

TERRAIN FOLLOWING TESTING FOR LOW ENERGY AIRCRAFT AND IMPROVED DATA ANALYSIS TECHNIQUES

Capt Russell E. Erb
Air Force Flight Test Center
412 Test Wing
Edwards AFB, California 93524-5000

Abstract

Changes in world climate and air defense systems have created the requirement for aircraft other than fighters and attack bombers to operate in the low-level environment. These aircraft, such as transports and helicopters, are characterized by lower energy operation when compared to fighter type aircraft. As such, some additional considerations are required when testing terrain following (TF) systems for these aircraft. A review is made of basic planning requirements and test execution methods for TF testing, with an emphasis on additional requirements for low energy aircraft. Data analysis is a major part of the TF test effort, since the quality of TF performance is not necessarily obvious from the cockpit. Classical methods of TF data analysis have relied heavily on statistical methods and isolated obstacle crossing performance. Additional data analysis techniques, relying on graphical methods and results over all types of terrain, are presented for better overall characterization of TF performance. These techniques are also useful for better communicating TF performance to managers and aircrews not intimately familiar with TF analysis techniques.

Introduction

With the current state of surveillance radar and air defense systems, virtually all nonstealthy aircraft have been driven to operating in the low-level environment. Fighter and attack bomber aircraft have operated in this arena for many years, but these aircraft have typically been characterized by high energy, with plenty of airspeed and power which could be converted to a zoom climb, if required, to avoid terrain in an emergency. With the increasing importance of low-intensity conflicts, many other aircraft, such as transports and helicopters, primarily in the Special Operations Forces, have been driven to operating in the low-level environment. While it is one thing to hang a TF radar on the nose of an existing aircraft, these aircraft typically have two characteristics that make them significantly different from early

fighter and attack aircraft TF systems. The first characteristic is these aircraft fly at a much lower energy state without as much excess power as most fighter aircraft. This greatly affects the planning required for setting up test runs. The second characteristic is the flight control systems typically lack the redundancy or other requirements necessary for installation of a full authority autopilot for automatic TF. As a result, the TF system is implemented through the flight director cue, requiring manual control inputs by the pilot to follow the cue. Integrating the pilot into the control loop adds a whole other dimension to developing the control loop.

This paper is written based on experience gained during the TF testing of the MC-130H Combat Talon II. The MC-130H is a modified C-130, fitted with a TF radar, integrated avionics, and a glass cockpit. It is used by Special Operations Forces for penetrating deep into enemy territory undetected at low level, day or night, in clear or adverse weather. The TF system is implemented as a manual TF system, with vertical movement of the flight director cue giving the pilot commands for TF flight. Lateral movement of the cue gives the pilot steering commands for navigation. The maximum sustained climb angle used in the MC-130H is 10 degrees, which is only possible at low weights and low altitudes on cold days. With a maximum normal speed for TF operations of 240 KCAS and a maximum usable climb angle of 10 degrees, the MC-130H is considered a low energy airplane.

The Air Force Flight Test Center published *A Handbook For The Conduct Of Terrain Following Flight Test* (TIH-88-003), which covers in good detail several aspects of TF testing. However, this handbook is based primarily on experience with TF testing of the F-111, F-16, F-15, and B-1 aircraft, all of which were fast, high energy aircraft. This handbook has a good description of the basic principles of TF flight, which will not be repeated here. The handbook emphasizes test execution, with less emphasis on technical and safety planning and data analysis methods. The data analysis methods that

are presented are heavily weighted toward use of classic TF parameters and statistical methods, which can be difficult to understand for someone without a significant knowledge of statistics. Additionally, the data analysis is heavily weighted to the use of simulations, which may not be available to the flight test team.

The objective of this paper is to go beyond this handbook and also to document lessons learned in the MC-130H TF program. Issues unique to TF testing of low energy and manual TF only aircraft will be discussed. Additional planning issues and test execution philosophy not covered in the handbook will be discussed. New methods of data analysis will be presented, with an emphasis on effective analysis when simulations are not available, additional TF parameters for more fully characterizing the TF performance, and returning the use of engineering judgment to TF data analysis.

Planning Issues

Performance Specifications

Performance specifications for a TF system are typically specified in terms of the classical TF parameters. In general, the aircraft should remain within some set deviation from the simulation results, such as +20/-15 percent of Set Clearance Plane (SCP). The hardness of the ride may be specified, such as no more than +0.7 g incremental or -0.5 g incremental.

The specifications most frequently concentrated on are for performance over isolated obstacles. Three parameters are typically defined. The aircraft should cross the peak within a given interval of the SCP. This is typically the same interval used for deviation from the simulation results. The aircraft should cross the peak within a given tolerance of level flight. Additionally, the aircraft should be pushing over (normal acceleration less than zero) as it crosses the peak.

While these specifications for TF flight are all characteristics of good TF performance, they do not fully define good TF performance with satisfactory mission utility. For instance, if the simulation does not give good TF performance, following that simulation closely will not give the aircraft good TF performance. Figure 1 shows it is possible to meet these three specifications and still have bad TF performance.

These specifications are tools that can be used to evaluate TF performance, but further analysis, coupled with sound engineering judgment and a mission utility

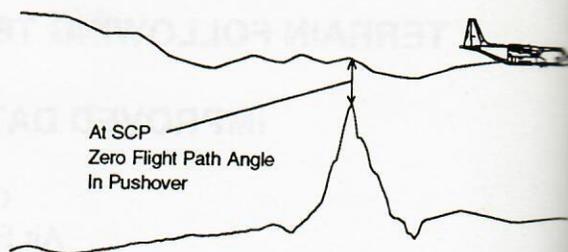


Figure 1. Poor TF Performance Satisfying Peak Crossing Criteria.

evaluation by the operational user, is required to determine if the aircraft has acceptable TF performance. The most important thing to do with specifications in the planning phase is to ensure that the appropriate data to answer the specifications can be collected during the flight test program. The specifications are of no use to the designers or the testers if they cannot be measured.

Instrumentation

TF flight is one area of flight test that is difficult to evaluate and virtually impossible to debug without some form of automatic data acquisition. Two major areas of instrumentation are required. The first area is instrumentation for measuring the performance of the TF system. These include:

- 1) Barometric altitude
- 2) Radar altitude
- 3) INS data (yaw, pitch, roll angles and rates)
- 4) Radar control point height and range
- 5) Cue commands
- 6) Flight control movements
- 7) Throttle movements
- 8) Airspeed/groundspeed

The barometric altitude and the radar altitude can be used to plot the profile of the aircraft flight over the terrain. The barometric altitude is plotted versus either time or distance down track to show the flight path of the aircraft. The terrain can be plotted by subtracting the radar altitude from the barometric altitude and plotting the result. Examples of these profiles can be seen in Figures 13 to 16. This profile will become very important in evaluating the overall quality of the TF flight.

INS data are needed to characterize the attitude and movements of the aircraft, and to evaluate the effectiveness of the antenna stabilization. Knowing the radar control point height and range will show what piece of terrain the radar is controlling on and when it stops controlling on that terrain. The ability of the pilot to follow the cue at a required level of performance can be determined by comparing the cue commands and the flight control movements.

The second area of data required is instrumentation for debugging purposes. No TF system has ever worked straight off the shelf, and some development effort will be required. Instrumentation will be required that records the data at various stages as it passes through the radar. The more in-depth the instrumentation, the easier it will be to detect the problems. However, the more in-depth the instrumentation, the larger the amount of data that will have to be recorded and stored. Perhaps the best solution is to have a system capable of recording just the radar data required for normal day-to-day analysis and debugging, with a capability to go more in-depth into the radar when the normal data are insufficient to determine the cause of the problem. The radar manufacturer will probably be the only one who knows enough about the radar to provide this kind of in-depth data acquisition.

An effort should be made to ensure that the instrumentation does not affect the performance or operation of the TF system. This may not be apparent until a production (i.e. noninstrumented) aircraft is flown.

Weather Testing

Testing TF systems in adverse weather presents two conflicting requirements. From a safety aspect, a TF system should not be flown on the cue into IMC until it has been thoroughly tested and evaluated in VMC. From a planning aspect, weather testing must be planned for early, since adverse weather is not a schedulable resource. Not only is its timing finicky, but also its location. The area around Edwards AFB sports a complete range of terrain types, which are good for VMC testing over different terrain types. However, this area is, not surprisingly, short on precipitation and adverse weather. Outside of Edwards AFB airspace, terrain following is only allowed on designated low level routes (IR routes), and these routes must be scheduled well in advance.

All of these factors work together to present two undesirable scenarios. In the first scenario, weather testing is put off until the end of the program, when the TF software is very mature. In this case, the program

may drag on for several years trying to match up planning with weather of opportunity. In between weather flights, very little would be accomplished since all other objectives were previously satisfied. This results in a very expensive program for very little data, and further delays the availability of the system to the user. Additionally, changes to the system to allow it to work in weather may affect the clear weather performance. This could result in re-flying the entire test program.

In the second scenario, weather testing is begun very early, in recognition of the fickle nature of weather. In this case, the risk involved in taking an unproven system into the weather is even greater, with a high probability that an unnoticed bug in the software or hardware could drive the aircraft into the ground. However, some weather testing can be accomplished with low risk by flying at Minimum Enroute Altitudes (MEAs) and not following the cue. The risk involved would be no different than flying IFR. This testing should highlight problems such as controlling on freezing layers, precipitation, or inability to discriminate terrain from weather. However, there is some risk of misleading conclusions due to flying high. For instance, the radar may be able to see the terrain better at lower SCPs when grazing angles are better and there is less weather to look through.

The best way to minimize the impact of weather testing is, like most other things, start planning for weather testing early. In this planning phase, it is important to note that there are typically two different types of adverse weather. It is very important not to confuse the two. The first type is fog and low clouds. This reduces visibility, but does not necessarily cause any false radar returns to confuse the system. The second type is precipitation. This may or may not reduce visibility, but will likely cause false radar returns.

Execute the weather testing plan as soon as the hardware/software configuration permit. This will give a better chance at finding weather while continuing development on the system. Flights at MEAs should be done early in the program to foresee future problems in weather. Deciding when to test in weather on the cue will involve getting enough time on a configuration to develop confidence that the system is working safely. The system may not be producing perfectly desirable TF performance yet, but it must be working safely and reliably. This decision must be an agreement between both the engineering analyst and the flight crews. Both groups must have veto power on this decision. When the decision to go is made, configuration control is absolutely critical to ensure no untested or incompatible hardware or software is introduced into the system, since the

number of available safety cross-checks will be greatly reduced. Be ready to change plans immediately and drop a planned test mission to execute a backup weather mission. Definite priorities regarding weather testing must be established, or weather testing will end up dragging on and on.

There are methods to increase the probability of finding weather when required. One option is to contact Air Combat Command Global Weather for climate predictions. They will develop long range forecasts to pinpoint areas of desired weather for long range planning purposes. Coordinate also with the National Weather Service to locate suitable test areas with good weather radar coverage. Ideally, the weather radar should have color coding or some other method to determine rainfall rates. Doppler radars are ideal for this determination. The Eglin AFB ranges have an ideal combination of weather radar and capability to measure actual rainfall rates on the ranges.

History has shown that all TF radars have multiple sensitivities or modes to deal with adverse weather operations differently than clear weather operations. A high sensitivity in the normal mode, while giving good terrain returns, will also see light precipitation as terrain. Some TF systems limit the top of the scan to the waterline of the aircraft in adverse weather. A lower sensitivity may be required for adverse weather operations to avoid problems associated with multi-path interference and reflections from the freezing level. A higher sensitivity may be required over low reflective terrain, such as sand dunes and snow. Expect to test and implement multiple sensitivities and modes. However, be aware of the hazards of multiple modes. For instance, a worst case scenario would be in a low sensitivity mode in heavy precipitation over low reflective terrain.

One method to help minimize the risk during weather testing is to use other sensors if available. One of the oldest and best is a well trained navigator with a ground map radar and a current chart of the area. Low level routes used for testing TF in the weather should be surveyed in VMC before starting weather testing and occasionally during the test program to ensure the charts are current, especially in the case of towers and other man-made obstacles.

