Case C: The Total Weapon System Concept

C(1): The Battle of Britain Command and Control System

Introduction

At the outset of the Battle of Britain in 1940, the German plan was to employ their superior air power in what has since become classic fashion: they would obtain air superiority over southeast England, as a prelude to land invasion — their Operation Sea Lion that, in the event, never took place. Gaining air superiority meant eradicating the RAF's fighter aircraft; this was to be achieved by a campaign of bombing airfields and factories, to destroy fighter aircraft on the ground, and to halt production. So, southeast England faced a bombing campaign aimed, not at centers of population, but principally at fighter bases and factories.

The UK was ill prepared and lacked experience with which to take on the might of the German war machine. Germany was particularly proud of its air force, the Luftwaffe, which had gained invaluable experience and become battle hardened during the Spanish Civil War. Germany had built up its airforce strength, so that at the start of the Battle of Britain the Luftwaffe could muster over 1200 aircraft on the French side of the English Channel, compared with only some 300 fighters, mostly Hurricanes with a sprinkling of Spitfires, facing them in southeast England.

The Royal Air Force (RAF) did have one ace up its sleeve: Chain Home radar. The UK had not invented radar — that seems to have been a German invention, but they used it at sea, where it seemed to work best. The German command apparently had no idea that the UK had a number of Chain Home radar stations dotted along the southeast coast, facing out to sea. The UK also found that their radar worked best over the sea, with land creating so many echoes that it was difficult to pick out and track an aircraft flying over land: however, detection and ranging over water were possible. . . .

Air Marshall Hugh ('Stuffy') Dowding had introduced the RAF's radar system in 1936 as part of his efforts to improve the UK's air defense system. Dowding also encouraged the formation of an extensive network of visual sighting posts of the Royal Observer Corps (ROC). Civilian

Systems Engineering: A 21st Century Systems Methodology Derek K. Hitchins
© 2007 John Wiley & Sons, Ltd. ISBN: 978-0-470-0585-5

volunteers manned the posts, which are still to be seen at high points all around southern England to this day. Each post was equipped with an optical sighting device mounted on a plinth, using which an operator could estimate the direction (track), range, height and speed of an aircraft, and could also identify aircraft type, how many aircraft in a raid, etc.

The manned ROC posts connected by a network of telephone lines (an intranet?) to a filter station, which received air intruder reports, often from several ROC sites at the same time or in quick succession. From these reports, filter stations were able to determine raid details, if there was more than one raid, and — if there were — where the raids were headed. The filter station then passed this raid information on to sector operation centers, which had large plotting maps of southeast England, on which operators placed and moved markers to indicate the progress of raids.

Air defense fighters were located at a number of airfields: aircraft could be ‘scrambled’ in response to an incoming raid, and given a vector to fly once airborne, so that they would intercept the raid. There were some 20 squadrons all told, and these could be rotated so that some squadrons were on duty, while others took time off for both pilots and ground crew to recover and repair. So, another factor in the equation describing the hoped-for balance between the opposing forces, was the tactics to be employed by overall commander Dowding, and by Air Vice Marshall Keith Parks, Air Officer Commanding 11 Group, covering southeast England — a New Zealander by birth, and, as it turned out, key to victory in the UK.

Figure C.1 shows the set up as it was in the summer of 1940, omitting the numerous ROC posts. The area was divided into sectors, each with its sector operation centers located at an RAF

