AIAA Guidance, Navigation, and Control Conference and Exhibit
11-14 August 2003, Austin, Texas

AIAA 2003-5762

Development of the F-117 Flight Control System

Robert Loschke*
Lockheed Martin Aeronautics Company
1011 Lockheed Way, B-608, P-10, D-6C3P
Palmdale, CA 93599-2526

* Retired Technical Fellow

ABSTRACT

The very unusual F-117 configuration has an equally unusual combination of aerodynamic instabilities, cross axis coupling, and directionally destabilizing exhaust nozzle effects that posed some unique challenges to the design of the Flight Control System (FCS). Radar and infrared signature considerations dictated the highly swept wings, swept canted fins, and the high aspect ratio engine exhaust nozzles. Installing two large internal weapons bays, the engines, exhaust nozzles, mission avionics, and fuel tanks into the defined planform resulted in the center of gravity (CG) being relatively far aft. The aft CG was desirable because cruise trim drag and the radar signature were minimized with essentially zero elevon deflection. Aerodynamic instabilities within the desired operational flight envelope were deemed acceptable. Although the F-117 design mission was precision bombing at night, the customer wanted the aircraft to meet the maneuvering and flying qualities requirements of a Fighter/Attack aircraft. To safely meet these requirements, the FCS had to provide a high level of stability augmentation and have the control power authority through the two fins and the four elevons to prevent departures during large amplitude maneuvering. Consequently, a full authority Fly-By-Wire FCS originally developed for the F-16 was modified for use in the F-117. New control laws were designed to provide the good flying qualities of a classic conventional aircraft and to allow the F-117 to be rapidly and safely maneuvered within the permissible flight envelope. The FCS also required conventional auto pilot relief modes and automatic navigation so that the pilot could concentrate on the tactical situation during a mission. The first flight of the F-117 prototype took place in June 1981 thirty-one months after program go ahead in November 1978. Refinement of the FCS continued in parallel with the development testing of the weapon system avionics and allowed limited Introduction of Operational Capability (IOC) to be achieved in October of 1983. High Angle-Of-Attack flight testing was completed after IOC and the fully developed Flight Control Computer was retrofitted into the fleet starting in 1986. The use of streamlined project management and the close coordination of customer and contractor teams during design, simulator evaluations, Iron Bird testing, and flight testing allowed the FCS to be developed quickly and at a relatively low cost.

1. INTRODUCTION

On 01 December 1977, the Have Blue low observable technology demonstrator aircraft made its first flight. The Have Blue program was a collaboration between the Defense Advanced Research Projects Agency, the United States Air Force (USAF), and the Lockheed Skunk Works to demonstrate those technologies which could reduce the vulnerability of combat aircraft to integrated air defense systems. (See reference 1 for details of the Have Blue Program history.) The Have Blue flight tests showed that a faceted aircraft with highly swept wings could be flown and that the Radar Cross Section (RCS) measured in flight was orders of magnitude smaller than that of conventional aircraft. Once it had been shown that a very low RCS aircraft was feasible, the USAF requested the Skunk Works to determine the best way to use these new technologies for a new attack aircraft. These studies led to the configuration that eventually became the F-117 and the Skunk Works submitted a proposal for a full-scale development program in February 1978. In November of 1978, the USAF issued a contract to the Skunk Works to develop the aircraft as rapidly as possible.

2. BACKGROUND

Flying Qualities and Maneuvering Requirements: The basic mission of the F-117 was as a "stealthy" night attack aircraft but the customer specified that it be designed to satisfy the requirements for a class IV Fighter/Attack aircraft of the then current Mil-F-8785B flying qualities specification. The low RCS of the aircraft would allow it to evade detection by Radar but it was possible that it might be spotted visually by a patrolling hostile aircraft. Since the F-117 was to be unarmed, its only defense would be to employ large amplitude "jinking" maneuvers to break contact and escape. As a result, it was required that the F-117 pilot be able to rapidly maneuver the aircraft to its controllability limits without concern of a possible departure. Minimum acceptable values of g-onset rates, roll acceleration, and time to bank were specified.

Aerodynamic Stability & Control Characteristics: The unaugmented F-117 airframe exhibits multi-axis instabilities in large portions of the permissible flight envelope. These instabilities are caused by the behavior of wing vortices that emanate from a junction at the side of the inlet and the leading edge of the wing. The resulting flow characteristics are typical of highly swept delta wings. For zero sideslip, the vortices effect each wing symmetrically but they lose coherence as the Angle-Of-Attack (AOA) increases resulting in pitch axis instability and a degradation in elevon effectiveness. In the presence of sideslip (Beta), increases in AOA cause asymmetric behavior of the vortices that result in lateral-directional instabilities. Initial wind tunnel tests of the proposed configuration had been conducted prior to the contract award and the remainder of the transonic and propulsion effects testing was completed after contract initiation. These test results defined the characteristics that formed the basis for the FCS design and are summarized below.

Pitch Axis Characteristics: Within the normal CG range, the pitch axis is stable at low AOA, then exhibits neutral stability and then becomes progressively more unstable as the AOA increases. Similar stability characteristics occur at negative AOA. Uncontrollable pitch up and pitch down occurs at critical positive and negative AOAs that vary with Mach number. The elevon effectiveness in pitch decreases as the AOA approaches the critical pitch up and pitch down boundaries and then disappears entirely at higher AOA. The pitch axis instability is exacerbated by a cross axis coupling from the directional axis where Beta causes a destabilizing pitching moment. Transonic effects begin to appear at about 0.7 Mach number and the pitch axis becomes increasingly stable in the low AOA range as the Mach number increases. With a forward CG location, the pitch axis static margin is about 13% stable for flight at 0.9 Mach number at the maximum permissible AOA. With an aft CG location, the static margin with the landing gear extended is about 8% unstable for flight at 0.25 Mach number at the maximum permissible AOA.