It will probably takes years to test all of the necessary perturbations. Consider the number of different modes to be tested over different types of terrain in different weather conditions, along with the unpredictable nature of the weather. Prioritize testing on the most likely

weather and terrain combinations that the aircraft will see in its intended mission.

Operational Considerations

Certain operationally significant maneuvers may require capabilities not apparent on normal developmental or qualification terrain routes. For instance, imagine an aircraft flying at military crest along a ridgeline (Figure 2). The pilot decides he wants to fly closer to the ridgeline. In a TF system with turning capability, as soon as the pilot starts a shallow turn toward the ridgeline, never intending to fly over the ridgeline, the TF system assumes he will continue on around the turn. This causes the TF system to look into the ridgeline, and issue a climb command, or possibly an obstacle warning. If this type of maneuver is operationally significant, the TF system must be designed to allow it without causing false obstacle warnings. Communication with operational aircrew is very important to ensure that the system being tested will accomplish the intended mission. Failure to do this will result in spending lots of money to develop a system that may not be operationally suitable.

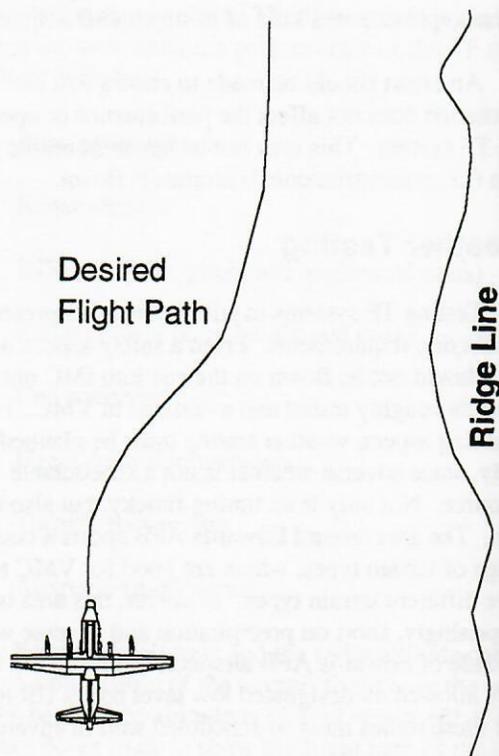


Figure 2. Example of an Operationally Significant Maneuver Not Normally Considered During DT&E.

Scheduling Considerations

Probably the most important realization to make up front in a new TF test program is that it will not be a short program. Do not be led astray by contractor's claims that everything is off-the-shelf and should work from the first day it is installed in the airplane. Every airplane is different, with different performance characteristics, flying qualities, and control systems. Installation of a new TF system on an airplane will therefore require some amount of integration and tweaking to make the TF system work effectively with the airplane. Therefore, some development time should be expected. The amount of development depends on the maturity of the system initially. The Combat Talon II TF system was adapted from a TF system developed for the HH-60 Nighthawk helicopter, and it took 5 years of flight test before the TF system was certified for use operationally. The initial testing of the Low Altitude Navigation and Targeting Infrared for Night (LANTIRN) TF capability was done on the F-16, and the manual TF testing of the F-16 LANTIRN lasted 2,000 flights. Integration of the LANTIRN on the F-15E benefited from the work on the F-16, taking only 20 flights to complete manual TF testing. The B-1 program, starting with a new TF system, took 55 flight hours of testing to get a usable system at 300 foot SCP. Only then was the TF system working well enough to start actually testing its performance.

Another reason that the program will not be short is the large number of different variables which must be tested over a given range. These include (but are not limited to) high to low SCP, level to rough terrain, straight flight, turning flight, altitude transitions, heavy to light weight, and high to low airspeed. Even with a sparse sampling of each of these variables, the test point matrix grows rapidly. For instance, consider the following test matrix: four SCPs, four terrain types, three weights, and three airspeeds. This yields a test matrix with 432 test points, after considering straight flight, turning flight, and altitude transitions. This does not account for multiple radar modes or sensitivities or adverse weather testing.

The third reason the program will be long is that TF testing is statistical by nature. This nature, along with the size of the test matrix, means that even with a lot of TF flying, there is a small chance that there is one or more areas in the envelope where the system could break down. The more flight hours amassed, the smaller this chance becomes, and the greater confidence the aircrews and engineers will have in the system. Additionally,

more flight hours give a better picture of the reliability of the system.

No effort is made in this paper to determine how long your program will be, other than to say:

1. The TF program will last longer than your contractor thinks.
2. The TF program will last longer than your program office thinks.
3. The TF program will last longer than your boss thinks.
4. The TF program will last longer than you think.

Test Execution

Prior To TF Testing

Prior to beginning TF, it is important to make sure all of the instrumentation, the antenna stabilization, and the Warnings, Cautions and Advisories (WCA) are working properly. Checking these items is not only important for ensuring the data are valid, but also for safety.

Other than the specific instrumentation added to the aircraft for TF testing, it is important to check the calibration on the radar altimeter. The radar altimeter will serve as the truth data for AGL altitude and crossing heights most of the time, since few TF routes have any other Time Space Positioning Information (TSPI) available, such as cinetheodolites. In addition, use of external TSPI can be very expensive. Additionally, the radar altimeter will serve as a "sanity check" for the aircrew for immediate feedback on TF performance. The radar altimeter can be calibrated by flying over a known course and tracking the aircraft with cinetheodolites. Of course, this method would be expensive. The radar altimeter can also be calibrated by any of the classical techniques used for calibrating barometric altimeters, as long as the elevation of the terrain is known. Probably the least expensive technique would be the tower flyby method. In this case, the altitude would need to be varied from one run to the next instead of the airspeed. The radar altimeter should be checked over a range from the lowest expected altitude below the lowest SCP to at least the highest expected altitude above the highest SCP. Another possible method would be to use a laser tracker for altitude information, similar to the tower flyby method.

Most current military aircraft use the Combined Altitude Radar Altimeter (CARA, AN/APN-232). There are several limitations inherent to the CARA that must be known while evaluating TF performance using CARA data. The first limitation is that the CARA has an inherent lag built into it. The indicated altitude is not the altitude measured by the most recent pulse, but a weighted average of the most recent altitude measured and previous altitude measurements. This averaging smooths the output of the CARA and normally has no impact operationally. Over smoothly varying terrain, the error resulting from this lag is very small. However, the lag can become significant over sharp peaks, where the radar altitude is decreasing and then increasing rapidly. The result is that the CARA will measure the peak at a lower elevation (higher radar altitude) than truth and at a time slightly after the actual peak crossing, as shown in Figure 3. This error will make the crossing heights appear higher than they are.

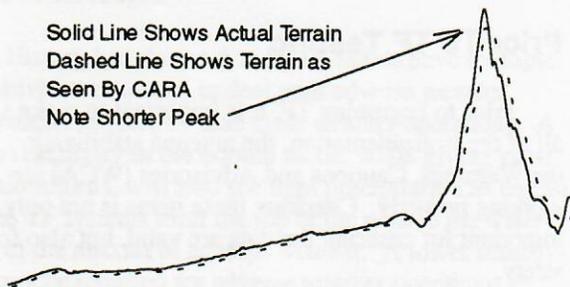


Figure 3. Exaggerated Effects of CARA Lag

A second limitation to be considered is the result of the beamwidth of the CARA. The CARA altitude reading is the slant range to the nearest return within a 120 degree cone (± 60 degrees from local vertical). While there is little error over flat terrain, in rugged terrain an error can be produced. If flying close to a canyon wall, the wall within the CARA beam may be closer than the terrain directly below the aircraft. Likewise, when flying toward rising or descending terrain, the shortest distance to the terrain will be at some other angle than vertical.

The third limitation of the CARA to be considered is a factor when flying over low reflective terrain, such as sand, water, or forests. The altitude reported could be the actual altitude above the nearest object, the altitude above the underlying hard terrain, or anywhere in between. Of course, there is no easy way to determine what the CARA is getting returns from.

Prior to TF flight, the antenna coverage and stabilization should be checked to ensure proper operation. Antenna coverage can be checked in flight test by

determining existence of accurate video displayed across the required area. Antenna stabilization can be verified by the accuracy of returns remaining constant as the aircraft changes pitch, roll, and drift angles throughout the course of each run. A Built-In-Test (BIT) function which operates continuously during the operation of the radar can be used to compare the actual antenna pointing angles with the commanded angles. Problems in satisfying pointing commands will show up as excessive pointing errors. Stabilization limits can be determined by loss of stable video on the displays where applicable and appearance of appropriate annunciators.

For safety reasons, it is important to check the WCAs before starting actual TF testing. The first efforts in developing the TF system should always be aimed at making the system safe, even to the detriment of its performance. After the system is safe, the system can be developed to give good performance.

Every TF system will have its own unique set of WCAs, but several are common to most systems. Common WCAs include an obstacle warning, a low altitude warning, stabilization limits exceed warning, and hardware fail warnings. Each of these warnings can be tested in a manner which will not endanger the aircraft or the crew.

The obstacle warning typically will indicate when the TF system declares that the aircraft can no longer clear an obstacle by a given amount using normal climb commands. To test this warning, fly the aircraft directly at an isolated peak at the same altitude as the peak or slightly higher. When the aircraft is close enough to the peak that the peak can no longer be cleared at the SCP using normal climb commands, an obstacle warning will be issued. Break off the run at a predetermined distance prior to the peak by climbing or turning away from the peak.

The low altitude warning can be tested by flying lower than the warning altitude over flat terrain in visual conditions. Stabilization limit warnings can be tested in conjunction with the antenna stabilization points by pitching or rolling the aircraft past the stabilization limits. Hardware fails can be simulated by pulling circuit breakers to remove power from the box which is simulated failed. However, this may not test all types of hardware fails.

After proper operation of all of the WCAs has been verified, the initial low-level flights should be made at a constant MSL altitude (without following the cue) over level or slightly rolling terrain. This altitude should be

high enough to clear the terrain by an acceptable margin, but not so high that the cue will command a constant full dive. The objective of these runs is to determine if the cue is reacting properly to the terrain as shown by the commands displayed to the pilot. This will probably be easier to determine during postflight data reduction than during the actual runs. Once proper operation of the cue has been verified, then the TF system should be ready to start actual TF testing.

Test Progression

In keeping with good flight test practice, the initial flights with the TF system should be accomplished in order from lowest risk to highest risk. However, since there are many dimensions to the test matrix, such as high to low SCP; level to rough terrain; straight flight, turning flight, altitude transitions; heavy to light weight; and high to low airspeed, the path from low risk to higher risk is not necessarily obvious. One possible method is suggested here, based on the MC-130H TF testing. For this testing, a multiple path strategy was developed, as shown in Figure 4. Based on the shape of this figure, this strategy was referred to as the "TF Waterfall."

On the left side of the waterfall chart are 70 combinations of SCP, maneuver, and terrain type. While these combinations are not all inclusive, they cover the primary areas that might be tested. All testing that is not on this chart, such as TF on the cue in weather, should be done at some time after the waterfall has been accomplished.

On the right side of the chart are 14 sets of boxes. Each of the boxes in a set correspond to the same maneuver and terrain, differing only by SCP. Within each box is a range of airspeeds from the minimum airspeed tested (such as 1.3 times the power off stall speed ($1.3 V_S$)) to the maximum airspeed tested and a range of weights from minimum to maximum gross weight.

The waterfall progression is started at the upper left corner at 1,000 foot SCP, straight flight over level terrain. After successfully accomplishing TF flight over a range of airspeeds and weights within a particular box, testing progresses diagonally down to the right to the next SCP in the same maneuver over the same terrain, i.e. "moving down the waterfall." The upper left box of each set can be reached after successfully accomplishing testing in any box directly above it.

