep Airport	Original Arrival Time	New Scheduled Arrival	Total Pax	Connecting Pax
CO31	DC10	EWR	3.10	5.20	300	180
CO42	737	SFO	2.45	4.00	120	40
UA66	767	LAX	3.55	6.15	240	180
UA82	737	SEA	2.10	3.20	120	35

Individual Gain

Flight	Orig Time	New Sched Time	Delay (mins)
CO31	3.10	4.00	50
CO42	2.45	CANCELLED	Equivalent 300
UA66	3.55	6.15	140
UA82	2.10	3.20	70

Mutual Gain

Flight	Orig Time	New Sched Time	Delay (mins)
CO31	3.10	3.20	10
CO42	2.45	6.15	210
UA66	3.55	4.00	5
UA82	2.10	5.20	130

This example shows the benefits of negotiation and search for options for mutual gain. Consider a Ground Delay Program imposed on flights into Denver. The arrival slots allocated to four flights are shown in the top table. Under the current system (no negotiation between airlines, slots may only be swapped if a flight is canceled first), only CO could make changes that provide benefit (the lower missed connection costs outweigh the cancellation costs). UA cannot make a similar change as UA66 cannot meet UA82's slot in any case (middle table).

If negotiation between airlines were possible, they could negotiate a slot swap to provide mutual gain (lower table).

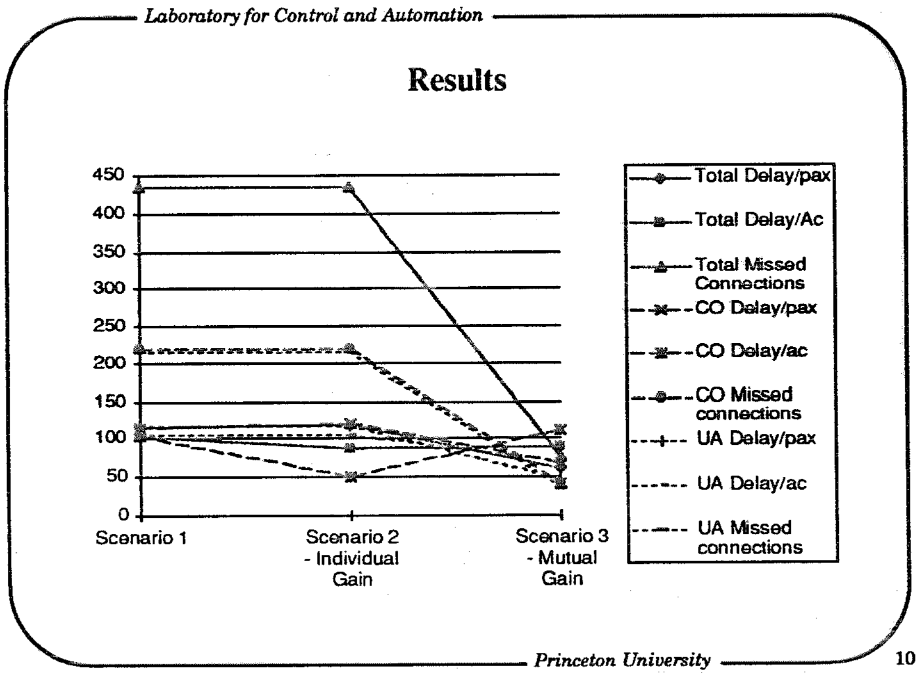

By almost every measure, the negotiated mutual gain solution is by far the best. Although CO's delay/aircraft is not as good as in the individual gain case, it does not have to cancel a flight. In terms of missed connections the mutual gain solution is the best. (Airlines incur costs from missed connections because aircraft, flight crew, and cabin crew all need to be rescheduled, and passengers may need compensation.)

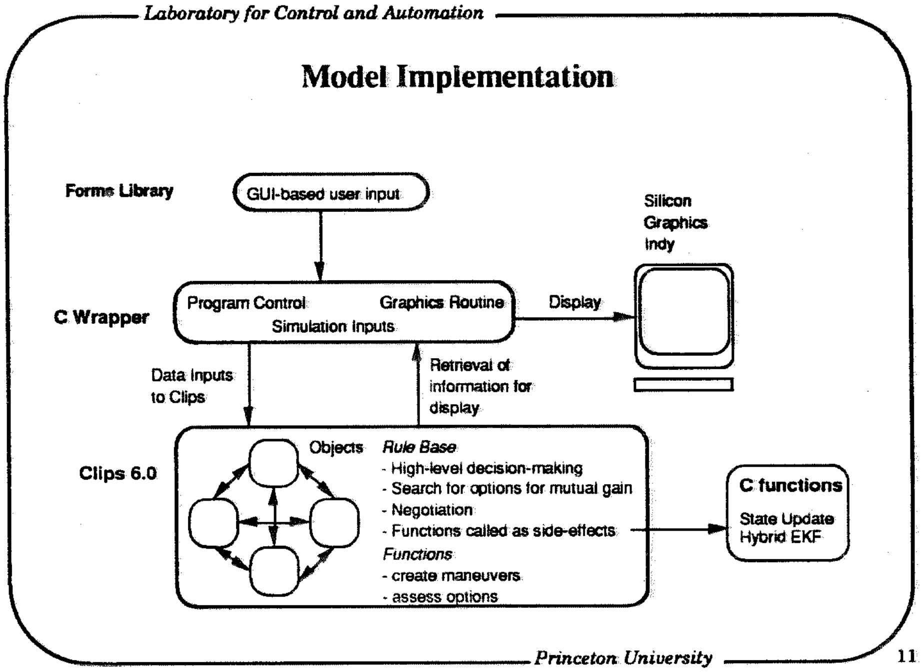

The computational model of an Intelligent Aircraft/Airspace System incorporating principled negotiation is being developed. The aim is to quantify the operational benefits to aircraft, airlines, and traffic managers from distributed optimization by agents coordinated through principled negotiation. The simulation is being implemented on a Silicon Graphics Indy in C and Clips (an expert system language).

Screen shots showing the overall output, the aircraft situation display, and negotiation are attached.

Figure 1. Overall view showing the aircraft situation display (1), Traffic Management Agent Window (2), Aircraft Status Windows (3), and a negotiation dialog box (4).

Figure 2. Close-up of the aircraft situation display showing eight aircraft, their flight IDs, and altitudes.

Figure 3. A negotiation dialog box. KL892 is proposing to turn immediately to heading 3.12 radians, followed by a climb 16 seconds later, leveling off after 166 seconds. This information is shown both in the TrMA Window (above) and a negotiation dialog box (below). The experimenter (acting as TrMA) is given the option to accept or reject the proposal.