The effective engine thrust vectors pass below the aircraft CG causing a small nose up pitching moment when the engine thrust increases. Other than an addition of a small nose down pitching moment and a small aft CG shift, extension of the landing gear has no effect on the pitch axis. A moderate nose down pitching moment appears as the aircraft settles into the ground effect just prior to landing.

Lateral Axis Characteristics: The aircraft exhibits essentially neutral lateral stability (dihedral effect) at zero AOA for all Mach numbers within the permissible flight envelope. As the AOA increases from zero, the lateral stability increases linearly with AOA as is typical of highly swept configurations. The aircraft is laterally stable in this low AOA regime through out the range of Mach numbers but abruptly shifts to strong lateral instability at a critical AOA that decreases as the Mach number increases.

Directional Axis Characteristics: At low Mach numbers with the landing gear retracted and the weapons bay doors closed, the directional axis is slightly stable at low AOA but abruptly shifts to strong instability when the AOA increases to a critical value. The directional axis becomes neutrally stable in the low AOA range as the Mach number increases. The abrupt shift to strong directional instability occurs at the same AOA where the lateral axis abruptly shifts to strong instability but is always less than the AOA where uncontrollable pitch up appears. The directional stability is decreased at all AOA when the landing gear is extended because of the large single piece nose gear door located well forward of the CG. With the gear down, the aircraft exhibits essentially neutral directional stability. Similarly, opening the weapon bay doors makes the directional instability much worse because the door area is mostly forward of the CG. The directional control effectiveness of the canted fins is retained through out the range of AOA for all Mach numbers.

Propulsion System Effects on Directional Stability: The inlets and the high aspect ratio engine exhaust nozzles are sources of two directionally destabilizing effects. The first is typical of any aircraft with the inlets ahead of the CG where the mass flow captured by the inlets is directionally destabilizing. This effect becomes negligible at airspeeds above 200 knots. The second destabilizing effect is due to complex flow interactions near the exits of the exhaust nozzles and it increases with Mach number. The nozzles and the platypus nozzle extension were designed to minimize the infrared (IR) and radar signature of the nozzles in the aft sectors. The exit plane of each nozzle is skewed relative to the aircraft centerline and results in the effective engine thrust vector being toed in towards the middle of the aircraft. The thrust vector acts through the nozzle centroid and the toe-in angle is a function of nozzle pressure ratio (i.e., engine power setting), nozzle manufacturing tolerances, AOA, Mach Number, and, most importantly, Beta. The directionally destabilizing effect can be visualized as a differential change in the toe-in angles when Beta changes, i.e., the toe-in angle of the upwind thrust vector decreases but the toe-in angle of the downwind thrust vector remains unchanged. Above 0.6 Mach, the directional axis is unstable at all

AOA because the destabilizing effects of the nozzles are greater than the low aerodynamic directional stability of the airframe. The toe-in angles of the two thrust vectors vary significantly and are almost always asymmetric. This asymmetry results in a net side force even if the thrust of both engines is the same. These directionally destabilizing effects and thrust vector toe-in asymmetries were first discovered during the Have Blue flight tests because a lateral accelerometer signal was used as a surrogate for Beta and was fed back to the FCS to augment the directional stability. If the true Beta was zero, the unbalanced side forces due to the asymmetric thrust vectors resulted in a false indication of Beta through the sensed lateral acceleration. The Have Blue FCS attempted to reduce the false Beta and thereby created a real Beta as measured by the Beta vane on the flight test nose boom. Analysis of the Have Blue flight test results showed that the thrust vector toe-in angles could differ by as much as 8 degrees at some flight conditions. This problem was eliminated in the F-117 by using the direct measurement of Beta to augment the directional stability. One advantageous effect of the trust vector toe-in angle is a smaller yawing moment transient following an engine failure.

Control Surface Sizing: The aerodynamic data base from the wind tunnel testing was put into a full envelope 6 Degree-Of-Freedom (6 DOF) simulation to determine if the control surfaces that had been tested would be effective enough for both normal and emergency situations. The basic rigid model wind tunnel data for all of the control surfaces was modified with estimated flex to rigid ratios to account for aeroelastic effects. An initial version of control laws was combined with the 6 DOF simulation for off line analysis and for the piloted flight simulator. Since the pitch and directional axes were unstable over large parts of the intended operational flight envelope, the two elevons on each wing and the two fins were sized by the requirement to prevent pitch up and yaw divergence during "jinking" maneuvers. Specifically, the elevons were dedicated to control of the pitch and roll axes and had to be effective enough to prevent pitch up with aft CG locations during abrupt symmetric pull ups, rolling pull outs, and approach to landing in turbulence. Similarly, the two fins were dedicated to control of the yaw axis and were sized by the specified maneuvering requirements with weapon bay doors open, turn entries at low speeds, engine failures at lift off, all combined with side gusts. An additional sizing requirement was that the pilot must be able to terminate any maneuver and land the airplane after any single surface actuator failed provided that the failed control surface remained near the neutral position. If the failed surfaces could be maintained near neutral, the design goal was that the pilot would be able to make an emergency landing with one operating elevon on each wing and one operating fin. Off line simulation of these various combinations showed that the directional control power of the fins seemed to be adequate but that the elevons were not large enough since pitch up could occur following abrupt pull-ups at low speeds. This was confirmed by pilot evaluations in the flight simulator. As a result, the chord of the elevons was extended one foot in the stream wise direction and the angular travel increased to + - 45 degrees. Because of the directional axis instability with the weapon bay doors open and because of uncertainties about fin flexibility, it was decided to reserve at least 25% of the available fin deflection as a control power margin during all large amplitude maneuvers at airspeeds above 250 knots.