Initial waterfall testing should be done at a medium weight. This will result in a medium range of climb

capability. This will probably be the lowest risk condition. A heavy weight will cause a low climb capability, resulting in climbs starting near the maximum range of the radar where radar estimates are the least accurate and climb angles are very small. A light weight will cause high climb angles, directing the aircraft closer to the obstacle before starting the climb. This can be a safety hazard if the system commands the climb too late.

Since there will inevitably be software and hardware changes required as development problems are discovered, guidance is required on how much of the testing needs to be reaccomplished. During the first march down the waterfall, a possible regression guidance states that a minimum of 25 percent of pertinent previously flown test points should be reflown. While the meaning of "pertinent" is open to engineering interpretation, normally this is taken as all previously flown test points. This is in recognition that any change in any piece of software in the control loop can have an adverse effect on the system, even in a supposedly unrelated area.

However, as a program progresses, repeating 25 percent of the previously flown test points can become an unbearable burden, especially when changes come out rapidly. It is possible that a change could be released before the 25 percent of previous points were completed on the current version of software. After the waterfall is initially completed, 25 percent of the previous points can be over 100 test points. As a result, another strategy is required as the TF system matures.

As the TF system matures, the test team's confidence in the system should increase. This is especially true if the contractor writing the software shows a history of successfully solving problems without creating new ones. Once the waterfall has been completed at least once, the ability to TF safely should have been demonstrated. Therefore, any new problems will arise from changes made to improve the system. For the initial flight of each new software patch, a standard regression profile should be developed. This profile should start with flight at the highest SCP over level terrain and progress in the same manner as defined by the waterfall. By using the same initial regression flight for each new software patch, a direct comparison can be made with results from previous software versions. Additionally, the engineers will develop a working knowledge of the expected performance and can quickly spot any adverse trends. Further testing on a new software patch will typically concentrate on areas that the patch was meant to change, but should also check across the spectrum of TF operations with sufficient magnitude to ensure the release has not adversely affected other areas of the system.

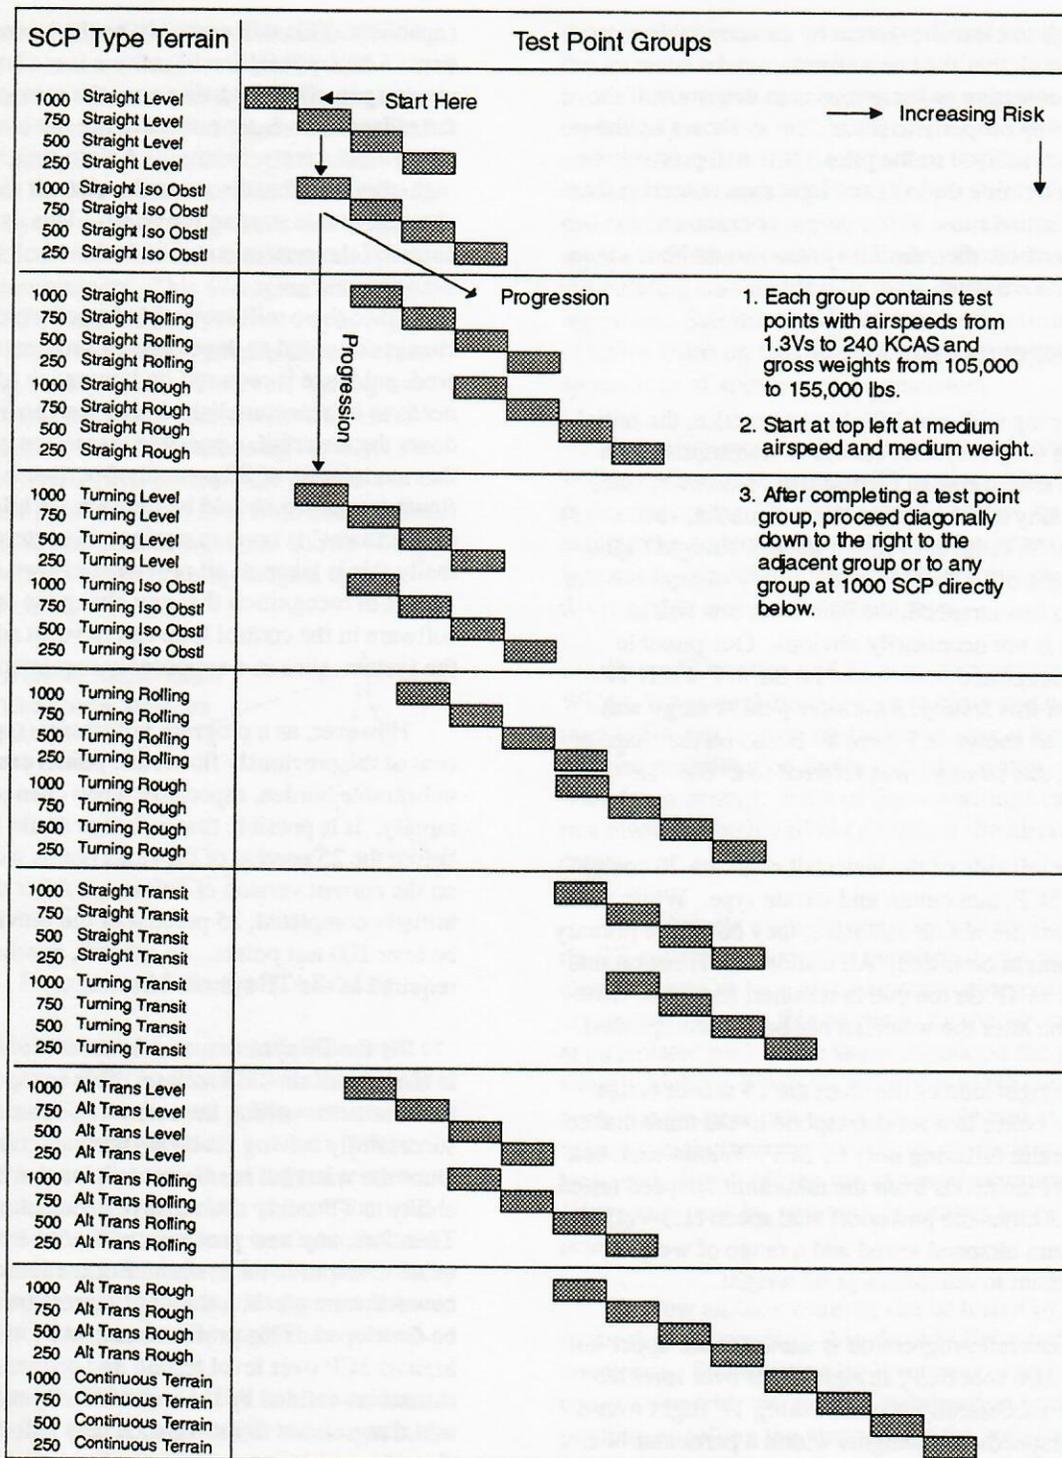


Figure 4. TF Waterfall Test Progression

This procedure will work during the development portion of the testing. However, for certification testing a test matrix should be developed to cover all aspects of the

TF spectrum. The entire range of each variable should be tested to minimize the possibility of not detecting a deficiency.

Configuration Control

Configuration control is the fancy name for keeping track of what software and hardware is loaded in the aircraft at any time. With good configuration control, the project engineer can determine what software and hardware was loaded on a given date at a given time at any point in the project's history. Since configuration control is basically a clerical job and rather unglamorous, it is easy to sluff off and ignore. Woe to the engineer who falls into this trap! Lack of configuration control can quickly become the biggest nightmare of the program. At the least, the engineers will not be able to track down problems because they will not know what software and hardware was loaded. At the worst, an incompatible combination of software or hardware will be loaded and result in the loss of the aircraft and crew.

With that said, how does an engineer maintain good configuration control? With a good string of documentation, of course. However, documentation alone will not work. A system which allows documentation error detection must be set up. Careless documentation is no better than no documentation, since inconsistencies will appear. Long after everyone involved with the project has forgotten what happened, an engineer will be reviewing the configuration control documents to research a problem and may find an inconsistency. At this point, the engineer will be in a dilemma. Which answer should he believe? Is either answer correct? This is why a system of error detection must be set up, since errors will inevitably be made.

When considering configuration control, it is important to realize that with integrated avionics, the entire set of software and hardware must be considered as a whole. Since all of these items work together, any change to any piece of software or hardware will affect the entire system, not only as intended but possibly causing new unintended problems. As such, the entire configuration of hardware and software should be tracked as a whole. Every time a piece of software or hardware is changed, the configuration number is incremented, even if returning to a previous configuration. This last point helps with the error detection. If an engineer is looking at data from two flights, and the configuration number on the flight with the later date is less than the configuration number on the earlier flight, the question is raised "Was the configuration number on the later flight recorded in error (a typo) or was the aircraft returned to a previous configuration?". If a policy of always incrementing the configuration number is strictly enforced, it would be

obvious that the configuration number was recorded in error.

Cue Dynamics

In a system only capable of manual TF, the flight director cue is the link between the radar derived flight commands and the pilot who actuates the control surfaces. For good flying qualities, the dynamics of the cue (how fast it moves, how much it moves) must be matched with the pilot's reaction time and aircraft dynamics. In general, this can be tested by defining a cue tracking performance criteria, flying TF over various types of terrain, and assigning Handling Qualities Ratings (HQRs) to how well the pilot can follow the cue. This testing would be done with different cue gains and filters if required. The cue gain would affect how much the cue moved. Filters may be required to reduce jittery motion in the cue, but must be tuned to ensure that no unacceptable lag is introduced. After making any changes to the cue dynamics, regression testing must be done to determine any effects on overall TF performance.

Some TF systems are designed with inputs to the flight director control loop other than the radar. For instance, flight control movement in pitch may be used to anticipate changes in vertical acceleration. This implementation assumes a change in vertical acceleration will follow a movement of the pitch controller. However, this is not always true. Pitch trim changes from throttle movement or flap extension require pitch controller movements which will not change the vertical acceleration. This effect is most pronounced at low speeds with large throttle movements and large angle-of-attack changes. In this type of system, testing should investigate not only how well the pitch controller movement anticipates changes in vertical acceleration, but if side effects, such as cue movement in reaction to throttle changes, are small enough to be easily compensated.

It is also possible that the cue gain will need to be scheduled with airspeed to maintain TF performance objectives. A higher cue gain at lower airspeeds can compensate for larger pitch controller movements required for greater changes in angle of attack.

Testing Over Multiple Terrain Types

TF testing should be done over all types of terrain, from flat to mountainous, since any type of terrain could possibly be encountered during a mission. One set of criteria for defining the roughness of terrain is:

Level - less than 30 feet rise/nm

Rolling - 30 to 600 feet rise/nm

Rough - greater than 600 feet rise/nm

These criteria are by no means a standard; each project may define its own criteria. These criteria work well for slow-moving (under 250 knots) transport type aircraft. TF routes in the Edwards AFB area which match these criteria are shown in Table 1. These TF routes are shown in Figures 5 and 6.

Straight Routes

In order to improve the quality of the data analysis, each TF run should be run over exactly the same route. Any deviation to one side will change the terrain profile and possibly the character of the entire run. The following methods are suggested as ways to maximize data repeatability.

Prior to the start of a route, the aircraft should be lined up with the route at the proper SCP and airspeed. Flaps, landing gear, and other parts of the aircraft should be configured as required. At the start Navigation Reference Point (NRP), the Flight Test Engineer starts the data recorders. The pilot follows the flight director cue to the defined performance criteria. Airspeed should be controlled using power to within a set tolerance (such as ± 5 knots for low speed aircraft, possibly greater for high speed aircraft), performance permitting. Power should be limited to maximum continuous power. If a power limit is reached, then the airspeed is allowed to bleed off. If the airspeed drops below a preset minimum speed, the run should be aborted.

The pilot should follow the lateral steering cue, if available, to keep the aircraft on the centerline of the run. Prior to each run, a flyover update at approximately 500 feet AGL over a specified point will help to keep the INSs accurate so that the guidance will cover the same terrain. This procedure will minimize the effects of INS drift and maximize the repeatability of the terrain data, thus improving the data analysis. Upon reaching the end NRP, the FTE turns off the data recorders.