<u>FIGURE 2: A/C SITUATION DISPLAY</u>

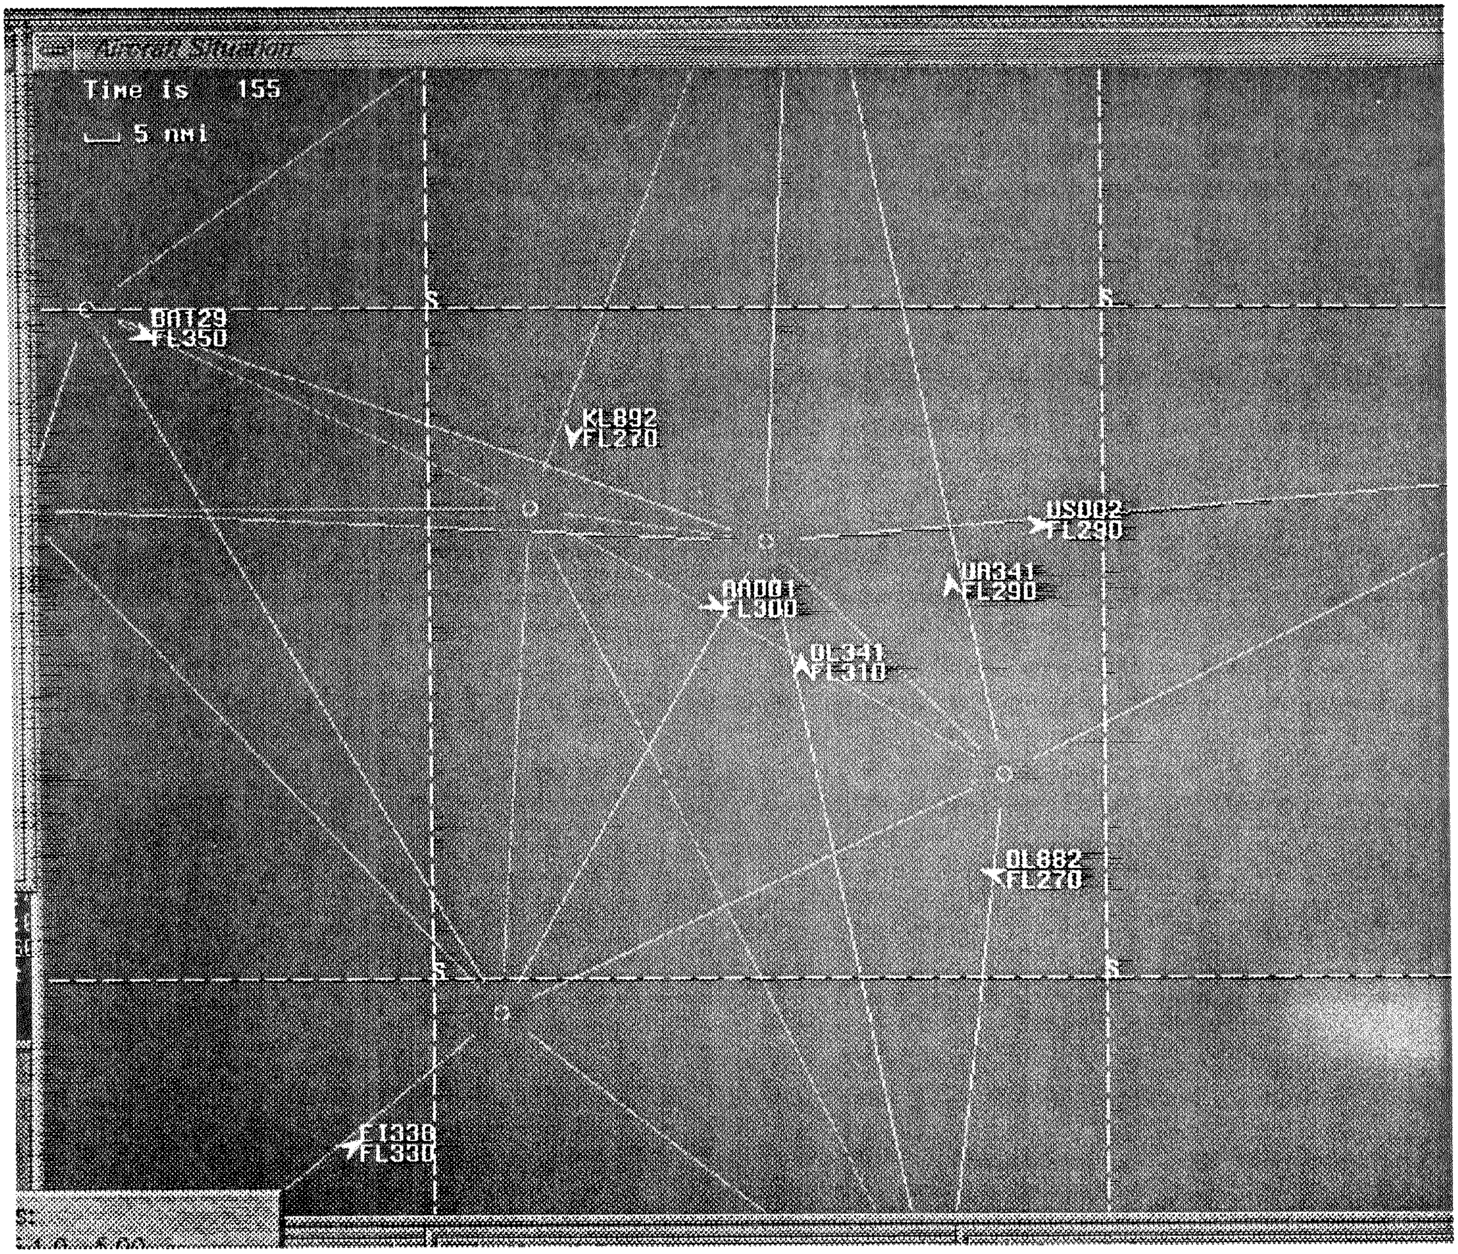

FIGURE 3: NEGOTIATION DIALOG BOX

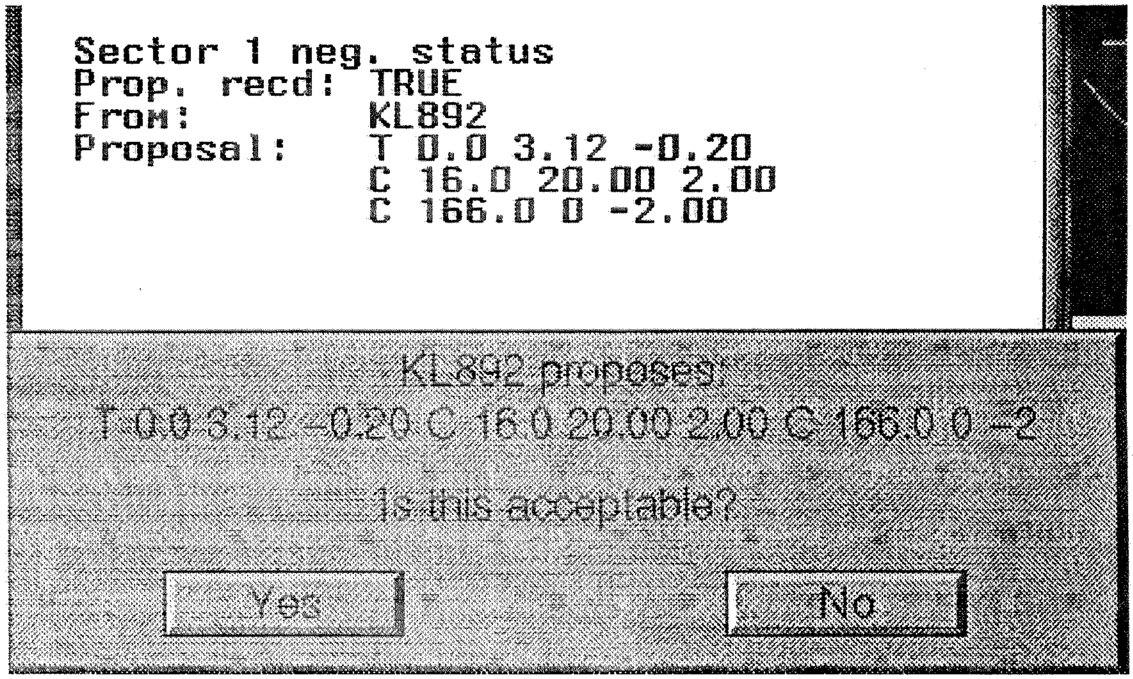

Laboratory for Control and Automation

Plan for future work.

- **Implement negotiation on end constraints/waypoints**
- **Add airline, airport TrMA models to simulation**
- **Refine rule bases for searching for options**
- **Conduct comprehensive performance testing**

Princeton University 12

Negotiation on required time of arrival for aircraft is currently being implemented in the computer simulation, along with the necessary models for airline and TrMA behavior. The rule bases which agents use to search for options are being refined. Once implemented, comprehensive testing using the computer model will begin in order to quantify the benefits of distributed search and principled negotiation on IAAS operations.

Failure Accommodation in Real-Time

Sai Manohar Gopisetty

Joint University Program, 1995

Present research focuses on implementation of a knowledge based architecture for failure accommodation in flight control. Emphasis is laid on non-linear dynamic models and research for appropriate design methods.

Need for failure accommodation

- **Complex control systems** consisting of a number of devices : sensors / actuators / transmitters. Ex: Aircraft.
- Human operator can be inefficient:
  - Mismanagement.
  - Slow in execution.
  - Stress.
- Automated failure tolerant systems can find use in:
  - Autonomous underwater/ space vehicles.
  - Chemical process control , nuclear power plants.
  - Robotics.
  - Continuous production line systems.

Modern aircraft are complex systems requiring continuous operation of a number of sub-systems. Failure of any part can degrade the system performance seriously. Pilot load in such an emergency situation is very demanding and through past experience as evidenced by various NTSB accident reports, many emergencies could have been handled better. This suggests the need for the automation of the overall control process which can, autonomously or with the aid of the pilot, greatly improve the system reliability.

This concept can in general be applied to a number other applications involving autonomous control.

Laboratory for Control and Automation

Requirements of a failure tolerant system

- Constantly **monitor** the state of the system.
- **Judge** whether the states are normal or faulty.
- **Diagnose** even unknown faults.
- Generate possible **repair plans.**
- **Validate** and **execute** the repair plan.

In short, *Failure Detection, Identification and Reconfiguration* in real-time.