Dynamic Model Tests: The aerodynamic data base for the piloted simulator included all of the static wind tunnel data and estimates of the primary pitch, yaw, and roll rate damping coefficients **Cm**q, **Cn**r, and **Cn**p respectively. Early piloted simulator evaluations showed that the aircraft could depart from controlled flight if the critical AOA was exceeded and violent, uncontrollable motions would follow. The estimates for the damping coefficients were considered to be valid for flight in the normal AOA range but were questionable after departure. Depending on the values assumed for **Cm**q, **Cn**r, and **Cn**p in the high AOA ranges, the violence of the out of control motions could be reduced and, in some cases, the pilot could recover control.

A series of dynamically scaled free flight model tests was made to determine what the actual departure motions would be and to see if recovery might be possible. Four unpowered 5.2% dynamically scaled models were constructed and launched by a catapult at various initial conditions from the fifth floor level inside a Skunk Works hangar in Burbank, California. The use of outside test facilities or a spin tunnel was not allowed because of program security constraints. Despite the low Reynolds number in the model tests, the results were considered to be valid because the flow was separated during the majority of the departure motions. One of the four models was radio controlled to determine if recovery was possible through control surface manipulation and/or drag chute deployment. The effects of different CG locations could be tested in the other three models and their control surfaces could be fixed in various positions prior to launch. After a horizontal launch with an initial velocity of 90 feet/second, the models could travel 250 feet down range and fall through a vertical height of about 50 feet before being caught in nets suspended above the hangar floor. Five high-speed cameras mounted at various positions in the hangar filmed the flights. The cameras were triggered when the catapult launched the model and time synchronization marks were printed on the film. This allowed the model position and attitudes to be determined for analysis. The model tests showed that there were no spin modes and that pitch axis departures resulted in very large amplitude pitch oscillations. If the model was launched with an initial Beta, the destabilizing pitching moment due to Beta caused a pitch axis departure. Based on these model results, if the full scale aircraft departed at an airspeed of less than 200 knots, the aircraft would oscillate in pitch at high positive and negative pitch rates through ranges of +- 120 degrees AOA and the pilot would be subjected to approximately +- 6 g's. The radio controlled model results also showed that once the departure occurred, recovery could not be made with the control surfaces but that it might be possible to recover if the drag chute was deployed as the aircraft passed through the low AOA range during the large amplitude pitch oscillation. Analysis of the model motions captured by the high-speed cameras allowed refined estimates of the various damping coefficients to be made and programmed into the flight simulator.

3. SYSTEM COMPONENTS

FCS Mechanization: Once the basic aerodynamic and propulsion characteristics had been defined, pilot evaluations in the simulator showed that it was not possible to maintain control for more than a few seconds without a continuous high level of artificial stability augmentation supplied through the FCS. To reduce the cost and to minimize the F-117 development time, it was decided to modify the Fly-By-Wire (FBW) FCS originally developed for the F-16. This system had been used successfully in the prior Have Blue flight test program and a good working relationship had been established between the Skunk Works and Lear Astronics who manufactured many of the F-16 system components. A further advantage was that the Lear Astronics plant was located in Santa Monica, California only 25 miles from the Skunk Works in Burbank. This close proximity combined with a streamlined procurement system permitted the F-117 FCS to be developed in the shortest possible time, at low cost, while maintaining high security.

The F-16 was just being introduced into operational service when the F-117 design began in late 1978. The F-16 FBW FCS had been under development for several years and all qualification testing (vibration, shock, temperature cycling, electro-magnetic, etc.) had been completed. The suitability of the basic technology for rapid modification and development had been satisfactorily demonstrated in the F-16 and Have Blue. Some F-16 components could not be used at all (e.g., air data probes and side stick) and some could be used without modification (e.g., rate gyros and accelerometers) but most of the F-16 components had to be modified to meet the special F-117 requirements. The availability of the F-16 FBW technology reduced the development time for the F-117 by at least two years.

Flight Control Computer (FLCC): The initial production F-16 FLCC was an analog computer so it was possible to use the chassis and power supplies as designed. New control laws specific to the F-117 had to be developed and mechanized in the FLCC to replace those developed for the F-16. The new control laws as well as the necessary air data computations were mechanized within the quad redundant FLCC to insure no degradation in flying qualities after two like failures in the electronic portion of the FCS. The possibility of using a digital FLCC instead of the analog unit was investigated but was abandoned because the fastest flight worthy digital machine available in 1978 only had about 20-25% of the through put required to meet the needs of the F-117. (The first digital machines capable of meeting the F-117 requirements didn't become available until the early 1990's.) In addition, the turn around time for hardware modification and validation testing of the analog FLCC during laboratory and flight testing was much shorter than that required for similar software modification and validation testing necessary for a digital FLCC in the 1978-1985 time frame. This was an important consideration because of the desire to achieve IOC as soon as possible.

Control Surface Actuators: A modified F-16 actuator controls each of the four F-117 elevons and the two canted fins. The modifications included different orientations of the actuator rod ends, different strokes, and different electrical plug configurations. These changes "Murphy proofed" the components so that regular F-16 components could not be accidentally installed into an F-117 and vice versa. Each actuator has a tandem cylinder, is powered by two hydraulic systems, and is electrically controlled by the FLCC through triple redundant electro-hydraulic servo valves. Any single failure in any actuator has no effect on flying qualities but two like failures in the electronic interface or servo valves could cause the actuator to revert to a fail safe mode where the actuator is secured at a specified control surface deflection.

Pilot Controls: An early production version of the F-16 side stick was used in the Have Blue aircraft but was not considered to be satisfactory so a conventional center control stick was designed for the F-117. The stick assembly contains spring cartridges and dampers to provide artificial feel. For pitch axis control, the stick grip moves 2.24 inches forward and 4.25 inches aft with a force gradient of 7 pounds/inch. For roll control, the stick grip moves 2.6 inches right and left from neutral with a force gradient of 4.4 pounds/inch. There is a small detent that identifies the neutral position in both pitch and roll. The stick natural frequencies are 33.8 radians/second in pitch and 25.2 radians/second in roll. The dampers are filled with hydraulic fluid and provide the stick with a damping ratio of about 0.6 for both pitch and roll. The stick moves when the pilot applies force to the stick grip and quad redundant stick position sensor signals are routed to the FLCC where they are used as pitch and roll maneuver commands.

