The Advanced Wild Weasel

Jerry Looper, 2000

This paper is for my children, so they can get a sense of what I did when I went to work each day when they were young. It is specifically dedicated to Bob Sumpter, who was my supervisor during much of this time, and is the best boss and the best friend I ever had.

Note: This is a scan of the paper Bob Sumpter provided to the Society of Wild Weasels on 8/24/18.
Victor Ballanco

FOREWORD

This paper describes, in some detail, one of the most interesting and important systems that I worked on during my career in aerospace engineering. It was from 1968 through 1976, at McDonnell Aircraft in St. Louis. The objective is to provide a little more insight into the actual workings of what was probably a fairly typical avionics system design, development, and flight test program during those years. During the development phase, I was responsible for "system integration", which at MCAIR means defining performance, aircraft installation, interfaces, and test requirements. During the flight test phase, I was responsible for the complete flight test program.

When I left the program in 1976, the team gave me a picture of the test aircraft which is annotated "For Jerry, who was present at the creation, and knows...". This was in recognition of the fact that I was fortunate enough to be part of the team from the first day. I was involved in almost all the fundamental design and mechanization decisions so I was one of the engineers who knew how and why the system was supposed to work.

HISTORY

The concept of "Weasel" comes from the Brits in WWII. The Germans had AAA (anti-aircraft artillery) radar on the coast of France and the RAF installed a detection system in some of their Mosquito fighters. The system had enough directionality that the signal got stronger when the aircraft was pointed at the radar. Having located the radar, they could bomb the radar site. They called the squadron "Ferrets" after the little animals that chase rabbits out of their holes, and a ferret is just a weasel.

RADAR AND COUNTERMEASURES

Typical ground-based radars of that time can be considered to be the electromagnetic equivalent of a flashlight, in that the energy is focused into a narrow beam, which is scanned around the sky in a pattern likely to intercept targets. When the beam hits a suitably reflecting target, like an aircraft, some of the beam's energy is reflected back to the radar. The elapsed time for a pulse of radar energy to travel to the target and bounce back to the radar is easily converted to target range. Noting the pointing angle of the focusing antenna of the radar when a target is detected provides the angle from the radar to the target.

Of course the RAF also used jammers to disrupt the radar signals, but jamming must be done each mission and radar techniques (counter-countermeasures) were quickly established which negated the jamming. Then the jamming equipment would have to be modified to counter the new radar techniques. It is more effective just to eliminate the radar site with bombing. One of our USAF pilots, who had flown Weasel missions in Vietnam said "Jam 'em today, jam 'em tomorrow. Kill 'em today, no problem tomorrow." It is, of course, also possible to defeat the Weasel concept by just turning off the radar, thereby eliminating the signal that the Weasel is intercepting, but this also renders the radar itself useless, which is counterproductive. So it becomes a back-and-forth game of measure, countermeasure and counter-countermeasure. Another favorite countermeasure is to use more than just one antenna, so that if one is taken out, it is easy to switch to another antenna. This is another reason for bombing the whole radar site, eliminating transmitters, receivers, displays, and operators.

EARLY ATTEMPTS

The USAF had three different "radar homing and warning" (RHAW) systems flying on two-seater F-100Fs, F-105Fs, and F-4Cs very quickly after the mid-1964 shoot down of an F-4 by the Russian-made SA-2 surface-to-air missile (SAM) in Vietnam. These anti-SAM systems were developed and deployed very quickly and were necessarily simple, although complex enough to require the second crewman to monitor and operate them. In fact, the first contract for such a system was a block diagram, a simple schedule, and a cost written on a blackboard and signed by USAF and the president of the company (Applied Technology, Inc).

USAF put Strategic Air Command (SAC) bomber electronic warfare officers (EWOs) in the back of the first two-seat aircraft that was available, ie, the F-100F. Using SAC EWOs was a logical move since these highly trained and motivated aviators had already been trained on the characteristics of the Soviet radar threat. Unfortunately, the F-100F aircraft was too slow to accompany a tactical strike mission, so they had to put the equipment into the strike aircraft, ie the F-105Fs and F-4s. These early systems provided a rough (within 15 degrees) azimuth to the radar as well as an approximation of signal strength from which the range to the radar could be estimated. The sensitivity of the system could only detect the main beam of the radar, ie it could only detect radars that were already tracking that aircraft.

"Weaseling" involved flying near the air defenses, typically a combination of SAMs and radar-controlled AAA guns, and luring the defenders into locking their radars onto the Weasel aircraft. The primary counter-SAM tactic was to launch the AN/AGM-45 Shrike ARM missile against the SAM radar site, but the early RHAW systems didn't provide range to the site accurately enough to prevent trying to kill SAMs at ranges farther than the Shrike could fly. What was obviously needed was a new system that could accurately determine the range to the radar, as well as detecting the radar signal when its beam was directed elsewhere, ie "sidelobe" detection. This meant that the angular accuracy and the system sensitivity needed to be dramatically improved. A better system, capable of finding targets accurately enough for ballistic bombing and to be coupled with a more capable ARM missile, was the objective of the upcoming effort.

REQUIREMENTS

In 1968, the USAF came to MCAIR to negotiate an engineering study to determine the feasibility of integrating the new AN/AGM-78 Standard Anti-Radiation Missile (Standard ARM) on the F-4D aircraft. During that negotiation, it developed that the USAF really wanted to address a completely new radar detection and location system fixing all the shortcomings of existing systems. They wanted to establish the performance requirements and to address integration into the aircraft, as well as cost and schedule issues. The original Weasel aircrews had served their tours in Vietnam and had moved on to positions of relative responsibility in the Pentagon and elsewhere and they had finally pushed the need for better equipment up the command structure.

During the first meeting, we outlined a set of 31 "performance requirements", ie what they wanted the new system to do. Some of these requirements were numerical and some were just functions they wanted included. Others were approximations and estimates of the technical performance parameters they thought they needed. Our task was to consolidate and refine these requirements into a framework around which we could mechanize the system and integrate it into the aircraft. The aircraft to be used was the F-4D, since it was the primary air-to-ground aircraft used by the USAF at that time.

As we were defining system requirements, we were visited by various combat veterans, and they impressed us with the need for the system as a critical element in attaining/maintaining air superiority over the battlefield. They also provided valuable insights into the operation of air defense systems as well as the tactical operations being used to counter them in the field. They were always thoroughly professional, and being fighter pilots at heart, they were a joy to be around. Their contributions were very valuable.

The mission was considered to be so dangerous that only volunteers were accepted. One of the SAC EWOs, Jack Donovan (called "Supermouth" for obvious reasons) said, upon being first briefed on the details, "You gotta be shittin' me!". Abbreviated 'YGBSM', this became the rallying cry of the "Tiger in the front seat and the Bear in the back seat" Weasel crew. But they also took great pride in being "First in - last out!" This means that they were the first aircraft over the target area and they stayed until all the strike aircraft were off, ie last out. The NATO code name for the Russian SA-2 radar was "Fan Song", and some of the Weasel aircrews wore shoulder patches: "Song Critic". They also had a distinctive flight scarf of black and yellow checks. I have one.