Princeton University

Given a set of noisy measurements form sensors and commanded inputs to the actuators, we first need to detect when there is an inconsistency in the data. Further we need to determine the size and type of the fault. For example, redundancy in a particular measurement can help to detect a malfunction in an associated sensor and isolate it. Failure detection and isolation can be at different stages or together depending upon the complexity and the application in question.

A crucial task to follow up is that of system reconfiguration. Identified faults can be directly accounted if they are expected. Unidentified / unanticipated faults require some kind of adaptation in the reconfiguration logic. All the above actions should take place in real-time, adding to the complexity.

Laboratory for Control and Automation

Intelligent Real-Time Control

- **Knowledge based** systems are an attractive platform for real-time control.
  - Provide hierarchical structure for model and information processing.
  - Combine declarative and procedural logic.
  - Amenable to parallel / distributed hardware.

- Previous KB related work at LCA:
  - D.A.Handelman: Anticipated faults in real-time.
  - Chein Huang: Causal reasoning approach.

Princeton University

The very complexity of the task suggests the need for an architecture based on artificial intelligence concepts. This would involve for example, incorporation of the knowledge of an expert in emergency procedure execution once a fault is detected. Previous work by Handelman (1988 , Princeton Univ.) showed how such a system can be implemented in real time. Huang (1988, Princeton Univ.) , in his work formulated a causal-reasoning approach to identify potentially unanticipated failures and a systematic approach towards reconfiguration.

The present research, as an extension to these works is intended to implement in real-time FDIR of unanticipated faults.

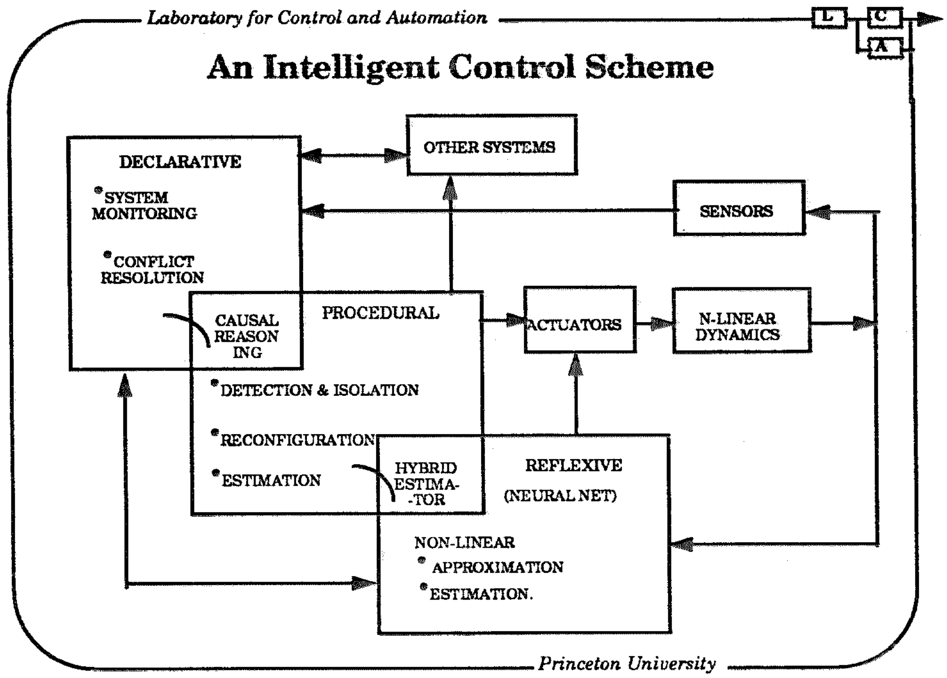

A knowledge based intelligent control scheme to serve as an architecture for the present problem is shown in the figure. It is made up of a hierarchical design, with declarative logic for goal setting, redundancy resolution etc., at the top level followed by procedural logic involving algorithmic execution for gathering information as well as computing actuator inputs, and finally reflexive actions like non-linear function approximation carried out by a neural network at the lowest level (in terms of command-action flow).

Integration of these modules also suggests synergetic components such as a hybrid estimator with combined procedural and reflexive components for on-line non-linear estimation and a heuristics aided causal search combining declarative and procedural logic for failure identification. Similarly declarative and reflexive actions can be combined for training a neural net / expansion of knowledge base.

Failure Detection and Identification

- **Parity Relations:** One elegant way to get exact sensor and actuator failure signatures through *analytical redundancy.*
- Model : (discrete, LTI)

$$x(k+1) = Fx(k) + Gu(k) + w + L_i m_i(k)$$
$$y(k) = Hx(k) + n + J_i n_i(k)$$

- A method to generate all possible dynamic relations between sensor output and actuator inputs.
- Formulation in brief :

$$\textbf{Define } P = \left\{ V \mid V^T \underbrace{\begin{bmatrix} H \\ HF \\ \vdots \\ HF^s \end{bmatrix}}_{T} = 0 \right\} \quad \textbf{Parity space of order s}$$

Residual generation:

$$r(k) = V^T \left[\begin{bmatrix} y(k-s) \\ \cdots \\ y(k) \end{bmatrix} - C \begin{bmatrix} u(k-s) \\ \cdots \\ u(k) \end{bmatrix} \right]$$

Analytical redundancy, which gives a means of relating various sensor outputs and actuator inputs using the underlying dynamic model of the system (aircraft). This is different from, and in the process reduces the burden of, physical redundancy of sensors to determine the type of failure. In the model shown, m_i, n_i represent actuator and sensor failure modes.

There are a number of techniques to exploit analytical redundancy for the purpose of FDIR. Parity space methods are particularly attractive because of their robustness to certain modeling errors and computationally effective nature. Parity relations for failure tolerance were first investigated by Potter and Suman in 1977 and later extended by Willsky and his group from 1984.

Parity Relations

Where

$$C = \begin{bmatrix} 0 & & & & \\ HG & 0 & & O & \\ HFG & HG & 0 & & \\ \cdots & \cdots & \cdots & \cdots & \\ HF^{s-1}G & \cdots & HFG & HG & 0 \end{bmatrix}$$

Using the above realtions , and with no failures,

$$r(k) = V^T \begin{bmatrix} H \\ HF \\ \cdots \\ HF^s \end{bmatrix} X(k-s)$$

Which is nominally zero.

V being the orthogonal complement of range space of T can be constructed by the orthogonal projection on to T from the Identity operator.

$$P \Leftrightarrow \{\text{independent rows of } I - T(T^T T)^{-1} T^T\}$$

By proper coding, since all possible redundant realtions are formed, any sensor or actuator failure can be identified. Further, it is possible to show that P can be chosen for robustness.

The residual vector, as defined above gives all possible dynamic relations between measurements and control inputs and is nominally zero (or zero mean with noise present) when no faults are present. A parity vector can always be determined by definition as shown in the construction above form the identity operator. Occurrence of a fault is now triggered automatically when the magnitude of the residual exceeds a threshold. By a systematic choice of the parity vectors used in the residual generation, specific failures can be identified as a bias of the residual in parity space. Moreover specific parameter modeling uncertainties could be accommodated with further refinement of design.

Laboratory for Control and Automation

Non-linear Dynamic Models

- Need for a good non-linear model first to avoid *linearization errors.*
- Consider the most general non-linear model:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x},\mathbf{u},\mathbf{w},t)$$

- No general methods to tackle such forms.
- Next best model: $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x},\mathbf{u},t) + \mathbf{G}(\mathbf{x},\mathbf{u},t)\mathbf{u}$, $\mathbf{y} = \mathbf{h}(\mathbf{x})$
- This form was shown to be amenable to Inverse dynamics approach.(Stephen Lane, 1988).

Failure representation :

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x},\mathbf{u}) + \mathbf{G}(\mathbf{x})\mathbf{u} + \mathbf{w}(t) + \mathcal{F}_1(\mathbf{x},\mathbf{u},t)$$

$$\mathbf{y}(t) = \mathbf{h}(\mathbf{x}) + \mathbf{n}(t) + \mathcal{F}_2(\mathbf{x},\mathbf{u},t)$$

Ideally, we would like to isolate $\mathcal{F}_1$ and $\mathcal{F}_2$ and base our failure analysis on these estimates - Extended Kalman Filter as a first step.

Princeton University

All the models considered till now are linear time invariant. When applied to aircraft control, even though these models are adequate for transport aircraft under nominal conditions, a failure may cause the linear model to be quite inaccurate. Since these situations are critical, poor failure accommodation due to linearization is an inherent danger. Since at present sufficient tools are not readily available, we started with robust linear schemes. Use of non-linear models during reconfiguration has the additional advantage of exploiting non-linear controllability where applicable.