The Harpers route is used for level terrain, and includes a section approximately 10 nm long of very flat terrain over Harper Dry Lake. This route can be flown in either direction.

The Red Mountain route is used for isolated obstacle testing. A climb through approximately 2,000 feet of

altitude is required, which is usually sufficient for low energy aircraft to reach a steady-state climb before cresting the peak. Haystack Butte can also be used for isolated obstacle testing. This route has the advantage of cinetheodolite coverage. However, with only a 250 foot change in altitude, the maneuver over the peak is very quick and dynamic, with the pushover beginning almost as soon as the pilot starts to react to the pullup command. As such, it was more of a perturbation in the flight path without any steady-state conditions. The effects of any lag in the system are amplified by the quick nature of the crest. Additionally, it is more difficult to analyze the radar data since Haystack Butte consists of two closely grouped knolls. It is possible the control point for the radar can move to the second peak before the aircraft crests the first peak, thereby further corrupting the data. While not ideal for isolated obstacle testing, Haystack Butte can be useful for testing short range TF performance. The nature of the obstacle will cause the pullup to start close to the obstacle.

The Modified Black route starts with flat terrain, running up to essentially an isolated peak, and then into rolling terrain. This run can be used for multiple purposes. The initial flat portion is excellent for gathering data over flat terrain, up to the point where the radar starts controlling on the first peak. Therefore, the length of the run usable for the level portion varies directly with the current climb capability of the aircraft. The back side of the initial peak is the steepest terrain of any of the runs, and is useful for checking a commanded dive angle limit. The remainder of the run is used to evaluate TF performance over rolling terrain. This run is only used in the west to east direction.

The Saltdale route covers rolling terrain, with two sections of terrain that vary at less than a 10 degree slope, which should be within the usual climb limits of most aircraft. These sections can be used to evaluate the ability of the aircraft to follow the terrain without having to account for steady-state climb sections. Although this route could be flown in either direction, it is most efficient to fly from east to west in an alternating sequence with the Modified Black route in a figure-8 pattern.

The Modified Rough route covers rolling terrain, transitioning into rough, mountainous terrain. This route is only flown from south to north.

The Rough 1 route covers mountainous terrain. Near the north end of the route is a large valley approximately 1,500 feet deep and 4 nm wide. This area is useful for investigating performance of the TF system in

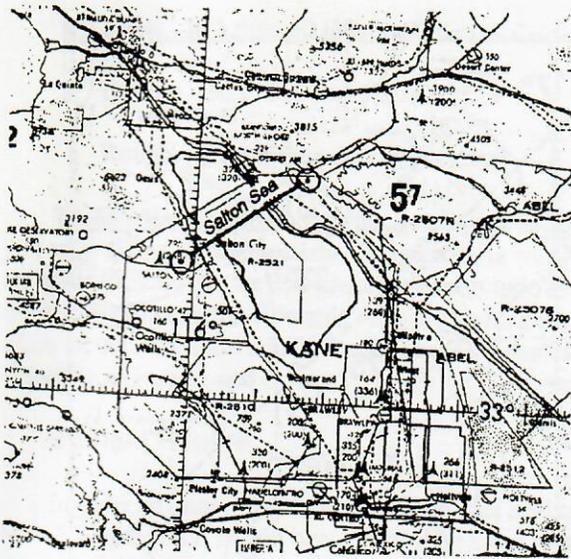


Figure 6. Salton Sea TF route.

transitioning from full dives to maximum climbs. This route can be flown in either direction. Flying from north to south presents the radar with terrain that is generally going downhill. Flying from south to north presents the radar with terrain that is generally increasing in elevation. This can have an effect on the TF performance at low climb angle limits, since at the north end of the run several peaks in excess of 9,000 feet elevation are within 20 nm.

Turning Routes.

The Harpers Turning route is defined by a point on the Harpers straight TF route, and is used for turning TF performance over level terrain. All turns are made to the south of the Harpers straight TF route. The Saltdale Turning route is defined by a point on the Saltdale straight TF route, and is used for turning TF performance over rolling terrain. All turns are made to the north of the Saltdale straight TF route. The test procedures for the Harpers Turning and the Saltdale Turning routes are similar. Prior to the start of the route, the aircraft is lined up on the corresponding straight TF route at the proper SCP and airspeed heading toward the start turn point. About 20 seconds prior to the turn point, the Flight Test Engineer starts the data recorders. Upon reaching the turn point, the pilot starts the turn. If the lateral steering of the cue does not command the turn, the lateral steering should be deselected. Otherwise, the cue may end up pegged to the side of the display and be difficult to follow. The direction of turn (left or right) is determined by the direction of entry (east to west or west to east) into the turn point. The pilot banks the aircraft to the turning

limits of the radar for the current conditions. For a turning TF system, the current turning limits should be displayed to the pilot at all times. Vertical guidance is by the flight director cue, with the same technique as in straight flight. As the turning corridor scan data are built up, the bank angle limit increases in the direction of the turn. As the bank angle limit increases, the pilot increases the aircraft bank angle, and thus the turn rate, to make the tightest possible turn. The turn is continued through 360 degrees, rolling out wings level parallel to the entry track. After the corridor returns to a straight flight corridor, the Flight Test Engineer turns off the data recorders.

This technique maximizes the data repeatability, but even so there are inherent problems. The radius of the turn increases as the airspeed increases, and the ground track is altered depending on the winds. As a result, each turning run has a different ground track. There is not the level of repeatability from run to run that is possible over the straight TF routes. Additionally, the analyst is required to have a working knowledge of the terrain surrounding the turning area. Unlike in straight flight, many times during the turn the radar can be controlling on a point outside or inside the turn that the aircraft will never fly over. Since the control point will not show up on the trace of terrain that the aircraft did fly over, it is more difficult to explain thoroughly cases where the aircraft would seemingly not follow the terrain shown on the data plots. The end result of these factors is that the turning TF data inherently has more scatter than the straight TF data, which had to be considered for proper interpretation. The amount of scatter will increase the more rugged the terrain is, but with a good knowledge of the terrain in the area and experience, a trained analyst will be able to determine good TF performance. This training will be greatly increased if it is possible for the data analyst to fly on the test aircraft during TF runs enough to become familiar with the terrain.

The Ridge Line Turning route is used for testing turning TF performance against an isolated obstacle. This route is defined by an initial point and a turn point. The line between these two points is parallel to the ridge line. The pilot flies the aircraft in TF from the initial point to the turn point at the specified test conditions. Upon reaching the turn point, the pilot deselects any lateral steering and makes a right turn into the ridge line. The pilot should roll into the turn using the same technique described for turns over level and rolling terrain. The aircraft is flown up and over the ridge line, crossing the ridge line on a heading roughly perpendicular to the ridge line. The turn is continued around 180 degrees of

turn, rolling out wings level parallel to the entry course. After the corridor returns to a straight flight corridor, the Flight Test Engineer turns off the data recorders. If this ridge line were perfectly straight on top, it would make an ideal isolated obstacle for turns. However, like most of nature, it is not straight. Data analysis must consider that the aircraft can be controlling on terrain that it will never cross, as previously discussed.

The Rough Turning route is set up differently from the other turning routes. This route consists of essentially an equilateral triangle with 8 nm sides (Figure 5). This route design is required, based on the available terrain, to allow testing turning TF in rough terrain without constant obstacle warnings from high peaks outside of the turn. The route is flown in a clockwise direction, starting at the southeast turn point (turn point 1). The flight director cue is used for vertical guidance throughout the run. The data recorders are turned on as the pilot flies the aircraft from turn point 1 to the west turn point (turn point 2). As the aircraft flies over turn point 2, the pilot starts a maximum rate right turn, staying within the bank angle limits. This turn is continued through approximately 160 degrees, then reversed to a maximum rate left turn, then rolled out on the next leg of the triangle. This procedure is repeated at the northeast turn point (turn point 3) and turn point 1. The maneuver is complete after 3 sets of turns are accomplished, and data recorders are turned off.

Altitude Transitions.

During operational flight, it may be necessary to transition from one SCP to another. Increasing the SCP is easy with a simple pullup away from the terrain. However, the TF system should allow for changing to a lower SCP with a smooth transition. This is necessary for operations such as letting down over unfamiliar territory in adverse weather. Altitude transitions can be done on any of the straight TF routes. Initial setup is identical to other straight runs. The run is started above the highest available SCP with the highest available SCP selected. The pilot follows the flight director cue down until the aircraft is stabilized on the SCP. The pilot selects the next lowest SCP and follows the flight director cue down until the aircraft is stabilized on the SCP. This procedure is repeated until reaching the lowest SCP. With four SCPs, typically three sets of altitude transitions can be done in the length of one TF leg. All altitude transitions within one set are done at a given airspeed. The airspeed is changed for each set of transitions on a particular leg.

Continuous Terrain

After completion of the segmented straight and turning routes, the TF system should be tested over a longer route encompassing as many different terrain types as possible. The purpose of the continuous terrain route is as a final check of the system over all types of terrain and transitions between each. This route should be set up similar to an operational mission intended for the aircraft. For Developmental Test and Evaluation (DT&E), this route should be flown on centerline as done on the segmented routes for data repeatability. For Operational Test & Evaluation (OT&E), it may be desirable to deviate off of centerline for terrain masking and to evaluate the capability of the system to maneuver around the centerline.

Land-Water Transitions.

The primary purpose of land-water transitions is to verify that the TF control loop can transition smoothly between TF based on radar altimeter data and TF based on forward looking radar data. This testing can be accomplished on the Salton Sea route (Figure 6). Straight runs start over land, fly toward the shore, continue across the water to the other shore, and up over land. This tests the transition in both directions. All other procedures are identical to those for straight TF runs. If the water is not smooth, it is possible that the forward looking radar will detect returns from the water, and the system will not transition to radar altimeter TF.

Turning transitions can also be tested. These runs would start over land flying toward the shore. Over the shoreline, the pilot starts the turn, following the same procedure as on other turning TF runs. The run is continued through about 1 1/4 turns to test the transition in both directions while in a full turning scan.

Low Reflective Terrain.

Sand dunes and snow covered terrain have a lower radar reflectivity than normal terrain. The performance of the TF system should be tested over both types of terrain, if possible. The Dumont Dunes route (Figure 5) is accomplished in the same manner as isolated obstacle runs. These sand dunes are approximately 400 feet high. TF over snow covered terrain is normally not possible in the Edwards AFB local area, and thus would have to be accomplished elsewhere.

Testing over low reflective terrain can be hazardous, especially if the radar does not detect the terrain. This may require pilot intervention to prevent ground clobber.

TF in Weather.

Routes intended for testing TF in adverse weather should be set up in areas with a high probability of frequent adverse weather, at least during a few months of the year. Two possible areas are the Pacific Northwest and Northwest Florida. The Pacific Northwest is good for low rainfall rates with low visibility. Northwest Florida is good for higher rainfall rates. TF will normally be restricted to existing IR routes, complying with any published restrictions. These routes should be surveyed by flying the route first in VMC, noting any towers or other obstructions not shown on the charts.

Towers.

No towers suitable for good TF testing exist within the Edwards airspace. The towers that do exist are either near populated areas, where TF flight would be prohibited, or are near rising terrain which the TF system might control on. The ideal tower would be a tower on level terrain, isolated from population centers, and in an area where TF flight is allowed. Level terrain is desired so that the radar control point is not ambiguous between the tower and the terrain.

Testing against towers follows the same basic procedures as for isolated obstacles. However, since the tower may not be picked up by the radar and controlled on properly, it is critically important to have a predetermined point prior to reaching the tower where the run will be aborted if no indications of controlling on the tower are received. Abort procedures can be as simple as turning away from the tower. Be sure you know if the tower has any guy wires and where they are, lest you strike them.