The directional axis controls consist of adjustable cockpit rudder pedals with a spring cartridge to provide artificial feel. A break out force of 20 pounds is required to start moving the pedals and a force gradient of 62.2 pounds/inch exists for the +- 1.125-inch pedal travel to the mechanical stops. The low inertia of the pedals results in a natural frequency of 158.5 radians/second. There is no purpose built damper on the pedals and the damping ratio of about 0.7 is entirely due to friction. Quad redundant pedal position sensor signals are routed to the FLCC to be used as directional maneuver commands.

Air Data Probes: The development of the air data probes was probably the single most difficult task in the entire F-117 program. During a five-year development span, more than forty major design variations were manufactured and tested before all requirements were successfully met. The FCS requires good measurements of AOA and Beta for stability augmentation and Mach number and dynamic pressure measurements are needed for scheduling purposes. Calibrated Air Speed (CAS) and barometric pressure altitude are required to allow the F-117 to safely participate in combined flight operations with tankers and other USAF aircraft. These measurements are continuously required in all weather conditions (icing, heavy rain) and any externally mounted air data probe must be able to survive an occasional lightning strike or an in-flight collision with a large bird. Many conventional air data probes had been previously developed which could meet these requirements but they could not be used on the F-117 because of RCS penalties.

The Have Blue technology demonstrator aircraft had successfully used a combination of three very small total pressure probes and a distributed set of static pressure ports mounted flush with the top and bottom surface of the nose to obtain the required air data. This system had minimal effect on the total measured RCS. Have Blue was never flown in clouds or in icing conditions because the air data system did not have anti-icing capability. An attempt was made to use a similar flush port system for the F-117 but it was found to be impractical due to excessive heating power requirements for operations in icing and heavy rain conditions. Tests in the icing tunnel showed that deicing the surface within a few inches of the static pressure ports with buried electric heaters was possible but any melted water would refreeze after running aft for a short distance. An ice dam would then form and cause significant static pressure measurement errors. The amount of electric power required to eliminate the possibility of the ice dams forming under all reasonably likely environmental conditions was greater than practical and this approach was abandoned in favor of low RCS air data probes which could be kept ice free with much less power. There was also a concern that heating large surface areas on the top and bottom of the nose would increase the forward aspect IR signature beyond desirable levels.

A minimum of four probes was required because the customer specified that the pneumatic system components dedicated to FCS functions had to sustain any two like failures without reducing the flying qualities below the Mil-F-8785B level three. Systematic surveys were made during wind tunnel tests to determine where four probes could be located to obtain the required measurements with acceptable accuracy. It was no surprise that the only practical location was at the nose. The testing at that location began with short probes which proved to be unsatisfactory. The probes were lengthened and re-tested until the test results showed that the desired measurement accuracy could be attained and that the flow interference effects near the nose could be adequately compensated over the AOA and Mach range. Six separate pressures measured at the probe tip pass through tubing approximately 120 to 130 inches in length back to the closest location in the nose where the pressure transducers could be mounted. Minimizing the tubing diameter inside the probe would minimize the external dimensions but would cause unacceptably large time lags in the AOA and Beta air data measurements. A simplified first order analysis showed that if the internal diameter of the straight tubes inside the probe was at least 0.050 inches for the first 66 inches, the lags would be acceptably small and it would be possible to attain the specified closed loop gain and phase margins. This 66-inch segment was followed by a section of airframe mounted tubing that varied from 54 to 64 inches in length with an internal diameter of 0.188 inches. The airframe tubes were terminated at the pressure transducers originally developed for the F-14.

Flight tests would show that larger tubes were required inside the probe for the first 66 inches. (See discussion in system testing.)

It was extremely difficult to design a faceted low RCS probe that could be internally heated without cracking the special low RCS materials on the probe exterior. The internal temperature had to be high enough to boil water since liquid water droplets could not be allowed to block the tubing during flight in heavy rain. The temperature gradient between the interior and exterior of the probe was high and differential expansions of the various materials caused cracking. This problem was not completely solved until the spring of 1983 and all of the initial flight testing was done with unheated probes.

Electrical and Hydraulic Power Systems: Each of the two engines provides power to an Airframe Mounted Auxiliary Drive (AMAD) with a mechanically driven power take off shaft. The primary generators and hydraulic pumps are mounted on and driven by the two AMADs. In the event that power can't be delivered to either AMAD, an Emergency Power Unit (EPU) can provide emergency electric and hydraulic power. The EPU can be driven by compressor bleed air from either main engine or from a conventional Auxiliary Power Unit (APU) which acts as a "third" engine if both of the propulsion engines fail.

The F-117 has two 3000 psi hydraulic systems each powered by two 45 GPM pumps. One pump in each hydraulic system is mounted on each AMAD so that both systems remain pressurized in the event of an engine failure. The system reservoirs were originally developed for the F-15 and incorporate fluid level sensing which is used to isolate leaks in hydraulic system subcircuits. The 10 GPM emergency pump powered by the EPU supplies power to only one of the two hydraulic systems.

Uninterrupted electric power is provided to the F-117 FCS from multiple power sources. The primary power sources are two 30/45 KVA generators. One generator is mounted on each AMAD. The 5 KVA emergency generator, the ship's primary battery, and dedicated FCS batteries provide emergency power. In the event of the loss of all other electric power sources, the dedicated FCS batteries alone can provide a minimum of ten minutes of operation starting at full charge. If all electrical power is lost and if it is assumed that at least one hydraulic system is pressurized, the elevon actuators go to a fail safe position of 2.5 degrees Trailing Edge Up (TEU) and the fin actuators go to 1.5 degrees TE inboard. These control surface positions bias the aircraft toward positive g giving the pilot the best opportunity to eject safely.