To start with, we can formulate the non-linear model in the form of the second and third equations shown. This form on further modification is amenable to no-linear inverse dynamics (NID) control.

These models also require non-linear estimation and as a first step, we use an Extended Kalman Filter (EKF) with full non-linear state and covariance propagation and linear update for filter gain calculation. Methods to extract $\mathcal{F}_1$ and $\mathcal{F}_2$, which represent unmodeled forces and torques, on-line are to be devised in future.

Failure Accommodation

- **Sensor Failures :**
  - Redundancy to be exploited.
  - Eliminate faulty sensor.
- **Actuator Failures :**
  - For identified failure, change in G is known.
  - A Proportional Integral Implicit model following law can be used (LTI).
  -Alternately, a control law based on ***Non-linear Inverse Dynamics*** can be used.

Here we have a non-linear system :

$$\dot{x} = f(x,u) + G(x)u$$

$$y = h(x) ; \qquad v = y^{(d)} = f^*(x) + G^*(x)u$$

$$u = -\left[G^*(x)\right]^{-1} f(x) + \left[G^*(x)\right]^{-1} v$$

Since this decouples the inputs and outputs, recovery to pre-fail status is possible.

Once a failure is identified, corrective action for expected failures is straight forward. Bad sensors can be eliminated and replaced by their redundant (physical or analytical) counterparts. If the linear model is still valid and the system still controllable, a proportional-integral implicit model following control can bring the desired reconfiguration. Alternatively a non-linear inverse dynamics approach which essentially decouples the required inputs and outputs as shown briefly in the equations can be used to attain close to pre-failure performance.

On the other hand, if the failure is unanticipated or full controllability is not assured, a combination of declarative logic and part controllability could be used. If unmodeled forces and torques are available, changing the non-linear model or corrected command inputs to NID control is another way.

Issues to be dealt with NID control

- Not every non-linear model of aircraft is realizable through inverse dynamics.
- Triangular form of non-linear dynamics expected:

(neglect unimportant/ weak controls)

$$\dot{x}_1 = A_1(x_1, x_2)u_1$$
$$\dot{x}_2 = A_2(x_1, x_2, x_3, u_1) + B_2(x_1, x_2, x_3, u_1)u_2$$
$$\dot{x}_3 = A_3(x_1, x_2, x_3, u_1, u_2) + B_3(x_1, x_2, x_3, u_1, u_2)u_3$$

- Use of EKF for state estimate rather than full state feedback.
- Specific application to FDIR is new.
- No proper theory for stabilization with non-linear dynamics.

There are certain restrictions that apply to NID control. Not every non-linear model is amenable to this procedure. In past studies a particular form of representation - the triangular form shown above - was necessary and certain insignificant controls were judiciously neglected. Further, singularities in G* should be avoided.

There are other issues such as inclusion of EKF instead of full state feedback and stability of the inverse system for non-linear models which have to get more attention in future.

Summary and Future Research

- An Intelligent framework for FDIR proposed.
- Use of full non-linear dynamics of the aircraft suggested.
- Issues for research:
  - Simulate and compare robust linear v/s non-linear based dynamics and control.
  - Develop methods to approximate non-linear faults. (unknown forces & torques)
  - Evaluate applicability of NID in failure tolerant control.

Main thrust in future is research toward working with the full non-linear aircraft model from detection to reconfiguration. Tradeoffs should be made somewhere due to real-time implementation and evaluation of robust linear models and non-linear methods in the presence of model uncertainty should be made. Non-linear approximation schemes for unmodeled forces and torques are to be devised. Incorporation of learning for such schemes will be studied.

The Use of Radial Basis Functions and Genetic Algorithms in Neural Networks

Ohio University,
June 22, 1995

Kristina A. Richardson

Our current research is on the use of neural networks for system identification of an aircraft lift coefficient. The use of radial basis functions as activation functions within neural networks offers improved convergence times due to their ability to produce a significant output throughout the training period by optimizing the function center and width. The use of a genetic algorithm to find the optimal set of network weights and biases is also demonstrated.

Research Objectives

Develop

- Improved training methods

Investigate

- Hybrid activation function and function error gradient-training
- Use of Extended Kalman filters, genetic algorithms, and Monte Carlo simulation
- Parallel processing

Long-term research objectives include developing improved training methods for neural networks. To facilitate this, we plan to extend Dennis Linse's work with hybrid activation functions and function error gradient-training. We will also investigate the use of Extended Kalman filters, genetic algorithms, and Monte Carlo simulation for determining optimal network weights and biases. The reduction in neural network training times offered by parallel processing will also be explored.

The Model

Non-dimensional lift coefficient

$$C_L(\alpha, q, \delta_E) = C_{L_{ST}}(\alpha) + C_{L_q}\frac{q\bar{c}}{2V} + C_{L_{\delta_E}}\delta_E$$

where,

α = angle of attack (-5°, 19°)
q = pitch rate (-20°/sec, 20°/sec)
δ_E = elevator deflection (-20°, 20°)
C_{LST} = static lift coefficient
C_{Lq}, C_{Ld_E} are constant lift derivatives
$\bar{c}$ = mean aerodynamic chord of the wing (11.2 ft)
V = constant flight air speed (400 ft/sec)

Lift derivatives from TCV manual, slope of the $C_{L_{ST}}(\alpha)$ curve from TCV data at the given flight condition.

The system identification problem is the nonlinear non-dimensional lift coefficient for the Boeing 737. This coefficient is a function of angle of attack, pitch rate, and elevator deflection. $C_{L_{ST}}$ is the baseline contribution due to angle of attack. The next two terms include the contributions due to pitch rate and angle of attack. A plot of the function is included in the next slide.

Laboratory for Control and Automation

Two Input Training Function

100 data points

input[α,0,δ_E]

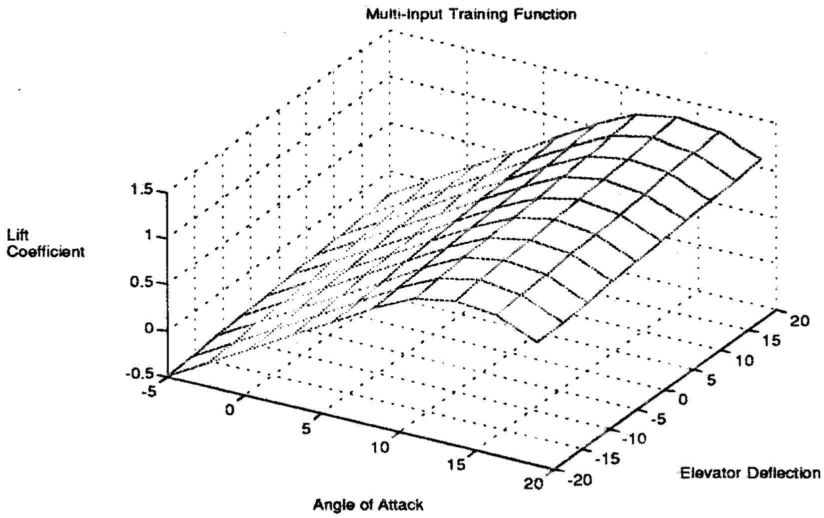

Princeton University 4

This is a plot of the lift coefficient versus angle of attack and elevator deflection. Angle of attack is varied from -5 degrees to 19 degrees while the elevator deflection is varied from -20 degrees to 20 degrees. Pitch rate is maintained at a constant 0 degrees/sec.