The tower chosen should be closely inspected for its radar reflection characteristics. If an antenna dish is mounted on the tower, the reflections from this dish may be the only part of the tower that the radar sees. From a safety standpoint, this is not a major problem if the dish is at the top of the tower, but it can be if the dish is only part way up the tower. Also, a tower with a dish near the top may lead to the false conclusion of acceptable TF performance over towers. The radar may interpret the height of the dish as the height of the tower. Beware! Not all towers have antenna dishes. A tower without antenna dishes may not be seen by the radar at all. Additionally, the reflectivity of the tower will change based on the run-in heading. The radar may also see the tower better at low SCPs than at high SCPs.

In the absence of an isolated tower, large power line towers may serve to give an indication of how well the radar can detect towers. Power line towers approximately 150 feet high run along Highway 395, and cross the Harpers, Modified Black, and Saltdale routes.

Certification Testing

When a decision is made to test a software configuration for certification to the user, it is likely that the latest configuration has not been tested against the entire spectrum of TF operations yet. During development, testing is typically aimed at detecting and solving problems.

For certification testing, the objective is to determine to the greatest extent economically feasible that no unknown software bugs or other problems exist, and to fully characterize any known problems. To do this, testing should be accomplished over all types of terrain, at all SCPs (not just the lowest SCP), and in all modes of the TF system. The software and hardware configuration must be frozen prior to starting the certification effort. Any change to the software or hardware after starting a certification effort will start a totally new certification effort, starting from the very beginning. Due to the nature of software changes, it is risky to assume that a new version of software will act just like the previous one.

When setting up the certification test point matrix, use parameter minimization to your advantage. Techniques for this are described under data analysis techniques. However, do not misuse parameter minimization. For example, assume that previous testing has shown that SCP error is independent of SCP selected. This does not mean that you can then test only at one SCP and extrapolate the results to other SCPs, since something may have changed since previous versions. What this does mean is that data can be taken at all SCPs. After reverifying the independence of SCP error and SCP, all of these data could be considered together to increase the number of data points, and thus increase the statistical significance. This way, for instance, a given airspeed would not have to be repeated at every SCP, thus reducing the number of test points.

Safety

Since TF operations are inherently more risky than up-and-away flying, several general risk minimizing procedures should be followed. The following is a list of procedures that should apply to most aircraft. Certain

aircraft may need additional risk minimizing procedures to these listed.

Prior to starting TF testing, ensure the accuracy of the radar data, the proper functioning of the antenna stabilization, and the proper functioning of all warnings, cautions, and advisories. Procedures for these tests were previously discussed.

Set a minimum airspeed. If the aircraft slows below this minimum airspeed, then the test point must be aborted, and the pilot will dive or turn as required to regain airspeed. A possible minimum airspeed to choose would be 1.2 V_{stall} . Another possibility is a minimum airspeed required to have sufficient energy to maneuver away from the terrain. A current value for this airspeed should be posted for immediate reference by the pilot at all times.

Know the patterns of bird migration for your test area. Chances of a bird strike are much higher in low level flight.

At the beginning of each test sortie, an operational check of the TF system should be accomplished. This is assuming any and all built-in-tests (BIT) have been passed successfully. This test should be tailored to the system under test. A suggested procedure would be to fly the aircraft visually (not following the cue) over level terrain to a radar altitude equal to the lowest SCP intended to be used. Climb to 150 percent of the SCP and verify that the system commands a fly down. Slowly descend below the SCP and verify a fly-up command. Continue to descend and verify that the low altitude warning and obstacle warning are functioning. Climb back to SCP and verify that the warnings reset.

One of the easiest and most important risk minimizing procedures is to have any crew member who is not actively flying the aircraft and has access to a window looking outside for other aircraft and ensuring proper terrain clearance. For a single-seat test aircraft, this could be done with a two-seat chase aircraft. The pilot of the chase aircraft will be watching the test aircraft, and the other chase crew member should be watching for other aircraft and terrain clearance.

Use all available sensors in the cockpit. The pilot flying should cross-check his altitude with the radar altimeter. If available, a crew member should monitor the ground map radar. Use any other sensors to your full advantage. This will become exceptionally important when testing in IMC, since the out-the-window view will no longer be available.

Familiarize yourself with the terrain ahead of time, if possible. This will help explain any anomalies that come up, such as climbs over level terrain in turning flight from controlling on distant high terrain.

Always leave yourself an out (escape maneuver). In the case of low energy aircraft, a zoom climb up and away from the obstacle is usually not an option. This is especially true if the system is designed to fly an aggressive profile, using the maximum climb capability available. Another suitable escape maneuver is to turn away from the obstacle toward lower ground. This should be considered even in the planning process, before ever getting into the airplane. Testing TF in a narrow box canyon would be stupid, unless there was a verified operational requirement to do so. Even then, additional procedures should be taken to minimize the risk. These procedures also apply to helicopters, since it is possible to fly the helicopter to a point in the envelope where forward flight is possible, but hovering is not.

For multiengine aircraft, the consequences of an actual engine failure must be considered and prepared for. For instance, over rough, high terrain, the aircraft may not have sufficient performance to climb to safety after an engine failure. If so, some action will be required. This could be a way to reduce weight quickly, moving the test site to another location, or waiting until another day when better aircraft performance is available.

For TF testing in IMC, additional procedures must be used. All flyups must be considered valid, regardless of the crew's first impression. While false flyups can be ignored in VMC where the pilot has a clear view of the terrain around him, this habit must not be allowed to cross over into IMC testing. Upon receiving a full flyup, climb immediately to the MEA until the flyup clears.

Data Analysis

Simulation

Simulations can be very useful for TF data analysis, if one is available. However, simulations have some inherent drawbacks which must be considered during the data analysis.

For manual TF systems, the simulation will only be as good as the pilot model. If the pilot model does not accurately match the way that pilots fly the aircraft, the simulation results will not match the aircraft data.

Most TF simulations are run using actual terrain height data as recorded by the aircraft. This may give acceptable results when compared to the same flight. However, if the simulation results are compared to another flight over the same route, the aircraft flight path may still be "correct" but not match the simulation results. The representation of the terrain in the simulation is two dimensional, but the aircraft flies in a three-dimensional world. Slight displacements to one side of the original ground track can change the terrain profile significantly. Additionally, the simulation may not account for controlling on terrain which is never overflown, such as in turns.

It would seem reasonable to compare the aircraft flight performance with the simulation by saying if the altitude of the aircraft was within a tolerance of the simulation value, such as +20/-15 percent SCP, then the performance would be acceptable. However, this is a trap, since it applies a one-dimensional test to a two-dimensional problem. For instance, if a pilot was a little late on the pushover over an isolated obstacle, he would overshoot the tolerances about the simulation. This is possible even if the flight path paralleled the simulation (Figure 7). In this case, what would otherwise be considered acceptable TF performance is thrown out as a failure due to bad analysis methods.

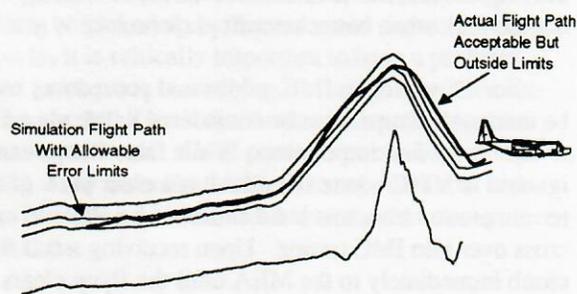


Figure 7. Acceptable TF Performance Fails Simulation Test

If a simulation is not available, the test team can "build their own" by maintaining data from past runs and overlaying the flight paths on a single plot, as described later in "Overlay Plots." After determining which runs show "good" TF performance by other means, these runs can then be used as a baseline for quick checks of subsequent runs. This method also will account for three-dimensional effects, as not all runs will be perfectly over the same centerline.

Minimizing Parameters

One of the biggest problems with keeping a TF program to a manageable size is the number of possible factors that can affect TF. Each additional factor increases the size of the test matrix geometrically, multiplying by the number of values that factor can take on.

But even as the test matrix grows geometrically with each parameter, it can also shrink geometrically by eliminating parameters. Parameter minimization works by determining by analysis of the system what parameters should not affect performance, assuming they do not affect performance, and testing the validity of the assumption. If the assumption is true, then no more testing is required. If the assumption is false, more testing is required and the test matrix is adjusted accordingly.

For example, peak clearance error may be a function of the airspeed (or groundspeed) that a leg is flown at. However, based on analysis of the system design, the engineer building the test matrix decides that the peak clearance error is probably not a function of airspeed. He assumes it is not, and builds a matrix flying peak crossings at four different airspeeds. When the data are plotted, the engineer looks for any trends with airspeed. If points of similar airspeed are grouped together and away from points at other airspeeds, the assumption is probably false, and additional testing is needed. However, if the points are randomly scattered together for all airspeeds, the assumption is probably true, and the amount of test points has been cut by a factor of four.

The key point is that data points were collected over different airspeeds, but can be considered as one group. This way, at any point in the testing the independence of peak clearance error from airspeed can be reverified. This is different from making the assumption and flying all points at a single airspeed. If all points were flown at a single airspeed, any changes in the software which introduce a dependence on airspeed will go unnoticed, and thus untested.

If a climb capability value is used in the TF calculations, the climb capability will likely normalize all data with respect to the factors on which climb capability is calculated. For instance, if the inputs to the climb capability equation are outside air temperature (OAT), altitude, and gross weight, then the aircraft will probably have the same TF performance for any combination of these factors which yields the same climb capability. Of course, this assumption must be verified by flying at dif-

ferent combinations of these factors. The only factor which is easily controlled by the test team is the aircraft gross weight. Testing should be conducted at various gross weights, to ensure there are no other effects of gross weight, such as changing the mass moments of inertia with different loadings.

Another possible parameter to eliminate is the SCP. If the calculations do not change based on SCP (other than the clearance altitude), it is possible that SCP error over level terrain and isolated peaks will be independent of the SCP. Make sure the SCP error is measured in distance, not percent of SCP, to apply this test.

Classic Data Analysis Techniques

The data analysis techniques in this section are called "classic" because they have been used on previous TF programs. The new data analysis techniques discussed in a later section were developed by the author.

G-Excursions

Specifications for a TF system will usually contain limits on the magnitude of the vertical acceleration for pullups and pushovers. While it is possible to look at the incremental acceleration at every data point recorded, it is usually more practical to extract the maximum and minimum incremental acceleration for each TF run. If these values are within the specification, then all other values for that run are within the specification. These maximum and minimum incremental g values can be collected for all TF runs over all types of terrain, and plotted in a histogram, as shown in Figure 8. The data points are placed in bins of a given width, such as 0.1 g, and the total number of points in each bin is plotted. The data will probably form two bell curve shapes, and the mean of the maximum and minimum values can be approximated from the histogram. Note that it is possible to have a few points outside of the limits and still have an acceptable system. These points could be caused by gusts or turbulence which momentarily increase the load factor on the aircraft with no command from the TF system.

SCP Error Over Level Terrain

Straight flight over level terrain is the most basic mode of terrain following. This condition creates the smallest demands on the system. All that is required of the system is to maintain a set aircraft altitude above the terrain. This mode permits determining the very fundamental functional relationships between the input vari-

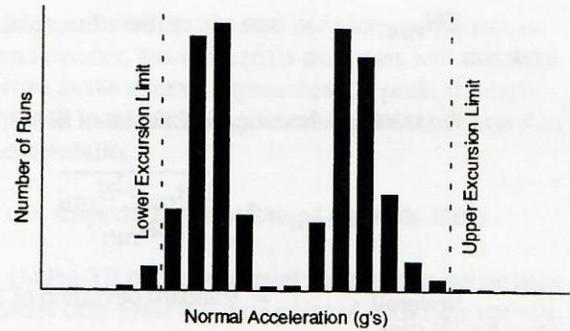


Figure 8. Vertical Acceleration Excursion Histogram

ables and the TF performance. This is the data where most of the parameter minimization can be accomplished. This mode is also the best for determining how closely the aircraft can follow the optimum flight path and how stable the performance is about the optimum flight path. This is because this is the only terrain where the optimum flight path is known confidently without a simulation model. The optimum flight path is simply the SCP above the terrain, with no sustained climb.