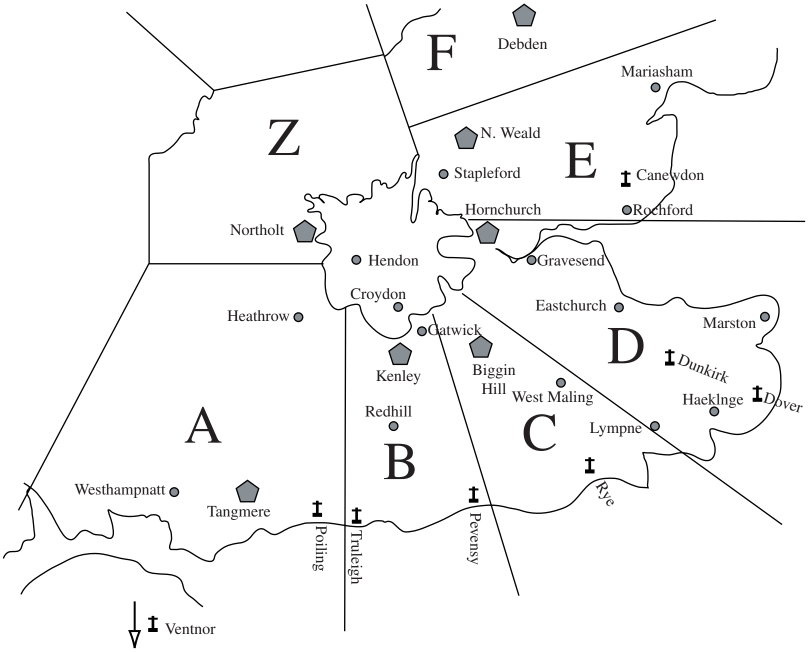

Figure C.1 Notional plotting map of southeast England, No 11 Group, RAF, 1940. Sectors are labeled A to Z. Pentagons are sector operation centers, one per sector. Other airfields are shown as small circles. Chain Home radars are shown as towers distributed along the coastline, looking over the sea towards France (bottom right, not shown).

station. Sectors also had airfields with fighters based at them, although the same fighters could be dispersed, usually forward towards the coast; this had the dual effect of making them harder to locate, so reducing the risk of damage due to enemy strafing while on the ground, and also making them closer to the coast; so, quicker to intercept incoming raids.

Interacting Systems

There were many interacting systems on both sides during the conflict. Figure C.2 shows a simplified view, with RAF No. 11 Group Fighters centre left, and the Luftwaffe with its fighters and bombers to the right. Clearly, they were the protagonists, but there was more to the system than that, else the Luftwaffe, with its 4:1 aircraft supremacy would have simply overrun the RAF.

Both sides expected to lose aircraft and crews during the campaign. For Germany, the way to replace losses was either to get new aircraft and crews direct from Germany, which was relatively slow, or to redeploy aircraft based in Norway, Denmark and the Low Countries, where their role was largely to harry allied shipping.

For the RAF, there were three possible resupply routes: direct from factories, from stocks of fighter aircraft that had been constructed and were held in reserve, and from other groups which were not so hard pressed by Luftwaffe attacks at this time. Obtaining replacement crews could be more problematic, however — replacement fighters from reserves or straight from factories did not come with pilots — these had to be recruited and trained, both of which were time consuming.

One important factor came to the aid of the RAF. The Luftwaffe adopted a tactic of assembling their bombers and their fighter escorts into airborne formations over France before crossing the Channel. While Chain Home radar was primitive by today's standards, it could pick out the large echoes created by these formations, which gave No. 11 Group early warning of an impending raid — also shown in Figure C.2.

The possession of radar turned out to be a 'force multiplier:' without the radar, and therefore without early warning, No. 11 Group would have been obliged to mount continuous daylight air patrols along the southern coast, ready to engage incoming targets; this would have reduced significantly the operational availability of RAF aircraft to meet unexpected Luftwaffe raids. . . .

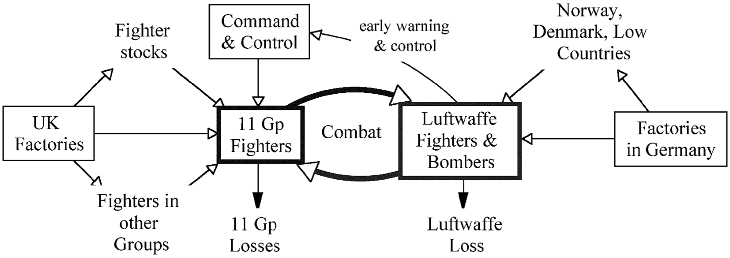

Figure C.2 Battle of Britain, 1940 — the systems interaction diagram (SID). Not shown, but a significant factor, was the weather — a highly variable, typically English, summer.

Working up the System — Operational Systems Engineering

The elements of the UK air defense system were in place before war broke out; the system worked, but it did not work well. In particular, the system took too long to respond to detected raids, so that enemy aircraft would have been able to deliver their bombs on airfield targets with relative impunity. In current parlance, the 'time around the loop' was too long, where the loop was the time to detect, report, decide on response, scramble defensive fighters, vector them on to target, engage and destroy/deter the enemy before bomb-drop. So, while the technological elements of the system were in place and performing, the overall system was nonetheless inadequate.

The Chain Home radar could detect Luftwaffe formations assembling over France some 20 minutes before the consequent raid arrived over the southern coast of England. The requirement for No. 11 Group was to meet those incoming German bombers, with their fighter escorts, as they crossed the English coastline. The key time-response factor, then, was 20 minutes.

Dowding set about improving air defense; his objective was to reduce the time taken by the people in the loop to perform their various tasks. ROC visual sightings, essential for tracking raids over land, were correlated in the filter station: one enemy raid might be reported by a number of ROC sighting posts, with reports being unsynchronized and offering estimates only of numbers, types, bearing, altitude, etc. Consequently, the filter station could be faced with an overload of information, from which to sort out the most likely situation. And suppose there were two raids, or one raid splitting into two sections. . . .