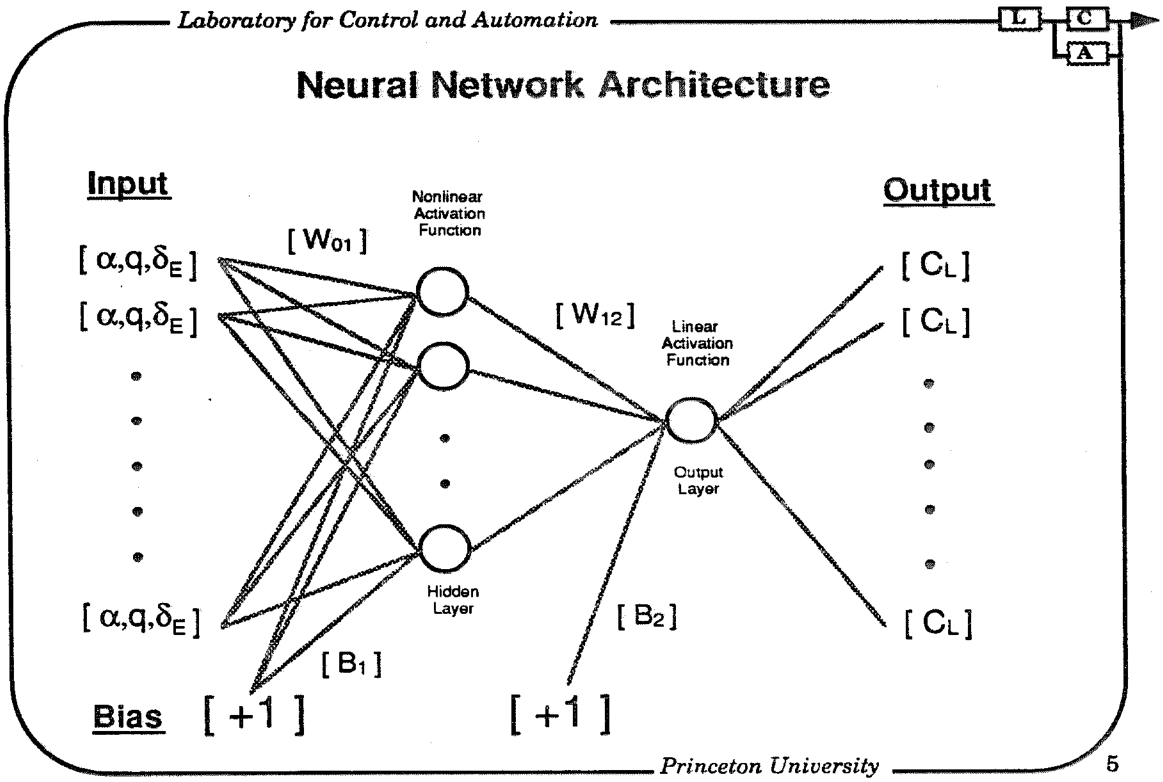

In our research we deal with a feedforward neural network architecture with one hidden layer. Neural networks include a weighted interconnection between layers. These weights must be optimized to produce the desired output. The optimal weights may be found through a training algorithm or an optimization algorithms. The key to the networks ability to make an approximation of a nonlinear function is the nonlinearity of the activation functions within the hidden layer.

Backpropagation Training Algorithm for Determining Weights and Biases

- Backpropagation, or the "generalised delta rule" was suggested in 1986 by Rumelhart, McClelland, and Williams

- Uses the differentiability of the activation function and the error between the target output and the net output to adjust weights and biases in the hidden layer (i.e. a gradient search technique)

- The error is "backpropagated" through the network to the hidden layers

The backpropagation training algorithm revitalized the neural network industry after its publication in 1986. It uses the differentiability of the non-linear activation function to 'train' the weights. Typically the activation function is chosen to be sigmoidal, hence the network is infinitely differentiable. Knowledge about the error is only explicitly available to the output layer during training. The backpropagation algorithm passes this knowledge back to the weights in the hidden layer.

Radial Basis Function Network

Consider a network with the hidden nodes computed as follows:

$$O_{pi} = \psi(\| net_p - c_i \|)$$

Where O_{pi} is the output of the hidden node i in response to the p_{th} input vector, net_p, and c_i represents the centre of a radially symmetric function, ψ.

ψ is chosen as Gaussian,

$$\psi(\| net_p - c_i \|) = \exp\left\{ -\sum_j \left[\frac{\| (net)_p - c_{ij} \|^2}{2\sigma_{ij}^2} \right] \right\}$$

where c_{ij} and s_{ij} are the center and width of the function respectively, $(net)_p$ is the jth input vector

In a radial basis function network, the non-linear activation functions within the hidden layer of the network are radial basis functions. This function is typically chosen to be Gaussian. Faster network training may be achieved by optimizing the centers and widths of the activation functions during training. Methods for optimization include clustering algorithms, least-squares methods, and orthogonalization.

Use of Genetic Algorithms for Determining Weights and Biases

A **genetic algorithm** is a search algorithm based on the mechanics of natural selection and natural genetics

1) **Generate** a random initial population of vectors

$$N_{pop} \begin{bmatrix} 1 & 0 & 0 & 1 & 1 & 0 & 1 & 1 & \ldots & 0 & 1 & 1 & 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 1 & 1 & 1 & 0 & 0 & \ldots & 1 & 0 & 0 & 1 & 0 & 1 & 0 & 1 \\ . & . & . & . & . & . & . & . & \ldots & . & . & . & . & . & . & . & . \\ 0 & 1 & 1 & 1 & 0 & 1 & 0 & 1 & \ldots & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & \ldots & 1 & 1 & 1 & 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 1 & 1 & 0 & 0 & \ldots & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

$$W_{12} \quad \ldots \quad B_2$$

2) **Evaluate** the performance metric with each row of weights and biases

Genetic algorithms are search algorithms based on the mechanics of natural selection and natural genetics. The initial step includes generating a random initial population of 1's and 0's in vectors. Every 8 bit segment in a vector represents a network weight. The binary number is converted to a real number and the performance of the network with each vector of weights is evaluated.

Perform the following steps $N_{pop}/2$ times:

3) Select two strings with the "best" performance metric using a probabilistic selection process, specifically tournament selection.

4) Create new strings with members from the selected strings by performing "crossover", swapping the heads and tails at a randomnly selected crossover point

A new population is now in place, the next step is optional:

5) Mutate (0 --> 1, 1 --> 0) a randomly selected bit with the string if the mutation probability is exceeded

Return to step 2) Evaluation until performance requirement has been satisfied

Tournament selection is then employed to select two random strings with the "best" performance. Two new strings are then created by performing "crossover", swapping the heads and tails, at a randomly selected crossover point. Once a new population is in place mutation may be implemented. If a mutation probability is exceeded a randomly selected bit within a vector is switched to a 0 if a 1, and to a 1 if it is a 0. Once this population is in place, return to the evaluation stage.

Real Number Genetic Algorithm

Real numbers may be used in place of "bits" in a genetic population.

N_{pop}

−0.7856	1.5432	0.5456	−2.7654	. . . .	−2.1395	−0.9320	2.7023	1.2391	−0.9392
0.5638	0.1284	19891	−0.9765	. . . .	−1.2437	1.8936	0.2902	1.6721	0.3159
.	.	.	.	. . . .	.	.	.	.	.
0.2134	1.3452	−0.9280	−2.3451	. . . .	0.3015	2.3109	−1.3926	0.6739	−2.4934
W_{01}	W_{02}	W_{03}	W_{04}	. . . .	W_{41}	B_{11}	B_{12}	B_{13}	B_2

The mutation step then involves incrementing a random member of a string by some amount if a mutation probability is exceeded.

A real number genetic algorithm employs real numbers in the place of bits. The steps remain the same until mutation. the mutation step involves incrementing a random member of the vector by some amount if the mutation probability is exceeded.

Conclusions

5 nodes in hidden layer, final performance metric 0.1

method	floating point operations
sigmoidal, backpropagation	7.9×10^9
RBF, backpropagation	1.9×10^9
8-bit genetic algorithm	40.0×10^9
real number genetic algorithm	8.2×10^9

- Locally tuned radial basis activation functions can converge an order of magnitude faster than sigmoidal backpropagation networks and are **also** universal approximators.

- Genetic search algorithms are able to solve the weight and bias optimization problem and are global but can be computationally expensive.

- Real number population genetic algorithms seem to be better suited to this problem and perform faster than bit genetic algorithms

Backpropagation with sigmoidal activation functions, backpropagation with radial basis activation functions, and genetic algorithms with both bits and real numbers were employed to solve the system identification problem presented earlier. All networks had 5 nodes within a single hidden layer. In terms of floating point operations the locally tuned radial basis function network performed most efficiently. Genetic algorithms were able to solve the problem and are global algorithms but can be computationally expensive. The 8-bit algorithm performed the binary to real number conversion on every iteration, hence it computational complexity.

Future Work

- Continue to investigate validity of previously developed algorithms for training neural networks
- Investigate training of neural networks using stochastic simulation and search
- Introduce massively parallel computation in the training of neural networks
- Compare our new results against existing methods

Future work includes continuing to apply previously developed network training methods to the aircraft dynamic system identification problem. Employ stochastic simulation and search methods to the weight optimization problem. Research advantages of massively parallel computation in the training of neural networks.

OHIO UNIVERSITY

INVESTIGATION OF AIR TRANSPORTATION TECHNOLOGY AT OHIO UNIVERSITY 1994-1995

Robert W. Lilley
Avionics Engineering Center
Department of Electrical and Computer Engineering
Ohio University
Athens, Ohio

The Joint University Program (JUP), in its twenty-fourth year, continues to provide students, faculty, and staff of the Avionics Engineering Center the opportunity to conduct basic research dedicated to improve the National Aerospace System. The 1994-1995 highlights include the awarding of an eighth RTCA Jackson Award along with three M.S. degrees for participants in the program.