4. DESIGN CRITERIA

Design Criteria for PIO Prevention: The USAF Aeronautical Systems Division flying qualities engineers were aware that previous aircraft had met all of the requirements specified in the then current Mil-F-8785B but still had experienced unanticipated and potentially dangerous Pilot Induced Oscillations (PIO) during flight testing. It was clear that Mil-F-8785B was deficient in defining the design criteria required to eliminate PIO and a great deal of research and flight testing with variable stability aircraft had been focused on the determination of basic causes. Consequently, the use of design criteria that might result from the on going PIO research efforts was allowed and encouraged.

The design criteria which was of most use to the F-117 development came from the Landing Approach High Order System (LAHOS) study which was published in March 1978 (See reference 2). Detailed analysis of the LAHOS results indicated that excessive time delay between the pilot input and aircraft response could lead to a PIO. The evidence was based on evaluations made by two well-qualified test pilots during flight tests in the variable stability NT-33 aircraft operated by Calspan under contract from the USAF. Although the data base was small and there was no independent verification available at that time, it was decided to design the F-117 control laws to minimize the time delay between the pilot input and the first perceptible aircraft response. As a result, the equivalent time delay in the F-117 pitch and roll axes is between 0.07 and 0.11 seconds. The correctness of this approach was partially validated when the HQ specification was updated to the Mil-F-8785C version published in November 1980. This updated version included a new requirement that the total time delay between the pilot input and aircraft response should not be more than 0.10 seconds.

Steps were also taken to eliminate control surface actuator rate limiting as a possible trigger for PIO. Since the elevons were used to control both pitch and roll, the requirement for large amplitude "jinking" maneuvers resulted in large angular deflections and very high surface rate requirements. A combined pitch and roll command from the pilot would add together for the elevons on one side of the aircraft and subtract from each other on the other side. The required actuator rate limit was established based on study of previous PIO time histories and an evaluation of maneuvers in the flight simulator. The design criterion that was developed from this analysis is that the pilot must be able to move the stick sinusoidally in both pitch and roll at a frequency of one Hertz with an amplitude of at least + - 50% of full travel without exceeding the actuator rate limits.

To maintain control of the unstable aircraft during large amplitude maneuvering, both the elevon and fin actuators must have large hinge moment margins to prevent the actuators from stalling even momentarily during maneuvers. Flight simulation showed that if the actuators stalled while the aircraft was maneuvering in a part of the flight envelope where it was unstable, a departure could occur in less than one second. The database for the flight simulator was upgraded to include control surface hinge moment coefficients as a function of deflection angle, AOA, Beta, and Mach number. Since the CG of each control surface was always aft of the hinge line, the flight simulator included the hinge moments due to normal and lateral accelerations, pitch acceleration, etc. in the total hinge moment calculations. With this simulator capability and the rate limit design criterion, it was possible to define the following flight control actuator design requirements.

With both hydraulic systems operating and delivering 2,450-psi pressure differential across each of the tandem actuator pistons, the attainable no load control surface rates must be at least +- 135 degrees/sec for the inboard elevons, +- 180 degrees/sec for the outboard elevons, and +- 75 degrees/sec for the fins. These no load rates ensured that the rates under aerodynamic and inertial hinge moment loads would be high enough to prevent rate limiting during anticipated maneuvers. In addition, the actuators must have enough hinge moment margin to allow the pilot to recover the aircraft from any permissible maneuver following the loss of either one of the two hydraulic systems.

5. System Design

Control Law Design: Three basic sets of control laws were developed for the F-117 and mechanized in the analog FLCC. The first set is used for take off and landing and is activated when the landing gear handle is in the gear down position. The second set is the normal "up and away" control laws that are used when the landing gear handle is in the gear up position. The third set is used for air refueling and is activated when the pilot opens the air-refueling door. Faders are used to eliminate switching transients during control law transitions. The control laws were configured to provide optimum handling qualities for the airplane in each of the critical flight phases. Because of the extremely unusual appearance of the airplane, a design requirement was to make the aircraft behave as much like a conventional aircraft as possible. The service pilots that would be transitioning into the F-117 would have obtained all of their previous flying experience in conventional high performance aircraft and it was essential that the aircraft respond in a normal fashion so that the pilot would be comfortable using normal piloting techniques.

Classical design methods were employed for the control law development and the gain and phase margin requirements of the then current Mil-F-9490D specification were satisfied. A brief description of the final version of the control laws at the end of the flight test development phase and the design rationale of various features are given below.

Pitch Axis Control Law: The up and away pitch axis control law is a proportional plus integral g-command augmentation structure incorporating an independent high gain AOA limiter. A signal linearly proportional to fore and aft stick position is passed through a small dead band at the stick neutral position and is combined with a limited authority pitch trim signal to be used as a pitch axis command. With this implementation, deflection of the stick from neutral is treated as an aircraft maneuver command which is satisfied by a specified blend of washed out pitch rate and the normal acceleration measured at the pilot's station. Full aft stick always generates a command to attain a positive 7-g but the AOA limiter can over ride the pilot input if the maximum allowable AOA is attained first. Similarly, full forward stick position generates a command to attain a negative 2-g. Although there is an uncontrollable pitch down instability at large negative AOA, a negative AOA limiter was not implemented because it would come into effect only at airspeeds below 200 knots and no tactical benefit could be established for large negative AOA maneuvers in that regime. If it should become necessary to maneuver in this regime, the pilot is cautioned to monitor the AOA gauge and not permit the negative AOA to exceed the specified safe value. The proportional plus integral structure generates an elevon command to reduce the difference between the stick position command and the specified blend of pitch axis feed backs. Because of the inherent pitch axis instability, there is no unique steady state relationship between stick position and the elevon position so the pitch trim is mechanized as a series trim, i.e., the stick is always neutral when the pilot trims to a zero stick force condition. Apparent speed stability is provided at airspeeds below 200 knots by the addition of a nonlinear AOA feedback when the sensed AOA exceeds 7 degrees. A term proportional to the square of sensed roll rate is fed back when the roll rate exceeds a specified threshold. This feed back prevents inertia coupled pitch axis departures if high roll rate maneuvers are performed at high AOA. The forward loop gains are scheduled as a function of the inverse of measured dynamic pressure.