To determine how close the aircraft is to the optimum flight path, calculate the arithmetic mean of the SCP error for that run. To determine how much the aircraft oscillates about the optimum flight path, calculate the standard deviation of the SCP error for that run. These means and standard deviations can be plotted versus climb capability and airspeed to determine if any functional relationships exist. Effects of SCP can also be determined from these plots.

To calculate an overall mean and standard deviation for the system, do not fall into the trap of averaging the means for each run and finding their standard deviations. Since the runs will normally be different lengths, it is meaningless to add up the means, divide by the number of runs, and call it the overall mean. Long runs will skew the mean toward themselves and away from short runs. Likewise with the standard deviation. Therefore, the mean SCP error for all runs is calculated as

$$\mu_{\text{overall}} = \frac{\sum \mu_{\text{run}} N_{\text{run}}}{\sum N_{\text{run}}}$$

where

$$\mu_{\text{overall}} = \text{mean SCP error for all runs}$$

$$\sum \mu_{\text{run}} N_{\text{run}} = \text{sum of products of individual run mean SCP errors and number of samples in each run}$$

ΣN_{run} = sum of number of samples in each run

and the standard deviation is calculated as

$$\sigma_{overall} = \sqrt{\frac{\Sigma(\sigma_{run})^2 N_{run}}{\Sigma N_{run}}}$$

$\sigma_{overall}$ = standard deviation of SCP error for all runs

$\Sigma(\sigma_{run})^2 N_{run}$ = sum of products of individual run SCP error standard deviation squared (variance) and number of samples in each run

ΣN_{run} = sum of number of samples in each run

If the points representing the mean SCP error are plotted on a chart showing the overall mean and a confidence interval (such as $\pm 2\sigma$), it is important to remember that the overall mean and confidence interval apply to all of the data collected, not to the mean and deviation of the means plotted. Be careful with this type of plot, as it may be misleading.

Isolated Obstacles

Peak Clearance Error

Peak Clearance Error is simply the difference between the altitude at which the aircraft passed over the peak and the SCP. Note that while SCP error over level terrain may be determined to be independent of SCP, it is still possible that peak clearance error may be dependent on SCP. In general, the peak clearance error in feet will be less at lower SCPs. Peak clearance error is strongly tied to the accuracy of radar estimates of the peak height, since the radar will command the cue based on how tall it thinks the peak was, not how tall the peak really is. Since the estimate of the peak height is based on an angular measurement, the accuracy of the estimate varied with the aircraft position relative to the peak. The two important factors are the distance to the peak and the lookup/lookdown angle to the peak. The closer the aircraft is to the peak, the smaller the beamwidth is in feet, and thus the accuracy was better. If the lookdown angle is zero, i.e. the aircraft is co-altitude with the peak, the estimate will be the most accurate, with the accuracy decreasing as the lookdown angle increases. Figure 9 shows the relationship of the aircraft to the peak at two points based on SCP. The higher the SCP, the farther away the aircraft is from the peak when it is co-altitude

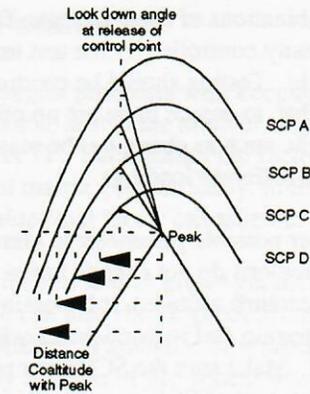


Figure 9. SCP Effects on Peak Height Estimates

with the peak. Thus, when the lookdown angle is the best, the accuracy is degraded by range at high SCPs. Regardless of SCP, the radar will stop controlling on the peak at some point prior to overflying the peak. At this point, the lookdown angle is greater at higher SCPs, thus degrading the accuracy of the peak estimate. In general terms, at lower SCPs, the aircraft is closer to the peak, allowing the radar to estimate the height of the peak better, resulting in better peak crossings.

Since the peak clearance error performance varies with SCP, a possible way to normalize the data is to look at the error as a percentage of SCP. If this normalization is perfect, the peak clearance error expressed in a percentage of SCP will yield the same value regardless of the SCP at which the data were taken. How well this technique works must be checked for each TF system to which it is applied. While using percentage of SCP may not be a perfect way to normalize the data, it may provide a better picture of the overall peak clearance performance. If the radar has very good, consistent peak height estimates, then measuring the error in feet may give a better picture of the overall peak clearance performance.

Flight Path Angle At the Peak

Another purpose of isolated obstacle testing is to determine if the TF system minimizes the exposure time of the aircraft above the terrain. This is evaluated by measuring the flight path angle of the aircraft as it crosses the peak. If the aircraft crosses the peak at zero flight path angle (level flight) at a constant pushover acceleration, then ballooning should be minimized. If the flight path angle was positive (climbing), then the aircraft will balloon on the back side of the peak and take longer to get down behind the peak. If the flight path angle is negative (diving), then the aircraft climbed too

much in front of the peak. This analysis is based on the assumption that the aircraft follows a nice parabolic pushover over the peak.

Results from flight testing may show that as the climb capability increases, the flight path angle at the peak increases. This can be explained by examining the entire run. As the climb capability goes up, the aircraft flies closer to the obstacle before pulling up. It then climbs at a higher angle. When the pushover begins, the aircraft has farther to pitch down to get to zero flight path angle. If the pitchover starts at about the same point each time, the flight path angle at the peak is tied to how high the climb angle is entering the pushover.

A primary reason that the flight path angle at the peak will tend to be positive is that the radar cannot see the terrain in the shadow of the peak that it is controlling on. Since it cannot see the terrain, this gap is filled with data based on certain assumptions. The problem that the radar has to solve when running in toward a peak is determining if it is truly a mountain peak, the front edge of a plateau, or a dam with a still lake behind it. The latter is the worst case, since when the aircraft rises above the dam, the still lake might not return any radar energy, which can cause the radar to interpret the dam as a mountain peak. Since it takes some time for the radar to scan the terrain after rising above the peak, the pushover will not be as aggressive as it would be if the terrain profile was known a priori.

Acceleration at the Peak

Measuring the aircraft vertical acceleration over the peak will give an indication of the trend of the flight path angle. If the acceleration is negative, then the aircraft is in a pushover. However, the noise level of the vertical accelerometer must be known. If the absolute magnitude of the noise is greater than the magnitude of the pushover acceleration, it is possible to read a small positive acceleration at the peak, even though the aircraft was actually in a pushover.

Staircasing

Staircasing refers to a characteristic of alternatingly climbing too early, then leveling off. A plot of the approach to the peak suggests a staircase rising to the peak. Staircasing is primarily caused by poor radar estimates of the peak. High estimates cause an early climb, which levels out as more accurate estimates are received. This condition is aggravated in low energy aircraft at low climb capabilities, because the TF system acts on radar data which are taken at a very long range, and therefore

less accurate. While this does not increase the risk of ground clobber, since the radar estimates will normally improve as the aircraft approaches the peak, the early climb will increase the exposure of the aircraft and thus its vulnerability.

Improving Statistical Sample Size

Unlike TF flight over level terrain, runs on isolated obstacles only yield one data point of each type per run. As such, it takes a long time and a large expense to collect a meaningful sample size of data. However, this sample size can possibly be improved by using data from previous software versions. This will be particularly true when the changes made between software versions do not address peak crossings. Data from each previous version, starting with the most recent, are compared with data from the current version. If the data prove statistically to be of the same distribution, then data from that previous software version may be used. This procedure is continued until a significant change is found in the data. At that point, only data from software versions after the change may be used.

New Data Analysis Techniques

General Curve of TF Performance

Figure 10 shows a notional curve which describes SCP error over all types of terrain for all types of aircraft. The basic shape of this curve will remain, although its specific shape and location will change based on the terrain in question and the capabilities of the aircraft. Understanding this curve will help the analyst know how to plan and how to analyze TF performance over different terrain routes.

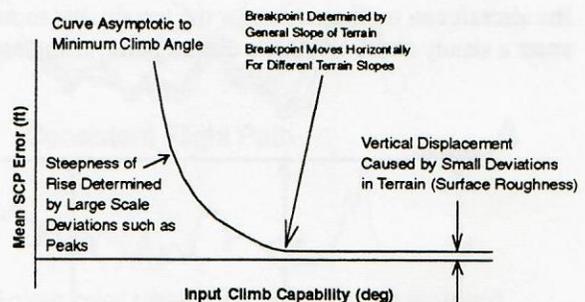


Figure 10. General Curve of TF Performance

The curve is characterized by two sections, above and below a breakpoint. For climb capabilities above the breakpoint, the SCP error is essentially constant. For climb capabilities below the breakpoint, the SCP error

rises sharply, going to infinity at the climb capability where the TF system stops working.

Consider the curve in terms of its different sections. The breakpoint between the sections will occur at a climb capability equal to the general slope of the terrain. For aircraft climb capabilities above the breakpoint (the upper section), the aircraft will have sufficient performance to follow the terrain contours. Below the breakpoint (the lower section), the terrain rises faster than the aircraft can climb. The aircraft starts a sustained climb earlier, and thus rises above the SCP, increasing the mean SCP error.

Above the breakpoint, the SCP error will be relatively constant. For perfectly smooth terrain and a properly functioning TF system, the SCP error should be approximately zero. This is the basis of comparing the aircraft altitude with the SCP over flat terrain. Over perfectly level terrain, any aircraft should be able to fly at the SCP regardless of climb capability. The curve describing performance over this terrain would consist completely of the upper segment, with the breakpoint at zero. Operations in this section lend themselves to easy engineering analysis by plotting mean SCP errors and standard deviations.

The mean SCP error will be positive for the upper section, and will be more positive the rougher the terrain is. This error is caused by dips in the terrain (depressions, crevices) too small for the aircraft to follow, but large enough for the radar altimeter to detect (Figure 11). These dips will increase the mean radar altitude, even though the aircraft is following the optimum flight path.

The SCP error of the lower section rises rapidly as the aircraft can no longer follow the terrain, but must enter a steady climb to clear a distant point, rising above

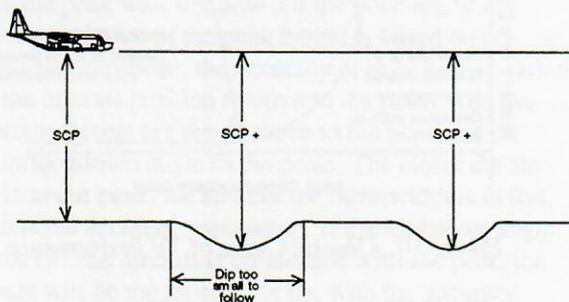


Figure 11. Apparent Increase in Mean SCP Error Due To Surface Roughness

the local SCP. The SCP error approaches infinity when the climb capability is so low that the aircraft no longer descends into a valley, but flies over it at a constant altitude. The steeper or more rugged the terrain is, the faster this section will rise. Operations in this area cannot be easily analyzed by plotting parameters on charts, and must use more engineering judgment to determine "good" TF performance.

Figure 12 shows how the general curve changes based on the terrain type. Over level terrain, the SCP error is very small over most or all of the range of possible climb capabilities. The breakpoint is close to or less than the minimum climb angle possible. As the ruggedness of the terrain increases, the breakpoint moves to a higher climb capability, and the SCP error at a given climb capability increases. The SCP error above the breakpoint will probably increase as the small variations in the terrain increase in number and size.