MECHANIZATION

We quickly determined the critical performance drivers and began to address the system mechanization. The three critical performance requirements were system frequency coverage (to cover the radars of primary interest to tactical air operations), sensitivity (to permit sidelobe detection) and target location accuracy (to permit the system to locate the radar on the ground accurately enough to enable ballistic bombing even if the radar turned off its signal). Shooting ARMs at the radar is easily countered by turning off the radar until the ARM has detonated, and the USAF pilots knew that since ARMs are expensive weapons of limited quantity it would eventually be necessary to drop "universal, all-frequency, 500-lb, lethal jammers", ie bombs, on the radar sites. Radar sites must have antennas above ground, so they are very vulnerable to bombing if they can be located through their camouflage. In Vietnam, the SAM sites were laid out according to Soviet doctrine with typically six prepared positions at a site, although all six were usually not active. Access roads made a very easily identified pattern, except that they were usually under the jungle canopy or similarly hidden.

Frequency coverage involves establishing enough receiver resources to detect the required frequencies. As the system requirements and design were firming up, we made many trips to WPAFB (Wright Field, Ohio) to discuss the characteristics of the Russian radar deployment. All this data was highly classified and changed as new Russian radars were found to be deployed in Vietnam. USAF had a secure facility at WPAFB where they documented the characteristics of all the Russian radars. Two (for security reasons) MCAIR engineers (usually Norm Koester and I) visited this facility repeatedly over many months gathering signal characteristics for the Russian radars we were expected to detect. We were careful, however, to tailor the system requirements to the Russian radar capability, not just to those deployed in Vietnam.

System sensitivity is set by balancing RF signal bandwidth, probability of detection, and receiver noise, so it is just a matter of being willing to spend the weight (a critical matter to aircraft since higher electronics system weight means less fuel and/or fewer weapons). RF signal bandwidth is determined by the characteristics of the radar pulses to be detected.

Target location is harder, however. To locate the target quickly (a radar might turn off after only a few seconds of operation), it is necessary to have much better angular accuracy and the angles must be in both azimuth and elevation. The earlier azimuth-only system doesn't tell us if the radar is below us or on the radar horizon, ie range measurement is very poor. Adding elevation measurement gives us a line through the radar, which by knowing aircraft elevation (assuming the earth is locally flat) gives us target range. This is combined with the multiple intercept lines through the radar which become available from subsequent measurements as the aircraft flies near the radar location to provide a good target range measurement. So the radar location problem becomes one of accurate azimuth and elevation angle measurement.

The classic technique for angle measurement is to compare the signal strength (amplitude) of the radar pulses detected by two antennas separated by some small distance on the skin of the aircraft (called 'amplitude comparison'). This mechanization is simple, but not accurate enough for good target location. So we selected phase comparison interferometry, in which the RF signal phase difference of the signals detected by two antennas on the skin of the aircraft is measured. Unfortunately, at radar signal frequencies, the 360 degrees of signal phase rotation occurs in a small number of physical degrees of azimuthal space, so it is necessary to add a third antenna in the baseline to resolve the ambiguity of measurement. The farther apart you install the first two antennas the more accurate the measurement, but the more ambiguities you get. So a baseline becomes three antennas in a line, with two mounted very close together.

The size of each antenna is inversely proportional to the frequency of the radar signal to be detected. We selected small circular spiral antennas, separated by about six inches, with the third antenna mounted between them. One baseline measures azimuth angles, so adding an orthogonal baseline takes care of elevation. Of course, what is required is orthogonality (not necessarily vertical/horizontal), so the whole antenna array is rotated 45 degrees to minimize the vertical dimension of the array when mounted on the aircraft.

PROGRAM ASSESSMENT

As we were defining the system structure and analyzing the system performance likely to be obtained, the USAF convened an official inquiry of experts to make a separate determination of whether it would be possible to install so complex a system in the F-4. After hearing our briefing and quizzing us individually, the expert committee concluded that we would probably fail. We finally convinced the USAF Program Manager that we had done enough analysis of the problem to convince ourselves that we would probably not fail, and since the need was acute, the program continued.

SYSTEM DESCRIPTION

We divided the system into hardware and software elements to facilitate procurement from our avionics suppliers. The system was made up of receivers/antennas, a signal processor unit, a central computer, aircrew displays and a recorder to provide a retrievable record of the characteristics of the signals detected. Since each antenna array can cover only 90 degrees around the aircraft, adding three more arrays covers the entire space around the aircraft. So we had one array facing forward and one looking to each side mounted in a chin location under the main aircraft fire control radome, and the fourth mounted on the tail looking aft. It is not necessary to point an array up or down, because there are no radars directly above the aircraft, and any radars directly below the aircraft will have already been detected by the forward antenna array as the aircraft approaches the target. We rearranged the main instrument panel in the aft cockpit to accommodate a PPI (plan position indicator, or 'God's eye' view of the radar situation), as well as small displays of time-based and frequency-based signal information. A simple pointing or 'homing' display was installed above the pilot's glare shield to give the front seater steering cues to the target. We also modified the pilot's gunsight to provide a red dot over the calculated position of the radar antenna. This also gave him a bombing reference point if he was unable to see the target.

We decided that the computer software should be written by the computer manufacturer based on specific requirements and algorithms provided by the receiver supplier. Procurement specifications were written and suppliers were competitively selected. The selections of suppliers involved an identification of companies who could conceivably design, build, and test the various parts of the system. We convened a supplier orientation meeting at which 50+ prospective avionics suppliers were introduced to us and to the system performance and schedule requirements as defined in the specs. Competitors were then given a number of weeks to submit their proposals, which we evaluated for performance, cost, and schedule.

ORGANIZATION

The first step in setting up any engineering organization is to procure coffee cups for the program. We did this, and then we were ready to really get started.

The avionics engineering effort was organized into three basic groups, reporting to Bob Sumpter, who, in turn, reported to a Project Engineer, who reported to the manager of all F-4 programs. I was in charge of the "Integration Group", whose responsibilities included defining the system performance requirements, managing the interfaces between suppliers' equipment, and defining and executing the system level testing of the hardware and software, whether in the lab, ie on the test benches, or in the aircraft. This included, of course, both ground and flight testing of the system in aircraft.

The second organizational group was the "Subsystem Managers" who were responsible for selecting suppliers and for the execution of the subcontracts with those suppliers. MCAIR, in general, did not make electronics equipment - we bought it from our supplier network. This is a technical management task, requiring close scrutiny of the suppliers' plans and performance, including technical characteristics, schedule and cost. They held formal design reviews at critical points in the development process to assess progress and to mandate improvements if necessary.

The third major organizational element was the "Software" group, responsible for design, coding, and testing the operational software to be run in the system computer to control the system sensors and process the resulting data.

Since I was necessarily much more involved with the Integration process, I will focus on that part, not to imply that Subsystem Management and Software were not equally important and complex. Integrators were not more important, just different.

STAFFING

Acquiring and retaining enough engineers to execute the program was a problem at MCAIR at that time. The F-15 engineering program was at full speed and they had, for all practical purposes, their choice of personnel. In fact, during a particularly hard period on that program, the president of the company got on the PA system and, as a morale booster for them, said that management had complete confidence in their team since "they have the top 750 engineers in the company working on the F-15" at which time, Charlie Tolson and I both jumped up and claimed to be "number 751".

SYSTEM DEVELOPMENT

Contracts were negotiated with the selected suppliers and the system was slowly developed. The management approach used was a series of Technical Coordination Meetings, held monthly and usually alternating between MCAIR and the supplier's facilities. These were formal meetings, including an agenda, action items, and minutes, concerning all aspects of system and interface development, schedule and cost. The suppliers built, tested and delivered to MCAIR a few prototype sets of equipment. (Prototype in this context means built by hand as contrasted with production rate manufacturing.) Some prototype systems were retained at the suppliers' facilities for qualification testing to assure that they would survive in the harsh environment of the fighter aircraft, as well as to assure they had the necessary reliability.