- Dr. David W. Diggle won the 1994 RTCA Jackson Award with his doctoral dissertation entitled, "Interferometric GPS Flight Reference/Autoland System." This work, initiated under the Joint University Program, resulted in the first real-time kinematic GPS autoland and has been expanded to aid in determining the feasibility of GPS for CAT III landings.

- AnnMarie Bizek Fink completed the M.S. program with her thesis entitled, "Investigation of Selective Availability in the NAVSTAR Global Positioning System." This research, done with Dr. Michael Braasch, investigated various SA models and further characterized this, the largest of all GPS error sources.

- Lap Ho, whose work was critical to the GPS autoland research, completed the M.S. program with his thesis entitled, "High Precision Short-Baseline Pointing System Using GPS Interferometry." This work describes the implementation and coding of many of the algorithms used in GPS interferometry.

- Sanjiv Koshal completed the M.S. program with his thesis entitled, "Hybrid System GMSK Digital Receiver Implementation in Real-Time." This project, which followed past JUP research in adding a data link to the existing VHF AM voice channels, focused on the receiver design and continued to show the potential of this unique modulation method.

With the increasing importance of GPS in the civil aviation community, the focus of many of the current research projects are GPS-related. Specifically, the effect of diverse GPS receiver architectures on system performance is being studied using both simulation techniques and real-time hardware. GPS/INS integration improvements are being researched and implemented along with a novel GPS/ILS integration. The issues associated with an automated landing control system using DGPS are also under investigation. Lastly, a TV technology data link has been proposed to fill the void for a data uplink to aircraft and is currently under design.

In addition to the quality research obtained through the support of the Joint University Program, it is also important to recognize the training and experience the program provides to the next-generation of engineers/scientists who will work in the aviation industry.

BIBLIOGRAPHY OF 1994-1995 PUBLICATIONS

1. Fink, A. B.: Investigation of Selective Availability in the NAVSTAR Global Positioning System. M.S. Thesis, Ohio University, Department of Electrical and Computer Engineering, Athens, OH, Nov. 1994.

2. Ho, L.: High Precision Short-Baseline Pointing System Using GPS Interferometry. M.S. Thesis, Ohio University, Department of Electrical and Computer Engineering, Athens, OH, March 1995.

3. Koshal, S.: Hybrid System GMSK Digital Receiver Implementation in Real-Time. M.S. Thesis, Ohio University, Department of Electrical and Computer Engineering, Athens, OH, Aug. 1994.

TELEVISION TECHNOLOGY DATA LINK

Damon E. Gura
Ohio University
Athens, Ohio

SUMMARY

The motivation for a television data link comes from the need to communicate and display aeronautical hazard information to pilots and controllers. The proposed data link work includes the design and prototype of complete transmitting and receiving units. So far the transmitter unit has been designed and tested. Therefore the transmitter will be the focus of this paper.

INTRODUCTION

It is important to supply real time weather data to pilots in order to avoid weather-related accidents. The proposed television-technology data link has the benefit of offering high capacity information transfer using commercial off-the-shelf (COTS) components. The reason that COTS components can be used is that the format for data transmission is the commercial television signal using standard TV channels. Through a small modification of the commercial TV signal a 30-frames-per-second data link is realized. The main idea is to combine single-frame graphics materials relating to hazards and warnings into the commercial 30-frames-per-second format suitable for standard TV transmission.

COMMERCIAL TELEVISION SIGNAL STRUCTURE

Before the transmitting unit can successfully be described it is important to review some of the standard TV signal structure. Please note that only relevant parts of the signal structure will be reviewed and the reader is referred to [1,2] for more information.

The National Television System Committee (NTSC) color TV system consists of three main signals. These three signals are formed from a linear combination of the red, green, and blue components of a particular scene. $M_y(t)$ referred to as the luminance signal corresponds to the gray scale or brightness of a scene. $M_i(t)$ or the in-phase component is combined with $M_q(t)$ or the quadrature component to form the chrominance signal which contains all of the color information for the particular scene.

The actual picture takes $^1/_{60}$ of a second from the top of the picture to the bottom and back to the top again. During this time 262.5 lines are transmitted, this is called a field. Since each field contains a half line, the next field lies in between the first field as shown in figure 1; the two fields are said to be interlaced. The complete picture is called a frame and consists of 525 lines. Interlacing is done to avoid flicker while accurately portraying motion.

SYNCHRONIZATION AND BLANKING

During the horizontal and vertical retrace, the picture tube is made inoperable by means of blanking pulses. Also for picture stability the scanning process at the transmitter and receiver must be synchronized. There are two types of sync and blanking pulses, horizontal and vertical. Horizontal sync and blanking occur at the end of each line while vertical sync and blanking occur at the end of each field. The composite video signal then contains video, blanking, and sync [3]:

$$M_c(t) = \begin{cases} M_s(t) & \text{during the sync interval} \\ M_y(t) + \mathrm{Re}\{g_{sc}(t)e^{j\omega_{sc}t}\} & \text{during the video interval} \end{cases} \quad (1)$$

where

$$g_{sc}(t) = [M_i(t) - jM_q(t)]e^{j33^\circ}$$

DATA LINK TRANSMITTER CONCEPTS AND IMPLEMENTATION

Figure 2 summarizes the original concept behind the television data link transmitter and receiver. The transmitter portion of the data link is depicted in the bottom portion of the figure and has been outlined. The data link transmitter operates as follows: First, standard video input is gathered and input into a switching device that allows easy access to any of the given input channels. Once an input is selected then a frame of video data is captured into its own particular frame store as labeled in figure 2. Thirty different frame store devices would result in 30 different frames of video data containing hazard and warning information. After the individual framed information has been captured and stored it is then multiplexed together and sent to a standard off-the-shelf television transmitter.

The implementation of this process has led to the replacement of the entire data link transmitter save the television transmitter and antenna with a slightly modified computer. This leads to a self contained transmitting unit using only a personal computer. The modifications done to the computer are increased memory, needed to handle video, and two video manipulation cards that serve to capture and store the video before transmission.

The video cards have different purposes; one is a 32-bit videographics card used for editing and overlay of captured frames while the other is a compression/decompression card for file size reduction and transfer speed. Some features of the videographics card are as follows:

1) Frame resolution of 720x486 pixels with 32 bit color on each pixel.
2) Up to three different standard video sources.
3) Includes a comprehensive development library compatible with most popular programming environments.

The compression/decompression card does automatic optimal compression on each incoming frame before storing it on the hard drive. The compression standard used is the Joint Photographic Expert Group (JPEG). JPEG works in the low frequency domain and uses the discrete cosine transform. Each frame is broken down into 8-by-8-pixel blocks and the discrete cosine transform is applied to each block. High frequency components are dropped and the rest of the information is Huffman coded [4]. The compression rate is directly related to the data transfer rate of the system. Transfer rates from 200 KBytes/sec to 2.5MBytes/sec are attainable.

Manipulation of the incoming video is important for many reasons. One obvious reason is to make the system versatile by giving the user freedom to edit the captured video easily. Another reason is for purposes of frame detection on the receiving side of the data link as seen in figure 2. The pilot or controller would command an input frame and expect only that commanded frame to appear on the display. Some type of encoding scheme must be implemented in the transmitter in order to control this receiver function. The encoding scheme uses $M_y(t)$ or the luminance signal to insert black and white bars of specific length and width, which can be thought of as bits, in the upper left corner of the captured frame. These bars represent a block code of the frame number to be transmitted. This was found to be a favorable scheme because of the ease of manipulating graphical data on the frame as well as the fact that white is the peak amplitude of the standard composite video signal while black is the lowest encountered value (except for synchronization levels).

The encoding scheme as well as the general operation of the transmitting portion of the data link has been integrated into a computer program. This program was coded using Visual C/C++ and resides on the transmitting PC. Figure 3 shows a screen capture of the actual program, allowing the reader to see some of its features. The program does the capture and storage, plus output control of the transmitting unit. In addition, the program encodes the black and white bars discussed earlier into the composite signal and inserts the actual frame number into the bottom right corner of the frame for verification on the receiving end.

CONCLUSIONS

The television technology data link uses the standard commercial television signal to achieve a low cost, high capacity visual method of information transfer. The data link consists of two parts, the transmitter and receiver both of which will be prototyped and demonstrated. The data link transmitter has been demonstrated to various audiences and hence was the focus of this paper. Current research involves the implementation of the receiving unit and finally the demonstration of the complete system. It is hoped that this new data link provide for the transmission of real time weather data to pilots in order to avoid weather-related accidents that have occurred and possibly serve to avoid them in the future.