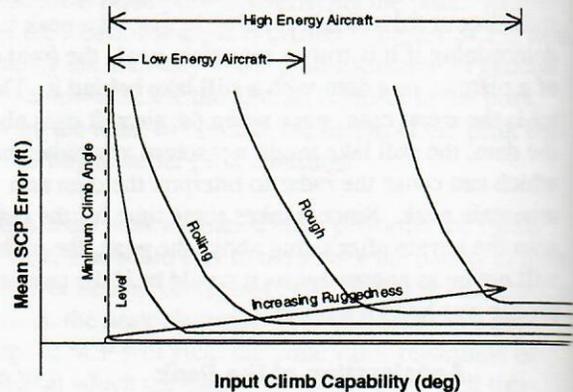


Figure 12. Changes in the Curve of General TF Performance for Different Terrain Types

Figure 12 also shows the areas of climb capability where low and high energy aircraft operate. One definition of an aircraft excess energy state is its ability to climb, which does not necessarily mean a faster aircraft. Therefore a high energy aircraft can operate at a higher climb capability and achieve a lower SCP error.

Figure 12 does not account for effects of speed or hardness of ride (maximum allowable vertical acceleration, both positive and negative). To follow the same flight path at a higher speed, a harder ride (higher vertical accelerations) will be required. For low energy aircraft with a small range of operational airspeeds (such as 100 knots from fastest to slowest), the effects of airspeed (groundspeed) will be small, and possibly will be hidden in the data scatter. The effects of airspeed (groundspeed)

will be further reduced if the TF system is designed to normally bleed airspeed while in a climb.

Overlay Plots

A good TF system should have the same performance over a given piece of terrain for the same climb capability, regardless of temperature and weight. The effects of temperature and weight should be accounted for in the climb capability calculation. Additionally, the performance should degrade gracefully as the climb capability decreases. Based on these premises, a wealth of information can be obtained by plotting on a single chart the flight paths of multiple runs at the same SCP over the same route. In the absence of a simulation, this is the best way to get a handle on the question "Is the TF system performing correctly?" Changes to software which make changes to the TF performance, especially undesirable changes, will show up quickly in this method as flight paths inconsistent with previous flight paths.

As mentioned above, for low energy aircraft, the effects of airspeed (groundspeed) will be small. For high energy aircraft operating over a large speed range, airspeed (groundspeed) will change the flight profile if the hardness of the ride does not increase with airspeed. In this case, the TF performance should degrade gracefully

as the airspeed increases.

To create these plots, the flight path must be expressed as an altitude for a distance down range. If distance is not available in the data, it can be determined with the time and the groundspeed. The distance should be incrementally calculated between each data point, since the ground speed will vary with time. Plot the first flight path on the chart with its corresponding terrain profile. This will be the baseline terrain profile. Plot the next flight path with its corresponding terrain profile. Shift the new terrain profile as required to get a best match with the baseline terrain. Shift the flight path the same amount as its terrain profile. After getting the best match possible between the two terrain profiles, remove the new terrain profile, leaving the flight paths and the baseline terrain profile. Continue this procedure for each run. Examples for the Modified Black, Red Mountain, Saltdale, and Rough 1 Routes are shown in Figures 13 through 16.

Examining the resulting plots should reveal certain characteristics. The flight paths over level and rolling terrain should be tightly clustered within a small altitude band and consistent with each other. Climbs toward peaks should start at increasingly earlier points as climb capability decreases. Regardless of the conditions

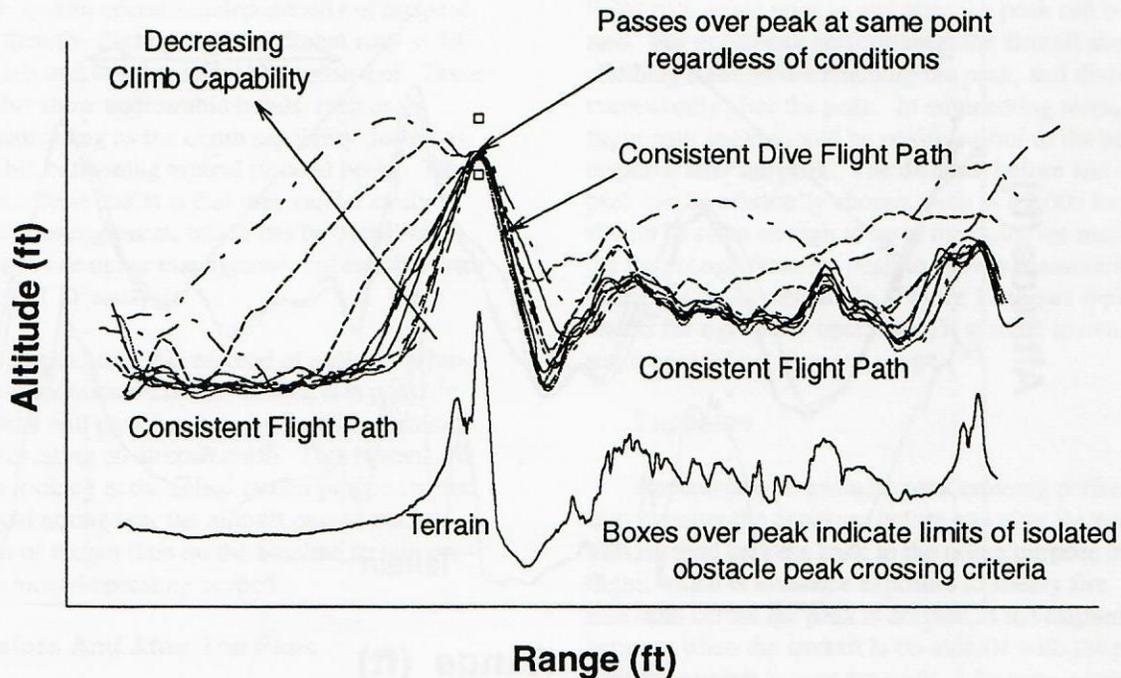


Figure 13. Modified Black Route Overlay Plot

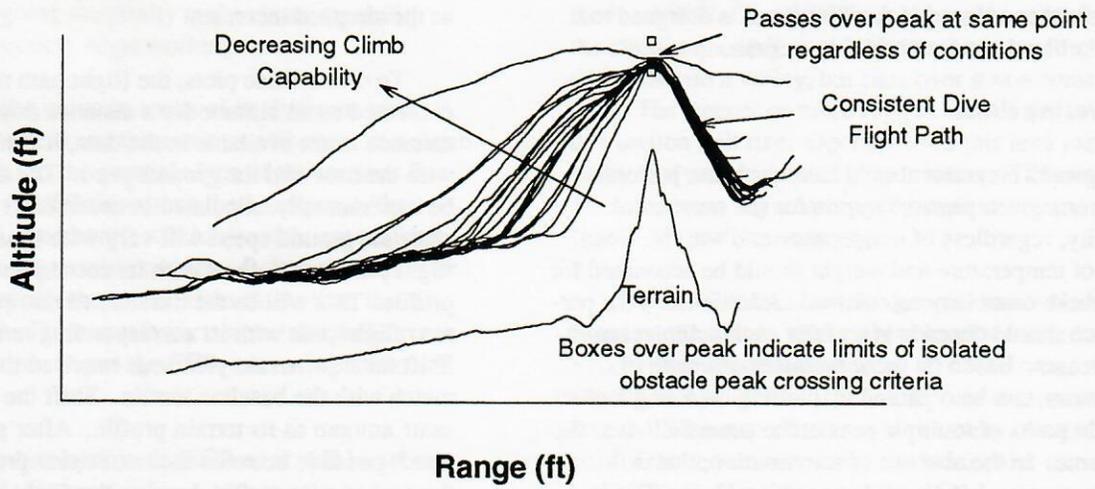


Figure 14. Red Mountain Route Overlay Plot

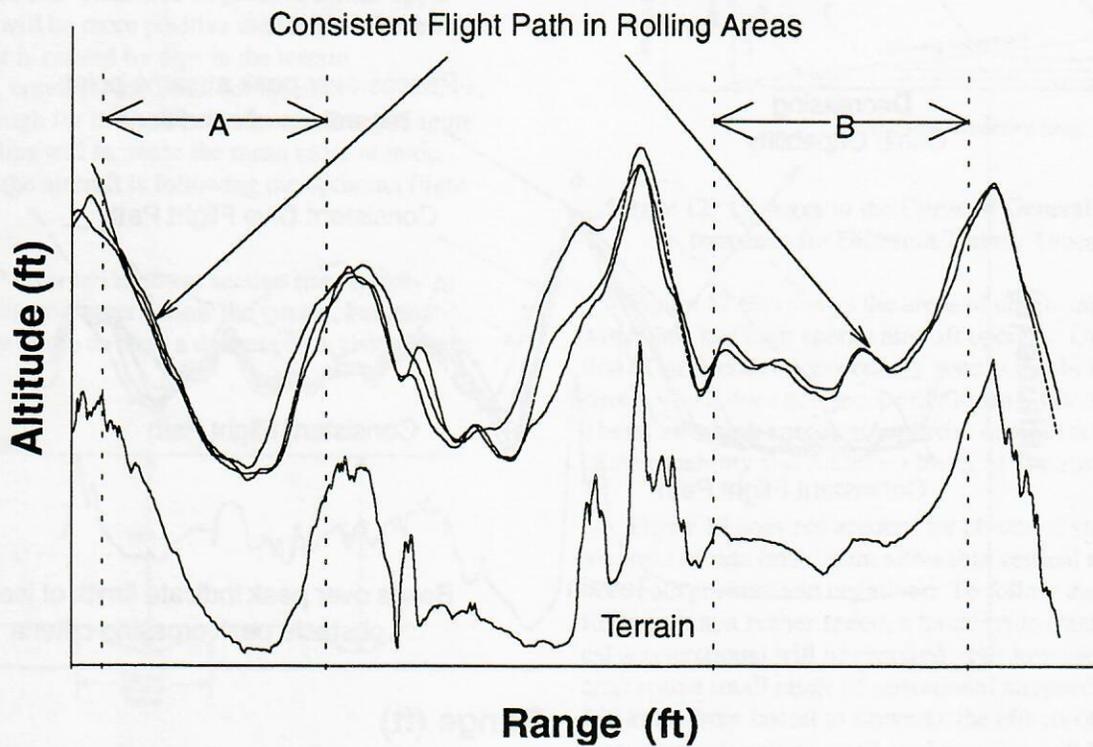


Figure 15. Saltdale Route Overlay Plot

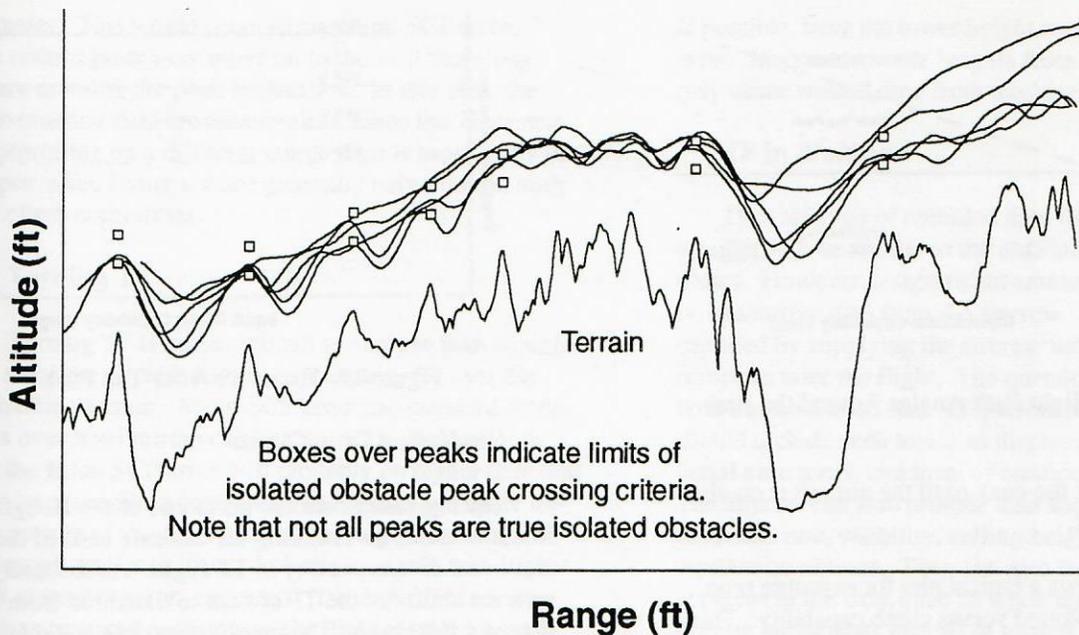


Figure 16. Rough 1 Route Overlay Plot

(climb capability, airspeed, weight), all flight paths should pass through the same point at the SCP over an isolated peak. Since the ability to dive is not dependent on the ability to climb, steady dives should be consistent regardless of flight conditions.