The dynamic model tests had shown that recovery would be very unlikely if the aircraft ever departed from controlled flight so major emphasis was placed on the design, implementation, and pilot evaluation of the AOA limiter. The very high gain proportional plus integral AOA limiter has more pitch command authority than the pilot and is activated when a non-linear combination of AOA and washed out pitch rate exceed a limiter threshold boundary. Following an abrupt full aft stick input, the washed out pitch rate input to the limiter threshold provides the anticipation required to begin reducing the nose up pitch rate before the AOA attains the maximum allowable value. The limiter threshold boundary is scheduled as a function of sensed Mach number and the pitch rate anticipation input is scheduled as a function of sensed dynamic pressure. This non-linear mechanization of the limiter prevented the pitch rate anticipation from making the aircraft too sluggish during the first part of abrupt pull-ups and allowed the specified maneuvering g-onset rates to be attained. The pitch trim authority is limited so that the pilot can always get to the AOA limit with full aft stick but cannot trim the aircraft at that AOA. The high AOA limiter gain and authority prevents the pilot from departing the aircraft by driving the pitch axis into resonance with fore and aft stick pumping at any attainable amplitude or frequency. It was recognized that a pilot could always depart the aircraft by allowing the airspeed to fall to zero in a steep climb so the pilot is cautioned to maintain enough airspeed so that the elevons can control AOA and prevent departure.

When the landing gear handle is placed in the gear down position, the g-command augmentation changes to proportional only but the AOA limiter retains the proportional plus integral structure. The AOA limiter is also biased upward so that the pilot can fly the aircraft to a higher AOA for lower approach and landing speeds when the gear is down. With the proportional only feed back structure, the aircraft has the same pitch axis characteristics as a classical aircraft exhibiting a phugoid mode and a well damped short period mode. The proportional only structure with the landing gear extended eliminates the need for any Weight-On-Gear mode switching at lift off or touch down.

When the Air Refueling door is opened, the g-command augmentation also reverts to the proportional only structure and the stick position input gain is decreased. In contrast to the gear down configuration, the AOA limiter onset boundary remains at the up and away values.

Roll Axis Control Law: The roll axis control law is a proportional only roll rate command augmentation structure with sensed roll rate as the only feedback. A signal linearly proportional to lateral stick position is passed through a non-linear parabolic shaping with a small dead band at stick neutral and is combined with a limited authority roll trim to command a roll rate. Like the pitch axis, the roll trim is mechanized in series. For take off and landing, the roll rate feedback gain is increased when the gear handle is placed in the gear down position to reduce roll response to turbulence. For air refueling, the lateral stick position input gain is decreased to maintain control harmony with the pitch axis.

Elevon Pitch and Roll Command Mixer: The elevon position commands from the pitch and roll control laws are prioritized, limited, and summed in the command mixer. Because of the pitch axis instability, the pitch axis command is always given priority and all four elevons operate between 25 degrees TEU and 37.5 degrees TED for pitch axis control. Each elevon is controlled by a dedicated actuator that can position the surface between 45 degrees TEU and 45 degrees TED. Any elevon travel beyond that required to satisfy the elevon pitch command is available for roll control up to the physical actuator stops. Because the elevons were sized to provide pitch axis control, they are much larger than they needed to be to satisfy the roll maneuvering requirements at the higher airspeeds. To avoid the structural weight penalty that would be incurred if all the available roll control power were to be used, the maximum allowable elevon roll command is progressively reduced as the sensed Mach Number increases. The mixer also ensures that the roll elevon command increments to each wing are equal in magnitude but opposite in sign to prevent the introduction of disturbances into the pitch axis.

Directional Axis Control Law: The up and away directional axis control law is a proportional only Beta-command augmentation structure with an automatic Beta trim. A signal linearly proportional to cockpit pedal position is passed through a small dead band at the neutral pedal position and is summed with a limited manual yaw trim signal to form a Beta command. The maximum allowable Beta command is scheduled as a function of sensed dynamic pressure to permit landing in specified crosswinds at low airspeeds but is progressively reduced to limit the maximum attainable Beta at higher airspeeds. At all airspeeds, the Beta command is satisfied by a specified non-linear combination of sensed Beta and washed out stability axis yaw rate. The Beta feed back gain is scheduled as a non-linear function of sensed AOA which increases the Beta feed back gain by a factor of two at the higher AOA. This is necessary to maintain tighter control of Beta at the higher AOA to minimize the destabilizing pitching moment generated by Beta. A dedicated actuator controls each fin and both actuators receive the same command. With this implementation, the fins are commanded to go to the position required to satisfy the command from the pilot. Because of the directional axis instability, there is no unique steady state relationship between the cockpit pedal position and the fin position so the yaw trim is mechanized as series trim. A product of pitch rate and roll rate is fed back to the yaw axis to prevent inertia coupled directional axis divergence during high roll rate maneuvers. The original directional control law design at the beginning of the flight tests included a so called "Aileron-Rudder-Interconnect" to help coordinate turn entries but it was removed after evaluation by both contractor and customer test pilots. Their collective opinion was that it did not provide much improvement in turn entries since the FCS was already responding to actual measurements of Beta and that it made precise pointing of the aircraft more difficult in wings level side slips.

The automatic Beta trim acts through a small dead band and is used to keep the steady state Beta within the dead band. The trim rate is proportional to the amount that Beta exceeds the dead band value. This mechanization accommodates any thrust vector toe-in asymmetries and prevents apparent roll mis-trims that would otherwise appear because of the rolling moment due to Beta (dihedral effect). The automatic trim is deactivated when the landing gear handle is placed in the gear down position or when the pilot applies more than 20 pounds of pedal force. If an engine fails at lift off, the pilot is instructed to retract the landing gear, keep his feet on the floor, concentrate on altitude and bank angle control, and allow the FCS to reduce Beta while the aircraft accelerates to the single engine climb speed.