It turned out that the computer supplier (Texas Instruments) couldn't reasonably accommodate the rapid changes in software requirements arriving from the receiver supplier (IBM), so MCAIR set up our own software team to handle the computer software. This was the most complex computer software problem ever attempted on a fighter aircraft at that time (1970), and we had concluded very early that for maximum flexibility as many of the system functions as possible should be computer-controlled. Keep in mind that this was fifteen years before PCs became popular on desktops.

AIRCRAFT INTEGRATION

As mechanized, the system included 52 antennas and weighted a little over 400 lbs. (Keep in mind that at that time, the typical avionics unit weighed 25-30 lbs.) It was a major avionics system and as such, was significant hardware to put in an aircraft which is already in high-rate production (indeed, the production rate for F-4s got up to 72 per month at one time). When it was decided to use the F-4D as the developmental aircraft, it became necessary to get permission from the Project Engineer responsible for USAF F-4 development. He was totally absorbed in producing the F-4E at this time, so when we assured him that we would only be working on one D-model aircraft, he agreed to make the necessary aircraft engineers available for the aircraft mod design. The aircraft itself was borrowed ("bailed") from the USAF. Eventually two F-4Ds were modified for flight testing.

The next problem was where to stuff the 400 lbs of equipment into the aircraft. The system antennas needed unobstructed views of the outside world, so the only reasonable locations were under the nose radome (the F-4C already was using that location for RHAW equipment, so since our equipment would be a direct mission substitute, the changeover was easy. We briefly considered using the wingtips for antennas, but there is not enough structure out there to support the equipment weight with the necessary rigidity. We decided to put the aft-looking antennas on the top of the tail after some discussion of maybe using the drag-chute door between the engines. The remaining equipment was located in a new avionics bay immediately behind the aft crewman. Of course, the controls and displays were located on the main instrument panel in the aft cockpit for the EWO.

INTERFACE MANAGEMENT

A complex avionics system has many connections among its own boxes and among the other systems on the aircraft. Signals and power for avionics boxes are distributed in wire bundles which are routed through the aircraft. Keeping track of these signals, and assigning them to specific wires, is part of "interface management". The AWW was among the first, if not the first, system on the F-4 to organize and document wiring information, ie type of signal, characteristics, type of wire, and destinations, using computer listing techniques. Prior to this system, interface wiring on the aircraft was documented in wiring diagrams, typically four feet high and fourteen feet long, completely hand drawn, and therefore error prone, not only in generation, but also in use.

TEST PLANNING

As the system design was coalescing and the suppliers were building the hardware, we also began to determine how we would test the system in the aircraft. It is possible to test many elements of the system in the lab, but since proper operation depended on measuring external signals as well as determining the specific location of the aircraft it is also necessary to have a flight test program. So the test program becomes a combination of supplier testing of each subsystem, then integrated lab testing of the combined system at MCAIR, then installing the system into the test aircraft and testing for proper performance within the limitations of ground operation, then testing the system on aircraft engine power, then testing it in flight.

The major issue with adequate system testing was that the signals to be detected, etc, were really from Russian radars, the specific characteristics of which were highly classified. We had established a complex system for getting the necessary classified design data from USAF, but the only cost-effective way to adequately test the resulting system in flight was to use a Gov't test range. The USAF had a full set of Russian radar simulations on Santa Rosa Island at Eglin AFB in Florida, but that facility was being used full time to train aircrews on the existing F-105 and F-4C radar homing and warning systems and it could not be spared for developmental testing. "Don't you know there's a war on?" Fortunately, the US Navy had a good set of Russian radar simulations at their facility at China Lake NAS in California, just north of the USAF flight test facility at Edwards AFB. Coordination with USN at China Lake quickly established that they would be glad to bail out the USAF in their hour of need. Inter-service rivalry in this case was a positive influence on the program.

TEST REQUIREMENTS

The major functions of the system (by now officially known as the Advanced Wild Weasel System) were RF signal detection, radar classification, and target radar location, so these were the focus of the flight test. We had to show that the system would detect all the necessary radar signals across the complete frequency spectrum of interest, that the system would then determine what kind of radar was being detected, ie Search (or early warning), SAM, AAA, AI (airborne intercept), etc. (and whether that radar was tracking this aircraft) and that the system would locate the radar position accurately.

TEST MANAGEMENT

I was responsible for the system test program, and it was very advisable to include the USAF in the testing as much as possible, so we agreed that the USAF could provide a system operator, ie a back-seater, for every other test flight although he would be required to fly the 'test cards' that MCAIR defined (with USAF help, naturally). In addition I made a deal with the MCAIR chief flight test engineer at EAFB. He would tell me if the aircraft was ready to fly and if not, why not, and when it would be ready. In turn, I would tell him the same about the electronics system. We agreed to tell each other everything and also not to withhold any information from the USAF reps. This policy was exactly right, and it carried us through some hard times.

LABORATORY TESTING

Each of the hardware subsystems was tested at its respective supplier's facility within the limits in simulated interfaces and simulated aircraft dynamics. But before we could put the equipment into the aircraft it had to be shown to work properly as a system. This is bench integration testing and was done on test benches in the MCAIR lab facilities in St. Charles (very convenient for me from a daily commuting traffic standpoint). This was when all the various subsystems were hooked together and the operational software was loaded and tested. This process always takes longer than management expects it to, and this program was no exception. We had many exciting meetings with program management wanting us to move to the aircraft, but there is no sense putting equipment into an aircraft, where checkout and maintenance are difficult, until you know it will work. It takes about three times longer to do anything on the aircraft than it does in the lab. Eventually the system was good enough to install in the aircraft.

AIRCRAFT INSTALLATION AND GROUND TESTING

Before we could bring the equipment from the lab and install it in the aircraft, we had to "ring out" the wires, ie be sure they were installed properly and that the connectors were wired correctly. This was done by the avionics lab engineers, then the equipment was installed. We tested the system in the aircraft, under aircraft power for the first time, to be sure that it didn't interfere with other aircraft systems, and vice versa. (You don't want fuel to transfer or the landing gear to start cycling when you turn the system on.) This "electromagnetic compatibility testing" takes a week or so to go through all the systems, but it was finally completed. During this time we were also trying to get the system to detect the various air traffic control and weather radars around the St. Louis airport. It was a big day when we got the software tuned up enough to begin to see the correct symbols for the airport radars on the displays.

FLIGHT TEST FACILITIES

When we had done all the supplier tests, integration and lab tests, and aircraft ground tests we could, we had the aircraft flown to EAFB, where we were assigned to a hanger on 'Contractor Row'. They also gave us a couple of offices, which we shared with the USAF guys. I didn't even have a desk, I just carried the few papers I needed in my briefcase. We also got some space in the USAF Avionics Lab down the road where we could set up our system integration test bench to work on problems separate from the aircraft. (It takes three times longer to isolate and fix a problem on the aircraft than on the test bench.) The test missions were flown over Echo Range at China Lake NAS, a few miles north of EAFB. Most of the Echo Range radars were located in a central valley in the desert, although a few were located in even more remote locations, including a hill top called 'Slate Ridge' and another location called 'Aluminum Ridge' because of all the aluminum foil chaff they dropped there during jammer tests.