If the TF system operates independently of airspeed and weight, then the flight paths of different runs at different airspeeds and weights should be consistent. These charts will also show undesirable trends, such as an increase in staircasing as the climb capability decreases, or unacceptable ballooning around isolated peaks. The best part about these charts is that they can be easily understood by nonengineers, which can be useful when briefing aircrews or upper management not familiar with the intricacies of TF analysis.

As with simulations, this method of analysis is limited by three-dimensional effects. In fact, it is possible that flight paths will pass through the baseline terrain, seemingly indicating an aircraft crash. This is normally explained by looking at the actual terrain profile for that flight path, and noting that the aircraft passed over a lower section of terrain than on the baseline terrain profile, and was indeed operating properly.

FPA Before And After The Peak

The classical method of evaluating ballooning by measuring the flight path angle at the peak is not really

sufficient to totally define good peak performance. As Figure 1 shows, it is theoretically possible to meet the peak crossing criteria and still have terrible TF performance. Additionally, a noisy vertical accelerometer can make it difficult to evaluate the pushover maneuver. To further characterize TF performance at the peak, the flight path angle prior to and after the peak can be evaluated. For good peak performance, the aircraft should be climbing right before reaching the peak, and diving immediately after the peak. In engineering terms, the flight path angle should be positive prior to the peak and negative after the peak. The distance before and after the peak can be arbitrarily chosen, such as $\pm 2,000$ feet. It should be close enough to be in the pushover maneuver, but far enough from the peak to show a measurable change in flight path angle. Figure 17 shows typical results for a properly operating TF system, assuming a small operational airspeed range.

Exposure

Another way to evaluate peak crossing performance is to measure the exposure before and after the peak. This method harkens back to the prime purpose of TF flight, which is to reduce exposure to enemy fire. Exposure time before the peak is defined as the elapsed time between when the aircraft is co-altitude with the peak until the aircraft is over the peak. Likewise, exposure time after the peak is defined as the time between when

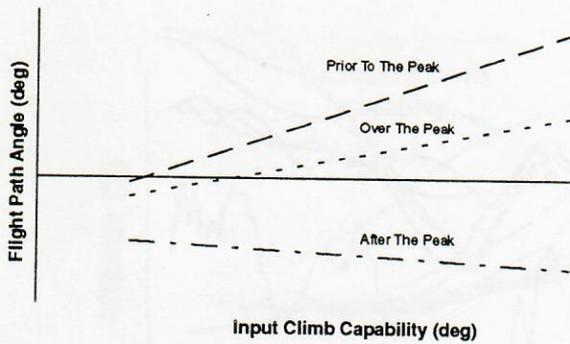


Figure 17. Flight Path Angles Around the Peak

the aircraft is over the peak until the aircraft is co-altitude with the peak.

Figure 18 shows a typical plot for exposure time prior to the peak plotted versus climb capability. This figure assumes a small range of operational airspeeds and ride hardness, such that airspeed effects are not seen. The exposure prior to the peak is dominated by the climb to the peak, and as this climb gets shallower, the exposure time goes up. Exposure time also increases at higher SCPs since the entire flight path is shifted up. Trigonometry shows that the exposure time varies with the cosecant of the actual climb angle used by the TF system (assuming constant speed), reaching infinity at the climb angle where the TF system stops working.

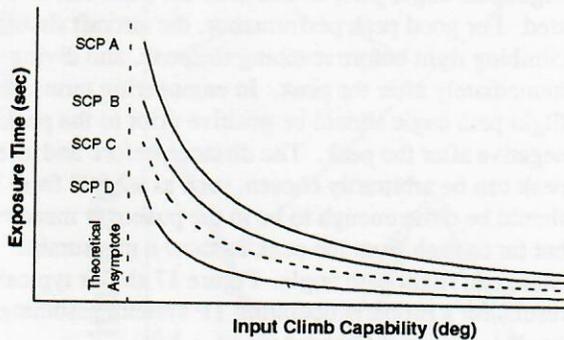


Figure 18. Exposure Before The Peak

Figure 19 shows a typical plot for exposure time after the peak plotted versus climb capability. After the peak, the aircraft pushes into a dive. The ability to reach the steady-state dive angle is not dependent on the ability to climb, so no variation with climb capability is expected.

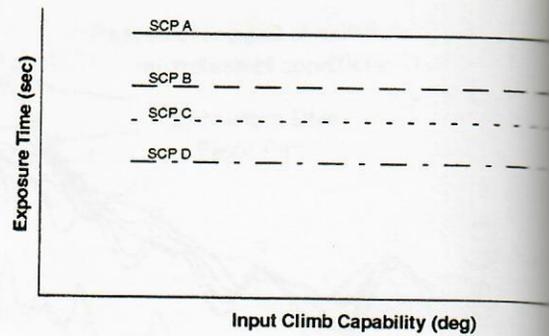


Figure 19. Exposure After The Peak

Big Valley Dive/Climb

The big valley near the north end of the Rough 1 Route is useful for checking the extreme ends of the climb and dive capability in TF flight. Additionally, it tests the ability of the TF system to transition from a full dive to a full climb. The steady-state dive and climb flight path angles are plotted versus climb capability (Figure 20). Dive data are compared with the line of the maximum commanded dive angle. Climb data are compared with the actual climb angle used by the TF system. If both sets of data match with these theoretical lines, then the TF system is performing properly. Note that at very low climb capabilities, the dive angle will approach zero, then be undefined as the aircraft does not dive into the valley.

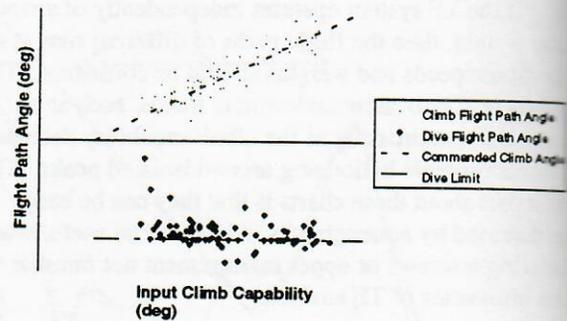


Figure 20. Flight Path Angles in Rough 1 Valley

Rough Terrain versus Isolated Obstacle

It is tempting to treat rough terrain as a string of isolated obstacles by looking at performance over selected "peaks" on the data traces. This is generally a fallacy, as the peaks are anything but isolated. Peak crossing data are only meaningful if the TF system is controlling on the peak up until immediately prior to crossing the peak. In rough terrain, the control point may be on higher terrain to the side of the aircraft, just far enough away that the controlling terrain is not picked up by the radar

altimeter. This would seem to give high SCP errors. The control point may move on to the next peak long before crossing the peak in question. In this case, the peak crossing data are meaningless, since the TF system is controlling on a different point. This is most likely to happen when flying toward generally rising terrain with low climb capabilities.

Turning TF

Turning TF is more difficult to analyze than straight TF, since the aircraft does not necessarily fly over the controlling terrain. Mean SCP error and standard deviations over level terrain can be used with caution. Note that the mean SCP error will probably be higher than that over level terrain in straight flight, caused by higher terrain outside of the turn. Turns into ridgelines can be evaluated similar to isolated obstacles, again with higher peak clearance errors expected. Turning flight results can be compared against straight flight results for similarities with explainable differences. As the terrain gets rougher, TF performance moves to the left on the General Curve of TF Performance (Figure 12). As with straight flight, as the turning TF performance shifts to the left, the analysis relies much less on plotting parameters and much more on studying flight path profiles and using engineering judgment.

Altitude Transitions

Altitude transitions should be tested over all types of terrain. However, overshoots of the new SCP caused by the transition can only be separated from normal climb and dive commands over level terrain. Over other types of terrain, look for a lack of transient behaviour at the transition and a smooth transition from one SCP to the next.

Towers

TF performance over towers can be analyzed similar to isolated obstacles. However, the tower will typically not show up in the radar altimeter data. It is critical that the exact time that the tower is overflown is recorded, so that the tower can be located on the data. The tower height will also be required to determine the SCP error.

If possible, have the tower height measured by a survey crew. Inaccurate tower heights from aeronautical charts may cause wasted time from tracking down nonproblems.

TF in Weather

Data analysis of recorded data for TF in adverse weather will be similar to the data analysis for other routes. However, a significant amount of extra data will be qualitative data from the aircrew. This can best be captured by supplying the aircrew with questionnaires to complete after the flight. The questions should address both human factors and TF performance. Questions should include such topics as displays useability, situational awareness, and level of confidence in the system. The aircrew can also provide data such as the approximate rain rate, visibility, ceiling height, and other weather conditions. They can also interpret details not apparent in the data, such as when the aircraft was controlling on weather instead of terrain.

Concluding Remarks

The nature and extent of any TF program requires that good planning be executed early in the program. A good set of specifications will guide the program, and proper instrumentation will allow for development and determination of specification compliance. Weather testing must be planned for, but will remain a target of opportunity. As with any program, good communication with the user will allow development of a system the user will be able to employ.

The TF testing can be executed in a safe progression with a well defined procedure for buildup in risk, for regression of software changes, and for good configuration control. Testing over all terrain types will ensure that the TF system is not designed to a single type of terrain feature. Parameter minimization techniques will allow a manageable test matrix to cover all aspects of the TF spectrum with acceptable statistical significance. In the end, these data, through use of the improved data analysis techniques presented, will allow the engineer to confidently answer the question "How well does the TF system work?"

Table 1 Edwards AFB Terrain Following/Terrain Avoidance Route Definitions

Route	Abbreviation	Type of Terrain	From	To	Straight/ Turn Point
Harpers	HRP	Level	N35 14.00 W118 00.00	N35 01.00 W117 07.00	Straight
Red Mountain	RMT	Isolated Obstacle	N35 00.57 W117 52.87	N35 20.93 W117 35.03	Straight
Haystack	HAY	Isolated Obstacle	N34 50.85 W117 49.20	N34 51.94 W117 37.34	Straight
Modified Black	MBK	Level to Rolling	N35 13.00 W117 45.00	N35 11.00 W117 02.00	Straight
Saltdale	SLT	Rolling	N35 19.00 W117 47.00	N35 03.00 W117 01.00	Straight
Modified Rough	MRH	Rolling to Rough	N35 15.00 W118 08.00	N35 44.00 W118 00.00	Straight
Rough 1	RH1	Rough	N35 15.00 W118 08.00	N35 55.00 W118 08.00	Straight
Dumont Dunes	DUN	Sand Dunes	N35 33.10 W116 11.35 (Approximate)	N35 40.75 W116 13.85 (Approximate)	Straight
Salton Sea	SLS	Land-Water Transitions	N33 17.50 W116 02.00 (Approximate)	N33 28.50 W115 44.00 (Approximate)	Straight/ Shoreline
Harpers Turning	HRT	Level	N35 14.00 W118 00.00	N35 01.00 W117 07.00	N35 04.90 W117 19.85
Ridge Line Turning	RGT	Ridge Line (Isolated Obstacle)	N35 15.60 W117 48.70	N35 18.50 W117 45.30	N35 18.50 W117 45.30
Saltdale Turning	SDT	Rolling	N35 19.00 W117 47.00	N35 03.00 W117 01.00	N35 09.70 W117 20.00
Rough 1 Turning	R1T	Rough	N35 27.86 W118 11.72; N35 30.42 W118 21.35; N35 35.05 W118 13.35	N35 30.42 W118 21.35; N35 35.05 W118 13.35; N35 27.86 W118 11.72	N35 30.42 W118 21.35; N35 35.05 W118 13.35; N35 27.86 W118 11.72