There are no changes to the directional control law for air refueling.

Automatic Flight Control System (AFCS): At the beginning of the F-117 program, Tactical Air Command (TAC) pilots worked with the system designers to define the AFCS capabilities needed for the mission. The pilots wanted an AFCS that could be engaged at any time to provide pitch and roll attitude hold, heading hold, altitude hold, Mach hold, and automatic navigation to pre-programmed way points. The defined hold functions included stick steering which allows the pilot to change attitudes, heading, etc. to new reference values without first disengaging the automatic system. This AFCS was mechanized as a dual channel analog system in the Navigation Interface and Autopilot Computer (NIAC). The AFCS was designed to automatically disengage after detection of any first failure and illuminate a warning light for the pilot. The pilot then has the option of engaging either single channel to determine which channel had failed and then re-engage the functional channel and continue the mission. The AFCS has less authority than the pilot so that either the pilot or the AOA limiter can overpower any single channel hard over malfunction. This basic system was used in the F-117 fleet from IOC in 1983 until 1990 when a major system upgrade began to be installed as part of the Offensive Capability Improvement Program.

The upgraded system contained all of the functions of the original but added a Flight Management System (FMS) consisting of an auto throttle, completely automatic 4-D navigation (latitude, longitude, altitude, and time), a coupled approach mode, and a Pilot Activated Automatic Recovery System (PAARS). A complete description of the FMS and PAARS is contained in references 3 and 4. The new FMS was mechanized in the modified NIAC as a dual channel digital system that permits more sophisticated navigation modes and greatly improved built in test capability. The pilot prepares his mission plan prior to flight and loads it into a module which is inserted into a receptacle in the cockpit. After take off, the pilot engages the FMS and it flies the aircraft according to the pre-recorded mission plan. The pilot has the ability to modify the plan in flight if mission requirements change after take off. The system does not have the capability to do automatic air refueling or completely automatic landings.

The PAARS is an all attitude capable recovery system that can be used to return the aircraft to a wings level climb if the pilot should become disoriented during a night mission and experience vertigo. The pilot can activate PAARS at any initial attitude or airspeed within the permissible flight envelope by a single depression of a dedicated PAARS switch on the stick grip and the system will then automatically choose and execute the best unusual attitude recovery technique based on the existing flight conditions. The system continuously computes and displays an estimate of the minimum altitude that will be attained during the recovery so that the pilot can monitor the progress of the recovery.

6. System Testing

Hardware Development Tests: A complete Iron Bird was constructed to validate the FCS design by testing the actual system components. The Iron Bird included the tubing of both hydraulic systems laid out and supported as it is in the aircraft. The hydraulic pumps and electric power generators were driven by large variable speed electric motors to simulate the effects of varying engine power settings. Force loaders simulated the aerodynamic hinge moments acting on the control surfaces to verify the hinge moment design margins of the actuators after hydraulic system failures. The FCS components and all of the interconnecting wire harnesses were installed and routed as in the aircraft. The 6 DOF simulation could be coupled to the Iron Bird which included a cockpit and a visual system. This allowed a pilot to exercise the actual hardware under the anticipated loads within the design flight envelope. The Iron Bird was used to investigate system interactions, validate the failure modes and effects analyses, evaluate FCS failure detection and redundancy management, verify the adequacy of the emergency procedures, and perform some of the initial fatigue life tests.

FCS Flight Testing: The first flight of the F-117 prototype occurred on 18 June 1981 with Hal Farley at the controls. A comprehensive history of the F-117 flight test program can be found in reference 1 and more technical details can be found in reference 5. Although the FCS was evaluated to some extent on every flight, several of the more interesting aspects of the FCS development have not been previously reported and are covered below.

Testing for PIO Tendencies: Very early in the flight testing, concerted efforts were made to see if there was any tendency for PIO to occur. Prior to attempting air refueling, a fixed pipper reflector sight was installed on the glare shield of the first two test aircraft and some limited Handling Qualities During Tracking (HQDT) testing was performed using chase planes as targets. Experience with other aircraft had shown that HQDT testing is a good way to expose latent PIO tendencies. In addition, sinusoidal stick pumping at various amplitudes and frequencies was regularly performed during envelope expansion tests to detect the presence of any unwanted coupling tendencies. There were several instances of low damping due to faulty system components (For an example, see below), but no overt or latent PIO tendencies were detected during flight when the FCS was functioning as designed. These results were validated when it was found that the F-117 generally satisfied the PIO criteria published by John Gibson in April 1982 (reference 6) at typical flight conditions within the design envelope. Flight conditions at forward CG locations generally did not satisfy all parts of the criteria. To the Author's knowledge, no operational pilot flying an F-117 has ever experienced a PIO. This may indicate that Gibson's PIO criteria in reference 6 are somewhat conservative.