We quickly learned that it was necessary for MCAIR to have an engineer on-site at Echo Range to be sure the correct radars were operated in the correct modes at the correct times and to make adjustments as necessary in real time if necessary due to equipment failures, etc. So an hour before each test flight, one of us (usually including one of the USAF test engineers) drove through the desert to Echo Range, using one of the company cars.

LIVING SPACES - EAFB

I had not originally intended to be at EAFB, but the engineer who had been brought on to the program for this assignment decided at the last minute not to go. The only logical solution seemed to be for me to handle it. Most of our engineers brought their families to EAFB. I did not, since Lura had taught in the California school system and we didn't want our kids' education to take a step back. Plus, moving a family with four young children for only a year seemed counterproductive. All our engineers and their families lived in Lancaster, Calif, which is a 45-minute drive to EAFB. I elected to live in a motel in California City, which was only 25 minutes north of the base. I did not get a phone, because I knew that MCAIR and USAF senior management would not be able to resist calling me at all hours of the morning from back East to "check on things". My supervisor Bob Sumpter had enough discipline not to do that, but the other managers had already shown that they did not.

Since we were strangers in a strange land (the California desert), we made our own entertainment. I found out that two of the engineers on the program also played guitar (not all that good, but better than I did), so on Friday evenings we would adjourn to someone's house for dinner and then "pick and grin" until we tired out. This practice culminated in the "Great Columbus Day BBQ". We decided to BBQ a 90-lb pig for Columbus Day, so we dug a pit in the back yard (dirt) of Don Funderburk's house, lined it with sheet tin, added 90 lbs of charcoal, spitted the pig with a metal rod, wrapped the pig with chicken wire to keep the meat from falling off, and turned the pig 45 degrees every 15 minutes for 12 hours. Then, at about midnight, we removed the meat from the bones and had a nice dinner. All the MCAIR, suppliers, and USAF people and families attended. At about 2am, we decided to have a music program, so we got out the guitars, amplifiers, mikes, etc and did so. That was a long day.

Another benefit to living in Lancaster was that there was a great pizza place at 20th and K street. Plus there was a bar in the desert halfway between the base and Lancaster serving cold beer. The name of the bar was "The Office". It was good to stop at the office after a 120 deg day.

PREPARING THE SYSTEM FOR FLIGHT

The system was either down for failure repair, or was being updated with system changes (usually software) from St. Louis, or was ready to go. Failure repair was done by lab engineers including some reps from the major suppliers (IBM and TI). We worked two shifts - the day shift to fly and the night shift for maintenance. Of course the shifts overlapped to the point that most were working 10-12 hour days. The major troubleshooting tools for the system were its own "built-in-test" functions which injected signals behind the system antennas, some small "squirt boxes" that radiated basic signals into the antennas, and a computer memory loader with which we could check the contents of any location in the computer software.

We had about one spare for each electronics box in the system installed in a test bench in the Avionics lab at the base. There was usually another unit on benches in St. Louis, being used to develop new software so we had to strike a balance between calling for spare hardware and fixing what we had on site. There are never enough spares in a developmental flight test program, so this was a continuing crisis. This is a universal problem with all flight test programs, since each system costs money and time and management is hard to convince that buying spare equipment is cost-effective relative to the cost of maintaining a test operation at a remote site.

If we needed to update the system mechanization, we would get new hardware or software from St. Louis. It was necessary that each change be successfully tested in St. Louis before it was sent to us. Most of the changes were computer software program changes, typically coming to us as "patch tapes", ie short punched Mylar tapes which we could load into our aircraft computer. Sometimes, due to overriding necessity of time or performance, we fell back to "hand loads", ie specific hexadecimal (base 16) numbers telephoned to us to be very carefully keyed into our computer software one number at a time. This was error-prone, naturally, and was frowned upon by all. I would make the decision if the system was ready enough for a meaningful test flight.

PRE-MISSION

When we determined that it would be productive to conduct a test flight, we would decide what specific system functions were to be tested, what Echo Range target radars we needed and what flight profiles we wanted flown. Then we notified the USAF at the base, all program staff, the aircrew, and China Lake. Then the specific flight instructions for the aircrew would be entered onto "flight cards". About an hour before take-off, we would hold the "pre-flight briefing" which is a semi-formal presentation of the flight cards to the aircrew, showing what data we needed to collect, the status of the system and the aircraft, complete with a description of all changes and maintenance done since the last flight, and describing any limitations to the aircraft operation, such as the prohibition from flying over certain locations for various reasons. After this briefing, a couple of us would get a company car and race up to China Lake like a scalded dog. We had a good enough relationship with China Lake management and security by that time that we could safely go through the "back entrance" to their facilities.

TEST FLIGHTS

It was normal to get 8-12 test points done on a given flight. We had only limited range time each flight since the Navy test facility was also used for other test and training programs.

Of course, the system was rarely completely "up", so it became a matter of balancing the operating functions, and the need to test them, against the value of continuing maintenance. Our original flight schedule objective was two test flights per week. We found at first that the system didn't work well enough against the China Lake radars to support that schedule, then when we eventually got to that flight rate, St. Louis could not keep up with their data reduction. So we probably averaged a little less than that. We had capability both from China Lake and from MacRadio at EAFB to talk to the aircrew in flight. We evolved some simple codes to enable us to discuss cockpit indications, or lack thereof. These aviators were professionals and took great pride in being able to get the test points called for in the flight cards with a minimum of input from the engineers.

We soon evolved a more or less standard set of flight maneuvers, eg straight passes either over or beside the target radar or "J-turns" which flew by and then maneuvered back over the target. This made it easier to communicate our needs with the pilot and by varying altitude we could cover the various angles between the aircraft and the target. We soon found that the two hardest problems were the transition of the incoming radar signal between adjacent quadrant antennas on the aircraft and the separation of pulse trains from different radars transmitting at nearly the same frequency. Both these problems were eventually solved by Bob Mitori's software team, but it took time.

We always said that a successful flight was one in which you got the aircrew and aircraft back. Beyond that, we tried to get valid test data. If the data showed good system performance, that was an even greater bonus. It took a while to convince USAF and MCAIR senior management that just flying does not mean you have a good system. But Brockhagen and I had a firm deal that unless the system was good and we had something to test (my call) and the aircraft was good (his call), we wouldn't fly. It makes no sense just to burn holes in the sky.

POST-MISSION

After the mission ("sortie") was completed, the China Lake monitors would race back to EAFB to try to get to the post-flight briefing, at which the aircrew described what they had done, or not done, and what results they obtained. The maintenance needed for the aircraft and the AWW system was also a hot topic at these meetings, since failures interfered with data collection. MCAIR test pilots go through lots of training, in addition to the training most got at USAF or USN test pilot school. They were very professional and took great pride in flying the cards and reporting exactly what happened. At the end of the meeting, we would make a tentative plan as to what we would do and when we would fly again. Then we would all adjourn to the video tape player to play back the video of the major system displays with commentary by the crew. Based on these indications, we would decide what to do next. If we had flown early enough in the day, I would call Bob Sumpter in St. Louis and let him know what happened and what the next plan was. Usually it was too late in the day before we had a plan, so I'd call him the next morning. He was always a voice of calm reasonableness, even when we were running around with our hair on fire.