ACKNOWLEDGMENTS

This work was supported by the Federal Aviation Administration's Aeronautical Hazards Program (FAA ARD-230) and is done in conjunction with other work that is supported by the National Aeronautics and Space Administration through the Joint University Program for Air Transportation Research. The author wishes to express his thanks to Dr. Robert Lilley for his support and guidance.

REFERENCES

[1] Terman, Frederick E., *Electronic and Radio Engineering*, McGraw-Hill Publishing Company, New York, 4th ed., 1955.

[2] Fink, Donald G., *Principles of Television Engineering*, McGraw-Hill Publishing Company, New York, 1940.

[3] Couch, Leon W., *Digital and Analog Communication Systems*, Macmillan Publishing Company, New York, 4th ed., 1993.

[4] Stafford, R. H., *Digital Television*, John Wiley & Sons, Inc, 1980.

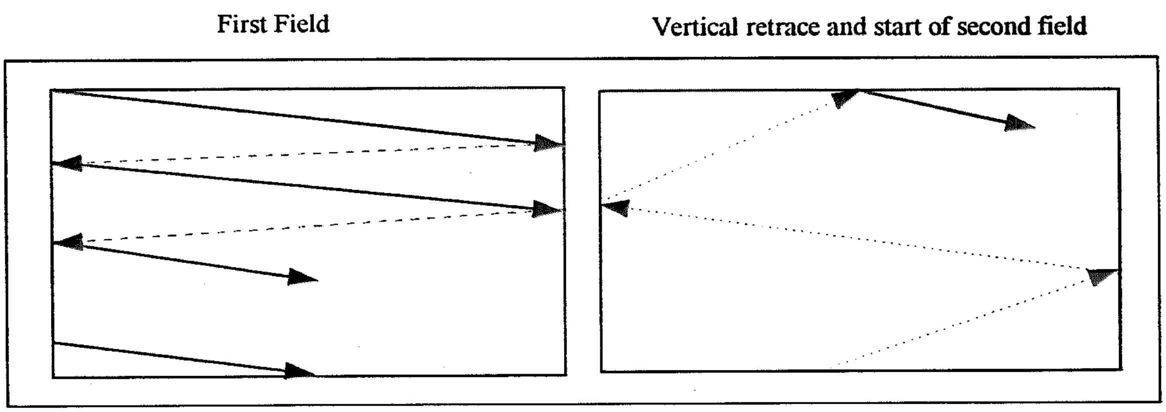

FIGURE 1: Standard Television Interlace Method

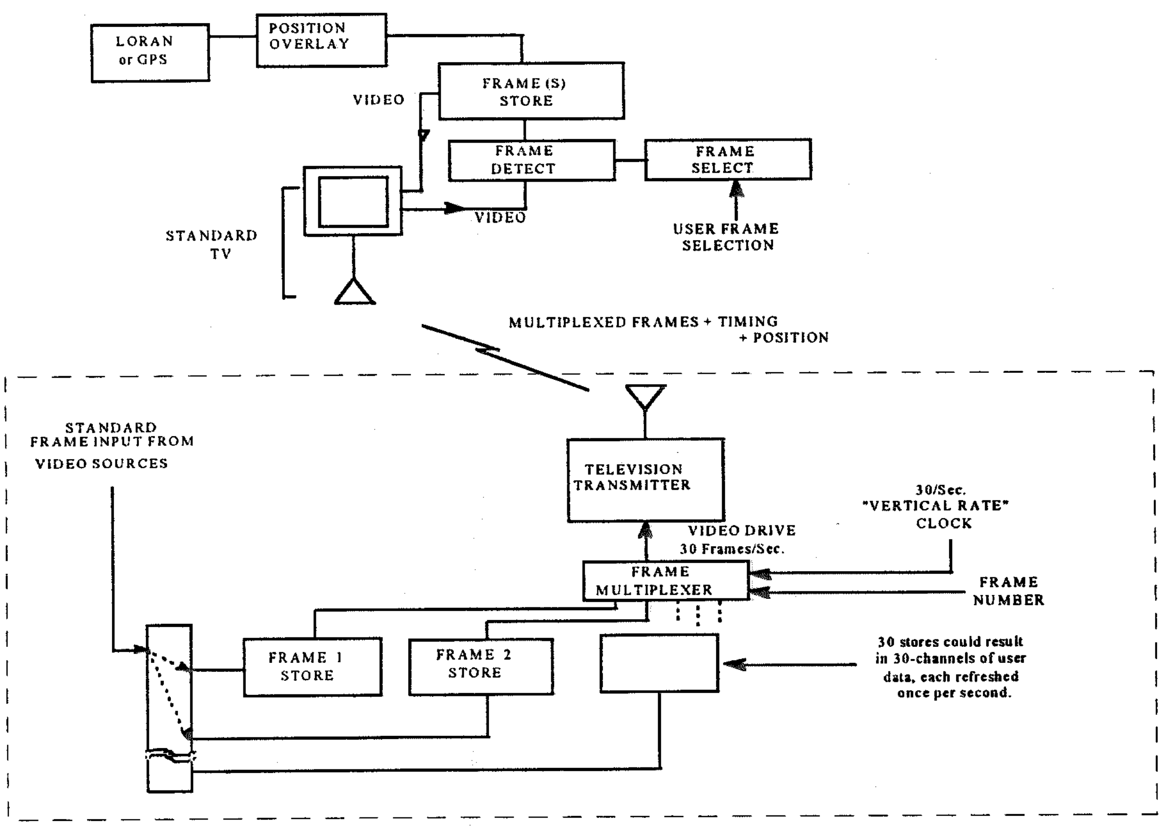

FIGURE 2: Conceptual Block Diagram of Transmitter and Receiver

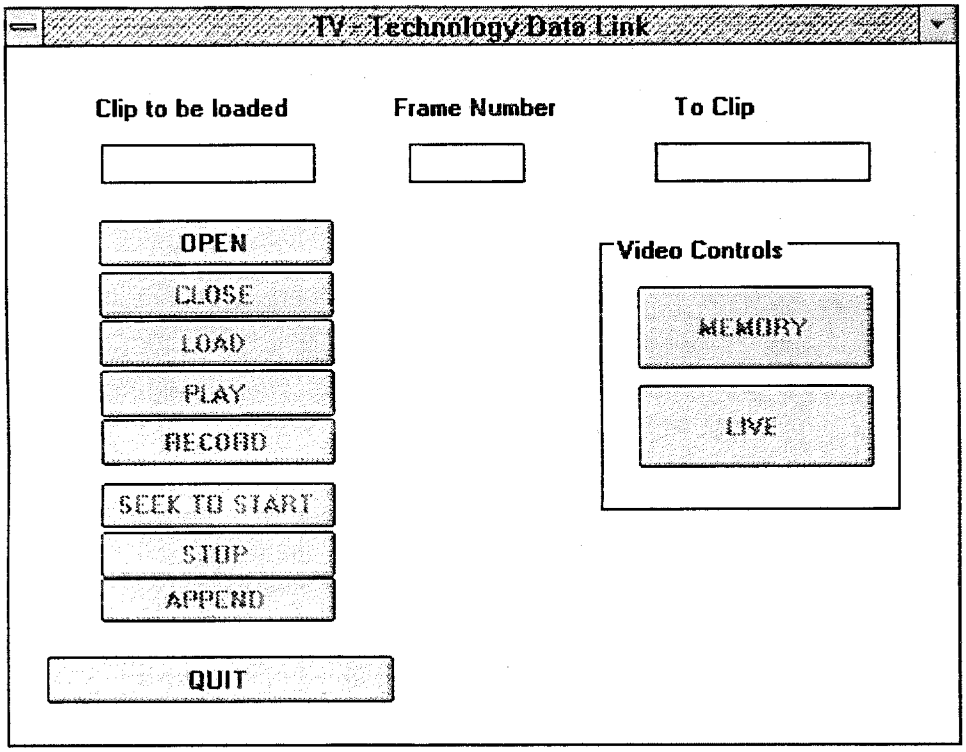

FIGURE 3: Screen Capture of Transmitter Program

DGPS/ILS INTEGRATION FOR AN AUTOMATIC AIRCRAFT LANDING SYSTEM USING KALMAN FILTERING

Steven J. Hill
Avionics Engineering Center
Department of Electrical and Computer Engineering
Ohio University
Athens, OH

SUMMARY

The Instrument Landing System (ILS) has served the aviation community well in the past 40 years that it has been in existence. It is very apparent, however, that due to the increase in air traffic and radio interference, that the replacement of ILS is nearly inevitable. A possible alternative is the Global Positioning System (GPS). It shows great promise in providing navigation for all phases of flight including precision approach and landings. Completely throwing away an older system in exchange for a new is not a prudent or cost effective decision. With this in mind, it would be wise to investigate a way to integrate the two systems and exploit the strong points of each. A GPS/ILS integration might be a possible solution.