Effect of Air Data Probe Heating: As previously noted, the air data probes on the F-117 test aircraft were unheated for approximately the first year of flight-testing. The performance of the air data probe measurements used in the FCS was considered to be satisfactory after calibration with chase aircraft and a nose mounted flight test boom. When the first probes capable of being heated for anti-icing were installed on the test aircraft in October 1982, they developed external cracks almost immediately and the heaters were disabled until an improved probe design could be made available. After the cracking problems were solved and the "final production" version of the probes had been installed on the test aircraft in February 1983, a new problem appeared when the probe heaters were turned on for a flight in March 1983. As long as the aircraft flew at altitudes less than 10,000 feet, little or no difference was apparent in aircraft flying qualities with the old unheated probes and the new probes. However, as the test altitude increased, the pilot noticed a tendency for the nose to wander slowly from side to side, something that had never been seen with the old probes. This was not a PIO since the pilot had his feet on the floor and was not attempting to control heading. During a level acceleration to 0.9 Mach at an altitude of 30,000 feet, the nose wander became a continuous oscillation of plus and minus three degrees of Beta with a six-second period. The pilot aborted the flight and returned for an uneventful landing. On the next flight, the probe heaters were turned off, and there was no evidence of nose wander at the 30,000 foot altitude. It was found that the nose wander seen in flight could be duplicated in the flight simulator if the theoretical lag value used in the air data probe simulation was increased as the altitude increased. As noted earlier, the theoretical model of the probe lag used for sizing the probe tubing was a simple first order model that did not account for the effect of air viscosity. The increased lag with the probe heaters turned on was caused by the viscosity of the air in the probe tubing increasing as the three quarter power of the absolute air temperature. (See reference 7.) When this effect was included in a more detailed probe model, it was found that increasing the internal diameter of the probe tubing from 0.050 inches to 0.061 inches would eliminate the increased lag due to the probe heaters. This was verified by laboratory frequency response testing of a probe having the larger internal tubing with the heaters on and off. When new probes with the increased tube size were manufactured and flight tested, there was no indication of nose wander at any altitude with the probe heaters on or off. Because of the high concurrency of the program, approximately 25 probes had already been built for the production aircraft and it was necessary to disassemble and rebuild them with the new larger tubes. This was the last FCS related milestone that had to be satisfied before limited IOC was achieved in October 1983.

User Pilot Evaluations: TAC pilots were an integral part of the F-117 Joint Test Force and they regularly evaluated the entire weapons system during simulated operational flights. After the new heated probes had been certified, the TAC pilots flew the first ten-hour night evaluation missions in weather and could provide assessments of the flying qualities during night air refueling from the point of view of a fatigued operational pilot. Such evaluations were a valuable check on the suitability of the FCS before IOC.

Most of the TAC evaluation missions were generally routine but occasionally a flight would provide more excitement than anyone wanted. For example, one TAC pilot was the first to experience a lightning strike on an air data probe while flying in bad weather. The probe measurements were degraded but there was no transient because the redundancy management in the FLCC prevented propagation into the system. After the pilot recovered his composure, he was relieved to see that the only abnormality was a warning light indicating that the heater circuit on one probe had failed. After an uneventful landing, the pilot was surprised to see the extent of damage sustained by the probe that had been struck. This was a good confidence builder for TAC because it demonstrated the robustness of the FCS and that the handling qualities were not affected by a lightning strike on one of the most critical sensors.

FLCC Modifications: Flight-testing began with the -3 version of the FLCC and it was modified several times as results dictated. Experimental even numbered versions were flown with pilot selectable variations to test different design options in flight. The odd numbered versions were releases for production. Thus the -5 FLCC configuration was the standard for the IOC aircraft in October 1983 and included the modifications that had been flight-tested in the -4. The –7 was the final production version of the FLCC released after the high AOA flight-testing was completed in 1985. (See references 1 and 5 for details of the high AOA flight tests.) The -7 was then retrofitted into the fleet starting in 1986 and it is still the fleet standard in 2003.

In reference 5, Hal Farley summed up the F-117 FCS flight tests as follows: "The net result of all this Flight Control Development is an airplane with comparable pitch and roll response to that of a conventionally shaped contemporary fighter and attack aircraft within certain boundaries. Despite press coverage to the contrary, the aircraft is very maneuverable and fully aerobatic." Given the singular appearance and unusual aerodynamic characteristics of the F-117, this is amazing but true.

7. Conclusion

The development of the FCS for the F-117 was successful for five basic reasons. First, the Have Blue flight tests revealed the directional axis effects of the engine exhaust nozzles which were circumvented by the direct measurement of Beta in the F-117 FCS. Second, the LAHOS test results provided important control law design criteria. Third, the early assignment of TAC pilots to work with the design team on a daily basis kept the focus on the final user's requirements. Fourth, the close coordination between the USAF and the contractor team permitted the use of innovative and streamlined management procedures in all phases of the full scale development program. Finally, the ready availability of the F-16 FBW technology allowed the F-117 FCS to be developed in the shortest possible time at low cost.

8. Acknowledgments

The following individuals have been of great help to the author in recalling the events associated with the F-117 development. All are now retired but the positions they held at the beginning of the Have Blue and F-117 programs are indicated.

Alan Brown, Skunk Works Chief Engineer for Have Blue and the first Program Manager for the F-117, later head of Special Technologies.

C. R. (Dick) Cantrell, Skunk Works Chief of Aerodynamics, later head of all Flight Sciences.

H. J. (Skip) Hickey, USAF Aeronautical Systems Division flying qualities engineer, later USAF Chief Engineer for the entire F-117 Program.

This paper was reviewed by the F-117 Security office to insure that it contains no classified information.

References

[1] Aronstein, D. C. and Piccirillo, A. C., Have Blue and the F-117: Evolution of the "Stealth Fighter", AIAA case study, 1997.
[2] Smith, R. E., Effects of control system dynamics on fighter approach and landing longitudinal flying qualities, AFFDL-TR-78-122, March 1978.

[3] AIAA 92-1077, Flight Management System Integration on the F-117A, S. Combs, A. Sanchez-Chew, and G. Tauke, 1992 AIAA Design Conference, 03 February 1992.
[4] AIAA 92-1126, Pilot Activated Automatic Recovery System on the F-117A, S. Combs, K. Gousman, and G. Tauke, 1992 AIAA Design Conference, 03 February 1992.
[5] Farley, H. and Abrams, R., The F-117A Flight Test Program, Proceedings of the Thirty-Fourth Symposium of the Society of Experimental Test Pilots, September 1990.
[6] Gibson, John C., "Piloted Handling Qualities Design Criteria for High Order Flight Control Systems", AGARD CP-333, April 1982, pp. 4-1 to 4-15.
[7] Truxal, John G., Control Engineers Handbook, First Edition, section 16, page 16-7, published in 1958 by the McGraw-Hill Book Company, New York, NY.