Most system failures could be isolated to a faulty unit while on the aircraft and then taken to the lab to isolate to fault to a circuit board, connector, or whatever. A good unit could be substituted in the aircraft and we were back in business, if we had a good unit at the base. If there was a good unit available in St. Louis, we could take advantage of the daily supply flight the F-15 program had organized. Sometimes we had to wait for good parts from the suppliers, in which case we either flew to test some other part of the system or stood down. Some failures were too insidious or intermittent for fault isolation on the aircraft by conventional means. In that case, we would pull the aircraft out of the shady hanger into the hot sun on the ramp, run up the engines, and see if we could repeat the fault on engine power. You had to wear Mickey Mouse ear covers and to take everything out of your pockets to be sure nothing got sucked up the engine inlets. This is hot, noisy, and dangerous. Did I mention noisy?

TEST DATA REDUCTION

We had the instrumentation tape recorder as well as the system computer recording flight test data. We also had over-the-shoulder TV cameras in the aft cockpit so we could run the video tape immediately after each flight and see what had happened. We had some of the aircrew's switchology instrumented also to be sure they did what they were supposed to do at the correct time. From the video we selected specific times that interesting things happened (or didn't) and then we sent the computer data for those time periods through the EAFB data system back to St. Louis, where the software developers could try to determine the results or the causes of failures.

INITIAL OPERATIONAL TEST AND EVALUATION

After we had done all the system testing that we could cost-effectively do and got the system operating as well as we thought it could, the USAF took the aircraft to Nellis AFB at Las Vegas, Nevada, for their own testing. This was done by USAF people, against their own radar simulations at their classified flight test facility. Nellis was really a training facility, not a developmental facility. The objective was to test the system under "operational" conditions, as contrasted with our "development" testing at EAFB. They could also use real bombs, missiles, etc, (although the real bombs were only 25-lb practice bombs). This testing was all done by USAF, with no input from MCAIR or our suppliers. The USAF was responsible for maintaining the aircraft. We took a few lab engineers to Nellis to be sure the system didn't hold up flight operations (it didn't).

The USAF was very touchy about their classified Nellis test range, so they were shocked when we printed out the locations and characteristics of all their radars (available in the system computer) and presented it to them. They still wouldn't let us into the test planning briefings, but at least we got some respect.

A highlight of the IOT&E evolution was that the system was accurate enough to permit us to put "Three in the shack", ie put three practice bombs in the exact center of the ground target, without visual reference to that target. This is "blind bombing", ie being able to hit the target radar even if it is so camouflaged that the aircrew can't see it.

Another highlight of IOT&E was when Supermouth Jack Donovan came out to Nellis and flew a mission over the range. When he got back, he just climbed out of the seat and was ready to "go back to war".

LIVING SPACES - NELLIS AFB

I didn't take many MCAIR engineers to Nellis since we were still flying at Edwards. They kept their families in Lancaster and drove up for a few days each week if necessary, depending on the system hardware's health. I was full time there, so I lived in the Desert Rose motel on the Strip. It was convenient to stop into a casino each morning, walk through the rows of slot machines, and get a cheap breakfast among the early gamblers.

AIRCRAFT CHANGE

About the time the USAF test team concluded that the system was ready for operational use, the F-15 program had firmed up to the point that the F-4E, which had been reserved for the air-to-air role, could be reassigned to air-to-ground (along with the F-4D which had been our designated aircraft). So USAF told us to make a few system improvements, install the system into a test F-4E in place of the existing 20-mm cannon under the radome and bring it back to EAFB, then to Nellis. The system, when deployed, would be known as the F-4G.

LAWSUITS

We were involved in two lawsuits as the program evolved. The first was with the Loral company in New York. Loral was the equipment supplier for the system's controls and displays. During the design phase, it became obvious that one of the interface functions was not being designed properly and would not meet its performance specifications. We pressed them for a redesign and they eventually responded with a lawsuit contending that our requirements were beyond the state-of-the-art. We contended that this was not true and that since they had taken a contract to provide the specified performance, they were liable for the redesign. The MCAIR Procurement Division, our contracting agency, convinced company management that it was necessary to fight the Loral lawsuit, so one of our engineers, Charlie Tolson, was detailed to provide our lawyers with the technical support and documentation required. After about two years of hearings, including deposing everyone, MCAIR finally realized that the lawsuit was costing more than was cost-effective, so the case was settled out of court.

The second lawsuit was with Applied Technology, who had contracted to provide a small signal analysis receiver for the system. As their design phase progressed, with very little design data or hardware, we finally concluded that they planned to present us with a lawsuit instead of the first hardware delivery, so we beat them to the courthouse steps, which got the lawsuit sited in St. Louis instead of California. This case was also eventually settled.

PRODUCTION DECISION

I decided not to accompany the F-4E back to flight test at EAFB because I had already done one tour in the desert and by that time I had been assigned as Program Engineering Manager of the whole program. It took a few more months of installation testing, developmental testing at EAFB, and operational testing at Nellis on the F-4E to get back to where we had been on the F-4D development. Eventually, however, we had gotten all the performance out of the system that the design would support, although the USAF still wanted more performance. I finally told them that even though they said it was "not ready to go to war", we had "all the performance the design would support, the system met all its specifications, and that they should either put this system into their F-4Es or go to war with what they already had in the F-4Cs and get their cans shot off." They decided to put the system into 112 F-4Es.

CHARACTERS

Any program the size of the AWW is bound to have its share of characters.

Barry Brockhagen was in charge of the flight test engineers. He sometimes wore a yellow suit, complete with tie, to formal meetings in St. Louis. He had spent some time in Saudi Arabia for the company and had many excellent tales to tell. He had the reputation in the company as being hard for avionics engineers to deal with, but we had a clean division of responsibilities and we showed him that we respected his opinion, so we got along just fine. Barry died of a heart attack at an early age, shoveling snow at his home in Lancaster.

Norm Koester was responsible for our knowledge base of the characteristics of Soviet radar systems. He was from Buffalo NY, so there never was a snow storm in St. Louis that he had not seen more snow in Buffalo. Norm's wife was from Latvia or Estonia, I forget, and I always wondered how he got a security clearance in those days. Norm and I did most of the coordinating about threat radar data with USAF, and when he left the company "for better opportunities", I inherited the job, so I had to know the frequencies, pulse repetition intervals, pulse widths, scan patterns, etc of about 300 foreign radar systems. There is nothing like trying to keep track of a bunch of numbers like that to give you a headache. Plus, every time a new system was deployed, I got a call and had to run up to Wright Field and convince the USAF that the system would handle that threat. Unfortunately, at one point the Russians deployed a new radar whose operating frequency was beyond our coverage, so we had to change the system to intercept those frequencies.

Alton Miilhouse was the TI (the computer supplier) rep on the program. When we would eat out, he would always order a baked potato with his steak. Then he would spend ten minutes mashing and stirring it into mashed potatoes. He said if he ordered mashed potatoes, he got dehydrated potatoes and he didn't like them. So we all finished our meals first, and then waited for him to finish. (His family name was originally spelled Mulhouse in the old country.)

Frank O'Donnell was a USAF major in charge of the EAFB contingent of USAF people. He flew every other mission as the backseater. Frank's degree was in mechanical engineering, but he was a good engineer and could tell a lot from the way the system displays behaved during test missions. Frank specifically requested that some specific elements of the system be instrumented for the USAF engineers. Brockhagen call this "FOD data" in honor of his initials. Frank retired from the USAF as a major and worked for TI in their commercial products operation.