BACKGROUND

The problem plaguing current ILS installations has to do with the localizer portion of the system. It provides lateral guidance for the aircraft. Vertical guidance is provided by the glideslope. The specific problem is the radio frequency interference with the localizer. It transmits on frequencies between 108 and 112 Mhz. These bands are shared with radio and television broadcasting stations which have one or two orders of magnitude more transmitting power than ILS. This interference makes it difficult to maintain current systems and nearly impossible, in some areas, to install any new ILS systems.

These problems are severe enough that FAA and ICAO are looking for alternative systems to replace ILS. There are two systems currently in development. They are the Microwave Landing System (MLS) and the Global Positioning System (GPS). The Microwave Landing System (MLS) was scheduled to replace ILS at all international airports by January 1, 1998. This was rescinded by ICAO at the Special Communications/Operations Divisional Meeting in Montreal in April 1995. The reasoning behind this was the prospect of using Global Navigation Satellites Systems (GNSS) instead of MLS for precision approach and landing. The (FAA) has already terminated all its contracts with companies for any further development of MLS. Other European countries still have an interest in MLS to replace ILS systems at selected airports.

The United States is actively pursuing GPS as the replacement for ILS. The only problem is that it will be some time [in the author's opinion] before GPS is completely certified to meet RNP standards for precision approach and landing. Since ILS is already certified, integrating it with GPS could create a RNP certified system that utilizes GPS. This integration serves as a transitionary system until GPS can be certified for precision approach and landing.

SYSTEM INTEGRATION OVERVIEW

The ILS glideslope will be used in combination with GPS measurements to calculate a position for the aircraft. The glideslope has very good vertical accuracy and does not suffer from radio interference like the localizer. The localizer will not be utilized in this integration.

GPS for the civilian user, has horizontal accuracies of 100 meters 95% of the time and 300 meters 99.9% of the time. This positioning accuracy is provided by the Standard Positioning Service (SPS) of GPS. SPS does not meet Required Navigation Performance (RNP) standards for accuracy. The Precise Positioning Service (PPS) of GPS, provides accuracies of 2-5 meters. Unfortunately, PPS is not accessible by the normal civilian user. A method must be devised that can use GPS in such a way that RNP criteria for lateral guidance will be satisfied. Integrity, continuity-of-service and availability are other ongoing research areas in GPS. These will not be discussed at this time. The focus is placed on obtaining an accurate aircraft position without PPS.

To improve the accuracy of GPS, different solutions have been devised that can position a moving aircraft in space with 1-2 meter accuracy. In some cases even submeter accuracies have be obtained. These methods fall under the term Differential GPS (DGPS). DGPS in its generic form, uses a receiver and datalink placed at an exact know surveyed point near the runway. The receiver "knows" its true position and can make corrections to the GPS measurements (ionospheric errors, clock biases, etc.). If the aircraft is within a radius of approximately 10 kilometers of the ground station, these corrections are valid and can be applied to the aircraft's GPS measurements as well. The corrections are uplinked to the aircraft and a position is calculated with approximately 1-2 meter accuracy. If DGPS is used, RNP criteria are met.

The DGPS lateral positioning will be used for the integration with the ILS glideslope. The DGPS vertical position will augment the glideslope and help in estimating the glidepath angle bias. The ILS Glideslope and computed DGPS measurements are not of the same type. The Glideslope provides angular deviations from a nominal glidepath angle. DGPS provides position and velocity measurements. To attain a position an Extended Kalman Filter is employed.

A Kalman Filter is an optimal state estimator. It was originally used in state variable control theory to optimally estimate state variables. In recent years, it has been heavily utilized for a multitude of integrated navigation systems. The filter takes measurements from different sources and optimally estimate states such as position, velocity, acceleration, etc. The filter performs well if the measurements are not too noisy and there are more measurements than states being estimated. Background for the Extended Kalman Filter can be found in [5].

IMPLEMENTATION

It is fortunate that ILS and GPS data have previously been recorded. This data will be used in simulating approaches and to judge the performance of the DGPS/ILS integration. The data was collected from the FAA/Ohio University/UPS Autoland Flight Tests done at the FAA Technical Center at Atlantic City, New Jersey [6]. The tests successfully demonstrated the ability of DGPS to perform precision approaches and landings within RNP accuracy standards.

The data collected consisted of GPS data from four different receivers. Ashtec Z-12 and NovAtel GPS receivers at both the ground station and in the aircraft. GPS time-tagged ILS data was also recorded. The ILS was GPS time-tagged so that all of the data could be matched during post-processing.

For the simulation, the Ashtec data will be assumed as the true position. The Ashtec Z-12 is capable of providing the needed accuracy by using PPS. Interferometry techniques employed in the Ashtec post-processing software allow a very accurate flight path to be generated. The path is accurate to within 10-20 centimeters. During simulation, the performance of the filter will hopefully follow this path closely.

The data collected by the NovAtel receivers will be used in the integration with the time-tagged ILS data. The simulation for the Kalman Filter will be written in C. The output position at each time step will be compared with the truth data in Matlab. Originally the entire simulation was to run in Matlab, but after initial testing it was obvious that the amount of data involved would not permit the use of Matlab. A simple Least Mean-Square integration has been implemented in Matlab. The initial results exhibit that it is possible to get a very accurate position using the ILS glideslope and only 3 satellites with differential corrections.

CONCLUSION

The NovAtel data is currently being processed for used in the Kalman Filter. At the time of this writing, the Kalman Filter equations and matrices associated with it had not been derived completely. More specifically, the equations to propagate velocity using carrier phase measurements have not been finished. After the equations are finalized, the filter will be implemented in C.

REFERENCES

[1] van Graas, Frank and Lee, Shane-Woei., *High-Accuracy Differential Positioning For Satellite-Based Systems Without using Code-Phase Measurements*. Proceedings of the National Technical Meeting of The Institute of Navigation, Anaheim, CA, January 18-20, 1995.

[2] Klass, Phillip .J., *ICAO Acts to Drop MLS Mandate*. Aviation Week & Space Technology, April 17, 1995.

[3] Bose, Keith W., *Aviation Electronics*, IAP Inc., Casper, WY, 4th ed., 1983.

[4] Fried, W.R. and Kayton, M., *Avionics Navigation Systems*. John Wiley & Sons, New York, 1969.

[5] Brown, R.G. and Hwang, P., *Introduction to Random Signals and Applied Kalman Filtering*. John Wiley & Sons, New York, 2nd ed., 1992.

[6] van Graas, F., D.W. Diggle, V. Wullschleger, R. Velez, M. Dimeo, G. Kuehl, and R. Hilb, "FAA/Ohio University/UPS Autoland Flight Test Results," Proceedings of the National Technical Meeting of The Institute of Navigation, Anaheim, CA, January 18-20, 1995.

CURRENT RESEARCH

For the integration, a low cost Systron-Donner 'Motion Pack' IMU is used with Novatel's GPS card. The IMU has an accelerometer noise < 0.002 m/s^2 over one second which transforms into a position error of 1 mm over one second and a gyro noise < 0.01 degrees/s which gives an error of 1.7 cm at a speed of 100 m/s over one second.

The cycle-slip detection requires inertial velocities which are accurate within 0.2 m/s during 0.5 second interval. This level of integration requires a synchronization between GPS and the IMU to the order of microseconds, i.e. the validity of GPS data and all the inertial data should be within a few microseconds.

A single board DSP data acquisition system is used for the synchronization and data collection. The board is based on Texas Instruments TM320C40 processor, and has a capability of simultaneous triggerable eight channel A/D converters. The GPS receiver gives out 1 pulse per second (pps) indicating that the GPS data is valid at that instant. This pulse is used to trigger all A/D converter 10000 times. A set of 100 A/D samples are averaged to form a total of 100 inertial samples per second per channel. The GPS pulse marks the beginning of the conversions each second, the samples are then averaged and stored.

CONCLUSION

This paper gives the background on the GPS/INS integration requirements for the currently running Ohio University projects.

REFERENCES

1. Van Graas, F.: Personal discussion with the author.
2. Farrell, J. L. : Personal discussion with the author.