John Sinnott was the Project Engineer, ie Bob Sumpter's boss. John wanted to spend tremendous amounts of time with the key engineers discussing the program, to no obvious benefit. He would just keep pushing for faster progress until we just told him "No!". That was fine. He didn't have any idea about what constituted reasonable progress, so he just kept asking for more until we stopped him. He was like the lady whose little boy was playing in the sand at the beach but was suddenly washed out to sea by a big wave. She implored God to return her son. The little boy was immediately washed back onto the beach. She then told God "he had a hat." Bob Sumpter made a major contribution to the program by keeping John off our backs at flight test.

Bob Mitori was the head of our software operation. He was ethnic Japanese and his family had suffered in the Utah camps during WWII. He was a pipe smoker in those days, like many engineers, and when asked a question that he didn't particularly want to answer, eg, when will the next software update be released, he would take out the pipe, clean it, find his tobacco, fill the bowl, tamp it down, find his lighter, then very carefully light his pipe. If, after witnessing this entire operation, you still remembered your question, he would give you a perfectly clear answer.

Don Funderburk was the classic engineer. He always looked for basic principles and exact information, such as when he tied a string to the hood ornament of his car so when he and Gerry Haworth drove from Lancaster out to the base across the desert, they could calculate the force of the wind by the angle the string make on the hood. He and, more accurately, Mrs. Funderburk were my hosts most Fridays for supper. Then Bernie Conway and Gerry Haworth would come by and we would play guitars for a couple of hours. Denise Funderburk had a lovely soft singing voice and we always threatened to get her a white cowgirl outfit and have her be the "girl singer" for us. She was decidedly not interested. Don retired from MCAIR in the mid-1990's.

Gerry Haworth was our software engineer. He had played guitar in a rock and roll band in college at Kansas State, so he had some interesting background. Gerry's wife was Patti. She didn't enjoy socializing with the other families. She was tired from "working on her tan". Gerry left MCAIR shortly after the program ended "to look for more satisfying work". I hope he found it, but I doubt it.

Bernie Conway was our displays lab engineer. He was from Gene Autry, Okla, and was, of course, a guitar player. His wife Darlene looked exactly like the movie star Natalie Wood. She had a twin sister who looked even more like Natalie Wood than she did. At one party, I was sitting on a sofa with a Natalie Wood on either side and someone said I looked pretty comfortable. After the program was in production, Bernie and Darlene elected to stay in California. He transferred to Victorville to help train Weasel aircrews, and later he transferred to the F/A-18 program at China Lake NAS.

Bob Sumpter was the Avionics Engineering Manager, ie he was the senior avionics guy. When the program started, he had been one of top group of avionics engineers working to define the F-4E avionics suite, but he was reassigned to the AWW program. Bob had the unique talent to see the big picture while being cognizant of, but not distracted by, the details. When he was younger and in the service, he had worked with atomic weapons, so he was a pretty calm person. He also had the patience to deal with the entire hierarchy of company managers, to represent us to them, and them to us, and to keep a good working relationship with our customers and suppliers. I don't think we would have survived without him.

DESERT STORM

Throughout the 1980s, the USAF trained Weasel flight crews at Victorville (California) AFB. During Desert Storm, ie the Iraqi war, the system was so well regarded that no fighter or fighter-bomber missions (except for the F-117 Stealth Fighter) would be flown without Weasel support. The Weasel aircraft flew more missions than any other aircraft. MCAIR got a letter after the war from one of the pilots expressing their appreciation for all the work MCAIR had put into the system.

FUTURE DEVELOPMENTS

Defense suppression is the first item on the agenda in any air campaign. The Germans tried and failed to suppress the RAF as a prerequisite to invading England in WWII. The USAF successfully suppressed the German fighters and AAA during the Battle of Berlin in that war. It is also clear that the first item in the air campaigns in Iraq and again in Yugoslavia was defense suppression. The US Navy never believed in the Weasel approach, preferring to launch ARMs into the target area, so no USN aircraft included Weasel equipment. They also believed that it was possible to jam the defense network's radars, either from beyond SAM range ("standoff jamming") or with aircraft accompanying the strike force ("escort jamming"). As the Navy's jammer aircraft, the EA-6B, wears out, it will be interesting to see how they handle the mission, since launching expensive ARMs is not very cost-effective in the modern times of cost-constrained warfare.

The USAF tried to get more modern Weasel equipment designed and installed in F-16s and F-15s. They preferred F-15s because of the two engines, but in the absence of a currently credible threat (a shooting war against a major SAM deployment), they couldn't muster the necessary resources. Recently, the USAF has evolved the Weasel approach to unmanned aircraft (small drones) that aren't as politically sensitive as live aircrews for this very dangerous mission. We'll have to see how that works. But it is safe to say that the era of manned aircraft, trolling in the SAM zone, daring the bad guys to come up, is a thing of the past. But when it had to be done, we were there and we did our part.

JLL

The Basic Math of the Advanced Wild Weasel System

In mid-1968, the USAF contracted with MCAIR for the design, development and test of a new electronics system for the F-4 aircraft. The new system, which we named the Advanced Wild Weasel, was to detect, locate, and enable weapons delivery on Soviet surface-to-air radars. The USAF had three earlier aircraft equipment sets deployed for that purpose, but all were deficient in various degrees and they wanted a system that would actually do the job. In fairness, the earlier systems had been designed, developed, installed and deployed in the panic of aircraft losses to Soviet surface-to-air missile systems in Viet Nam in 1964. The USAF had not anticipated the deployment of Russian air-defense systems (See Note 1.) in Viet Nam. The name "Wild Weasel" was derived by the USAF as a US version of the British WW II radar-hunting aircraft that the RAF called "Ferrets". So this new system was to be "Advanced". Our system was to prove indispensable in Gulf War I some years later, as desired.

So, here we are in summer of '68, meeting with the USAF (mostly majors or less in rank – all former aircrew of earlier unsatisfactory systems). They provided 31 tactical and technical requirements, but some of the technical requirements were conflicting, some impossible, and in a couple of instances, just silly. This paper is a summary of the basic analysis we used to set the system design requirements and to define the eventual mechanization. There are no numbers here because many of them were, and may still be for all I know, classified. But the equations are not.

The purpose of the system is to detect and locate Soviet ground-based radar systems and then to classify them as to basic purpose, ie surface-to-air missile (SAM) control, anti-aircraft artillery (AAA), search, or airborne intercept (AI), etc radars. SAM fire-control radars are to be classified as to type, SA-2, SA-3, etc. Suitable cockpit displays are to be included. So it is immediately clear that we also need antennas, receivers, displays, an analog-to-digital converter, a computer for control and processing, and the software to run it, and interfaces with ownship navigation, weapons control, and display systems.

The first problem is to specify the spectrum of radar frequencies to be covered by the new system. Considering the frequency coverage of previous systems and adding in our company's assessment of foreign radars, we can establish a preliminary specification. Comparing this estimate with the USAF foreign technology database sets the required coverage. This process is repeated periodically to insure that any new foreign radar system developed during our development time is appropriately accommodated.

The next problem is to specify enough receiver and antenna sensitivity in our system to detect the radiated signal from the target radars at desired ranges. This is further complicated by the need to detect the enemy radar when it is looking elsewhere, ie when we can detect only the target radar's sidelobes or backlobes. At ownship aircraft, we will have available from each target radar a signal of S, where S= Pt * Gt / 4pi * R^2. Pt is the radar's transmitter power. Gt is the radar's effective transmitted antenna gain, ie mainlobe gain minus side/backlobe losses. (See Note 2 for a discussion of "gain.") So we need the Pt* Gt, ie the effective radiated power ERP, of the target radars and a specification of the range at which they are to be detected. (See Note 3.) We can get the ERP from the customer's foreign technology experts and we can specify the detection range considering the effective range of the associated weapons plus some safety margin. Thus we can solve for S at the various frequencies we need to cover.

But there is no point in specifying an impossible parameter, so before moving on we need to check if system sensitivity S is actually achievable. We remember that S is the system sensitivity, so we need to check the contribution of the receiving antennas as well as that of the receiver itself. We need to cover 360 degrees of azimuth, ie all around our aircraft. So to minimize the number of antennas (and the associated complexity of installation on our aircraft) we want wide beamwidth antennas. And because gain G = 4 pi / beamwidth squared, we see that our antennas are not going to provide much help in detecting the incoming signals. Also, to minimize the number of antennas all around the aircraft, we want wide frequency coverage, ie we want wideband and wide beamwidth antennas.

The circumference of such antennas is approximately equal to its longest needed operational wavelength, so knowing that wavelength equals the speed of light (constant) divided by the lowest frequency we want to detect, we can define the antenna configuration. Such antennas typically operate over a wide, but not infinite, frequency bandwidth, so we must also consider adding more antennas and associated receiver hardware to cover the complete frequency spectrum of the target radars, thereby adding yet more cost and weight!

We would like our receivers to cover as much of the radar frequency spectrum as possible as quickly as possible, so we need to look at receiver frequency bandwidth, which is a measure of how much of the appropriate frequency spectrum can be checked at any one time. The detection criterion, ie the ratio of received signal S (which we have already calculated) to system noise N, has been thoroughly studied and is readily available in radar textbooks. We can select our S/N criterion by considering the probability of detecting a signal in the presence of noise and also considering the probability of a false alarm.

The noise component N = k* T* B* L, where k is Boltzmann's constant, T is the system noise temperature, which can also be taken as a constant, B is the receiver bandwidth, and L sums the system receiving losses, which can be readily estimated from current radar technology. So the issue is bandwidth B. Too much receiver bandwidth increases the noise, thereby decreasing sensitivity, ie a low noise component N wants a narrow bandwidth B, so we take B to be the reciprocal of the narrowest target radar pulse width, itself set by that radar's range resolution requirement. We conclude we need a narrow bandwidth system that must then be tunable throughout the target radars' frequency spectrum. We effectively trade target detection time for detection range.

The accuracy of location of the target radar is another driving requirement, since the tactical requirement was to locate the target radar accurately enough for potential visual identification even through a triple canopy jungle or to enable weapon delivery even if the radar site is not seen. The simplest way to estimate range to the target is to compare its detected signal strength with some standard metric, and recognizing that radar power is a function of the range squared (R^2), estimate R. Remembering that we need to detect the side/backlobes of the target radar's radiation pattern, this is too inaccurate for target location calculations so it is necessary to consider angular mechanizations. If we can calculate the elevation angle from ownship to the target radar, we can use ownship altitude to estimate range, assuming the ground between the aircraft and the target is fairly level. Unfortunately, that approach (even after correcting ownship's pressure altitude measurement for local deviations in atmospheric pressure) is still not accurate enough.

So we move to azimuth triangulation, ie measuring a line through the target, then flying a little distance, taking another line through the target, and solving the resulting triangle for range. The error of the target's location calculation can be readily calculated from the errors associated with the measurement of the angles. The distance flown between measurements is readily available from ownship navigation systems. The eventual system mechanization used the altitude and el-angle ranging technique to set a quick first range estimate that was then refined by az-angle calculations.

Angular accuracy is thus the driving criterion. We can estimate the angle to the target by comparing the amplitudes of the radar's transmitted signal as received in two adjacent antennas. This is called "amplitude comparison interferometry" and it isn't accurate enough, either. So we turn to "phase comparison" interferometry. We measure the difference in phase of the received signal through adjacent antennas and thus estimate the arrival angle of the target radar's signal. The operative equation is sine theta = phase difference times signal wavelength divided by 2pi* d, where theta is the difference between the angle of arrival of the signal and the angle normal to the plane of the antennas and d is the distance between the antennas.

We can get better angular accuracy by moving the antennas apart, but there is only so much space available on a fighter aircraft. And if the antennas are separated by more than half the signal wavelength, we get ambiguous solutions, ie there are several solutions to the phase-to-angle equation. We can resolve the ambiguity by adding another antenna between the two longer baseline antennas, ie by creating a "long baseline" and a "short baseline" antenna set. By comparing the results of both baseline antennas, we can solve for the arrival angle of a plane to the interferometer. (It would be simpler to use the amplitude difference between the shorter baseline pair to resolve the long baseline phase ambiguity, but unfortunately that technique is not accurate enough because of inconsistencies in antenna amplitude patterns, so we elect to use phase differences on both baselines, ie a "phase/phase" system. Because we need a line, ie not just a plane, to the target, we add an orthogonal set of antennas. So the issue left to control is the error in measuring the phase difference between the two signals, which is itself set by how accurately we match the phase performance of the aircraft cabling between the antennas and the receivers.

This cable phase matching, ie ensuring that each cable contributes essentially the same phase change, is so critical that it had to be demonstrated by the equipment supplier early in the design part of the program. The USAF design review team sent in early in the program to assess the likelihood of program success was skeptical that the necessary performance could be maintained in flight. (They were wrong.)

Having determined that we want a frequency-scanning receiver of bandwidth B, we need to determine how long it would take to scan through the required spectrum, pausing (dwell time) at each frequency width B for at least one pulse repetition interval, PRI, of the signal likely to be detected at that frequency. So we sum the dwell times to determine how much of the frequency spectrum we can cover in how much time, ie spectrum width divided by bandwidth equals number of dwell times and that number times the average time per dwell equals total minimum scan time. We won't always detect each pulse because pulses from different radars could probabilistically arrive at ownship at the same time, so it becomes necessary to adjust the dwell times and revisit individual frequencies at different time intervals according to a priority scheme. For example, we don't want to spend so much time looking for search radars that we miss a weapons control radar preparing to shoot. Fortunately, we can take advantage of the fact that as radar technologies developed after about 1935, operating frequencies tended to cluster into specific frequency bands, ie HF, VHF, UHF, L, S, C, X, Ku, Ka, etc. (See Note 4.)

Target prioritization is needed to support receiver frequency scanning and to enable a useful display for the aircrew. The system would detect many radars when flying in the tactical environment and displaying all of them would unnecessarily complicate the aircrew's task. The radar horizon in miles equals 1.23 times the square root of the altitude in feet or equals the square root of twice the altitude in feet, depending on your preferred source, which means that an industrial area or transportation chokepoint protected by significant missile and flak weapons will include many radars (think Kammhuber Line in WW II.) And any radar mainlobes pointed at ownship will be detected all the way to the radar horizon. So we need to estimate how many pulses, ie how many radars, might be detected per unit of time. This involves investigating the USAF database of radars, whether friendly, enemy, or neutral, that are assumed to be detectable by our system at a given altitude when operating at the more stressing locations in potential operational areas. Additionally, it is necessary to establish a map location of each detected signal, since it may be of interest even if it stops radiating (or radiates into a dummy load) for self protection purposes. (See Note 5.) More radars are detected than can be handled by our processing system, so a prioritization scheme enables lower interest radars to be considered in an organized way and even systematically eliminated.

We can establish a prioritization mechanism based on the degree of threat a particular detected system represents to ownship, based primarily on the range of the associated weapons and how quickly they could be brought to bear, as in the case of AI radars. But to do that, we also need to classify each detected system. Classifying each radar signal as to function, ie SAM, AAA, AI, etc is a matter of measuring its frequency, pulse repetition interval (PRI), pulse width, and scan modulation and comparing the results with a look-up table derived from an analysis of the USAF foreign technology database. Fortunately, different kinds of radars have become increasingly specialized and unique in that search radars tend to be lower frequency, with longer PRIs, and to scan through their assigned airspace in a regular periodic manner. Targeting radars tend to have higher frequencies, shorter PRIs, and to scan across their airspace quickly, and AI radars are moving. It is therefore possible to define an acceptable classification table and a prioritization mechanism.

The size of the display of the target situation was necessarily limited by the size of the aircraft's rear cockpit main instrument panel. Testing by pilots in the company flight simulator soon determined the maximum number of prioritized targets to be displayed, including a mechanism to limit that number at the discretion of the operator.

Our aircraft typically carries one or both of two kinds of weapons for lethal defense suppression. The favored weapon is an anti-radiation missile (ARM) that uses its own antenna and receiver system to steer itself to the target radar's antenna. It can be launched at a substantial distance from the target, which is obviously safer for the aircrew. If ARMs are not available or the target radar stops radiating, the alternative is the ballistic bomb. The F-4 aircraft has an excellent set of bombing modes, the accuracies for which have been thoroughly studied by MCAIR. Most of the associated errors, including system timing uncertainties have been fundamentally accommodated. These bombing modes are readily available if the pilot can see the target, but the desire here is to locate the target radar accurately enough with the system to enable successful weapons delivery even if the target is not visually available. So we use our target map to calculate range to the target and to show the pilot where to aim. Errors associated with this technique add some inaccuracy to the standard bombing modes, but in practice, the resulting accuracy as measured by the radius from the target to the circumference of a circle containing half the bomb hits, ie the circular error probable (CEP), is quite satisfactory.

It would be pleasant and satisfying to report that we got our requirements right, but in the event, we found two significant and related problems that had to be corrected, one during flight test, and the other before full production of the system. The first concerns multiple targets and the second the functional capacity of the computer.

During flight test, we found that we routinely detected and displayed two of the same kind of radar targets when only one radar was operating. We had thought that the antenna array patterns of adjacent-sector antennas would have sufficient null between them such that a signal would enter the system through only one array at a time and that coverage of the pattern null could be accommodated as ownship maneuvered. But it did not work that way, so the system could present two targets, separated by the typical error of the location mechanism. It proved to be much easier to correct this situation in the processing software than by redesigning the antennas. So the computer processing capacity began to grow.

The AWW system was the first fighter aircraft system to use a digital computer in the modern sense of software control and data manipulation. Earlier systems had used analog computers for many flight and weapons system functions. We determined to use software for as many functions as possible to take advantage of the flexibility such a design promised. It is easier to change software than hardware – not cheap or easy, but cheaper and easier than hardware. So we identified and sized each function, estimated its bit/byte size, and thus estimated the program size. As we worked through the multiple target problem, as well as other issues, we learned that the computer memory capacity was too small, and the computer had to grow, which it did. Twice. The phrase "software grows to fill available space" is true.

In addition to these technical missteps, we made some management errors. In estimating the time duration of the flight test program, we combined the number of specific functions to be tested with the various target systems to be worked, then factored in the number of specific test point frequencies, angles, ranges, flight altitudes, airspeeds, attitudes, etc. We then estimated how many test points could be accomplished per flight. We then estimated the number of test flights expected per week and added a safety or contingency factor to get a flight test duration estimate. Unfortunately, many of these parameters necessarily had to be only estimates since we had never tested such a complex electronics system in flight, and they tended to be underestimated. We also did not anticipate the problems involved with detecting and displaying multiple targets (see above).

An added schedule killer turned out to be the USAF insistence on our trying to improve the performance of some specific system functions, using software changes, beyond what the associated hardware could support and beyond what the performance specifications stipulated. We also underestimated the time required to get flight test data back to St. Louis, to locate (in test time) the data points to be investigated, and then to design and test the required software changes. In the event, we were not able to approach the traditional flight test management metric of flights per week.

We got some numbers wrong – but not many. This is certified by a letter from a pilot who said that in Gulf War I, no attack mission could be flown unless the Weasel was first in and last out.

Notes

1. Most Soviet radars of that era were conventional pulsed systems, ie not pulse Doppler. They tended to be standard long PRI designs using simple reflector or array antennas. So the usual simplifying conventions were appropriate, ie
Detection range is proportional to antenna diameter.
Unambiguous detection range is proportional to pulse repetition interval PRI.
Antenna azimuth or elevation beamwidth is proportional to antenna diameter.
Range resolution is proportional to transmitted pulse width.
Accurate angle tracking requires very high az and el scan rates or maybe even conical scan.
Tracking radars tend to operate at higher frequencies, ie at shorter wavelengths.

2. Antenna gain is a measure of the compression of the radiation of a sphere centered at the radiator into a beam. It is calculated as G = 4pi / theta squared, where theta is the width of the compressed beam in either the azimuth or the elevation plane. A radar antenna beam is analogous to that of a flashlight. The narrower the beam, the "brighter" it is, or the farther it sees or can be seen.

3. Radar technology math is done using decibels, ie dB. A decibel is one tenth of a bel (Alex. G. Bell), which comes from sound pressure technology and is the exponent by which the number 10 must be expanded to give the number being specified. It is one tenth of the logarithm of the number in question, eg the log of 100 is 2 (because 10^2 = 100) so ten times that, ie 20, is the dB equivalent of 100. 30 dB is 1000 times, etc. Adding the logs of numbers is the same as multiplying the numbers, so we use dB because the arithmetic is much simpler. Multiplication would be further complicated by the wide range of numbers involved. An antenna gain of 100 times is 20 dB. Line loss of a factor of –50 times is -17 dB. Receiver sensitivity of 0.001 Watts is –30 dB relative to one Watt. So the combination is readily 20-17-30 = -27 dBW. One regularly hears the phrase "3 dB down", which because 0.3 is the log of 2 means that 3 dB down represents the original value divided by 2, or the half-power point in a circuit.

4. The radar frequency band designations were created during WW II for security purposes. Modern electronic warfare usage prefers to divide the spectrum into A, B, C, D, etc nomenclature. This serves no particular purpose to the practitioner other than some confusion, which may or may not be desirable.

5. From the earliest days (in WW II) of counter-radar operations, the standard technique for radar self-protection is to stop radiating. This of course represents a temporary victory for the hunter. But a radar transmitter produces considerable heating (necessitating a cooling system). This equipment takes some time to attain operating temperature, so a better self-protection technique is to keep the radar operating, but to radiate its generated energy into an absorptive (dummy) load. Full

operation can then be resumed quickly. Depending on the efficiency of the dummy load and upon the sensitivity of our receiver system, we may be able to detect even the dummy load radiations. If not, it is convenient for our system to create a geographic map of the targets if they become temporarily undetectable. A more sophisticated self-protection technique for the radar is to have a separate alternate antenna conveniently available if the Weasel seems to have located the primary antenna. But transmitter-to-antenna cabling losses, as well as test and maintenance issues, tend to force the radar operators to keep the separate antennas fairly close to the transmitter. The Weasel aircraft can approach the victim radar indirectly, using terrain features for concealment.

JLLooper