Dowding introduced reporting procedures, to reduce the time taken to describe a raid or enemy position. It is from this period that expressions such as 'Angels 15' come, meaning 'at an altitude of 15 000 feet. By introducing simple verbal codes for reporting various situations, Dowding not only accelerated procedures, but also reduced the propensity for misunderstanding. With practice, the mean time taken to filter and report a raid was reduced to 4 minutes — still a relatively long time, but workable.

The sector operations centers were similarly scrutinized and procedures improved. One of the problems with the map, on which counters were positioned to represent incoming raids and outgoing fighters, was that raid reports might be few and far between, so that the position of a counter may not be updated for several minutes. This could lead to misunderstandings about the 'state of play.'

Dowding and his team improved the Operations Room Plotting Clock, Figure C.3, adding different colors, red, blue and yellow, around the rim of the face, at five-minute intervals. So, the first five minutes after the hour were in red sector, the next five minutes were in yellow sector, and the third five minutes were in blue sector. The colors then repeated in the same order around the clock. When a new plot was received in the sector operations center, its marker was put on to the plotting table (the map) with a color label on it corresponding to the color that the clock minute hand was pointing to when the plot was received. So, a plot received at 3 minutes past the hour would be marked red on the table, another received at 8 minutes past the hour would be yellow, and so on. In this way, operations staff could tell at a glance whether a plot was current or stale. (This was an early example of the management of latency.)

At airfields, aircraft were positioned on dispersals, with their crews sitting beside them, listening to landline messages broadcast over loudspeakers direct from the corresponding sector operations center. Aircraft could be plugged into ground power trolleys, which were anchored so that aircraft could start up and taxi, automatically disconnecting from the trolley cable as the aircraft moved forward.

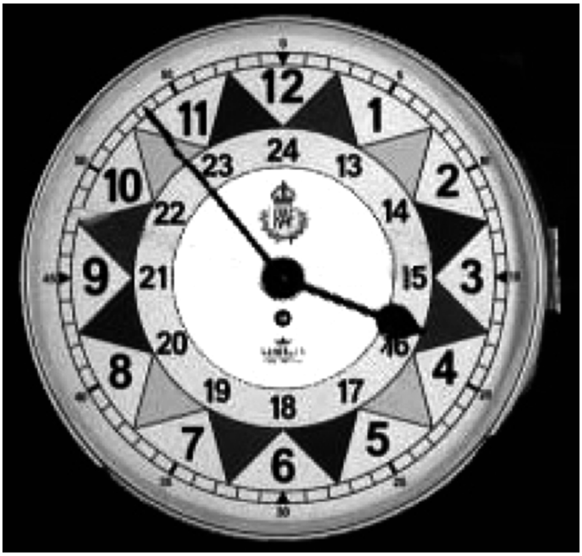

Figure C.3 Operations Room Plotting Clock Face, showing the minute hand in a ‘yellow’ sector

Once airborne, sector operations staff could vector fighters towards an incoming raid using the simple radios fitted into RAF fighter aircraft: (German attacking aircraft had no such communications with their bases). RAF fighters had to climb to altitude as quickly as possible, so that they could engage enemy aircraft with at least height parity, if not height advantage. On average, climbing to operational height took some 13 minutes. Together with the 4 minutes for target filtering time, this left some 3 minutes out of the 20-minute early warning period to transit to the coast, see, and engage the enemy. It was still tight for time. . . .

So, Dowding and his team gradually worked up the overall command and control system, with its sensors, processors and effectors (ROC, Chain Home radar, filter stations, sector operations centers and fighters) by taking advantage of the adaptability of the human elements within the system. Over a period of months, they gradually ‘raised their game’ until the air defense system as a whole could match the Luftwaffe raiders — just. Although the term was not in use at the time, they had developed a ‘total weapon system concept,’ in which all the parts of the of the air defense system operated organismically, as a unified whole, to achieve a singular, specific purpose — the neutralization of Luftwaffe bomber aircraft attempting to bomb fighter airfields and factories.

Let Battle Commence

The Battle of Britain was brief, but bloody. Graph C.1 shows the relative aircraft losses for Luftwaffe and the RAF on a day-to-day basis.

While it can be seen that the daily pattern of losses is not dissimilar, at least in terms of peaks and troughs, it is also evident that the Luftwaffe were losing many more aircraft than the RAF. And, since many of the Luftwaffe aircraft were bombers with several crewmembers, the loss of

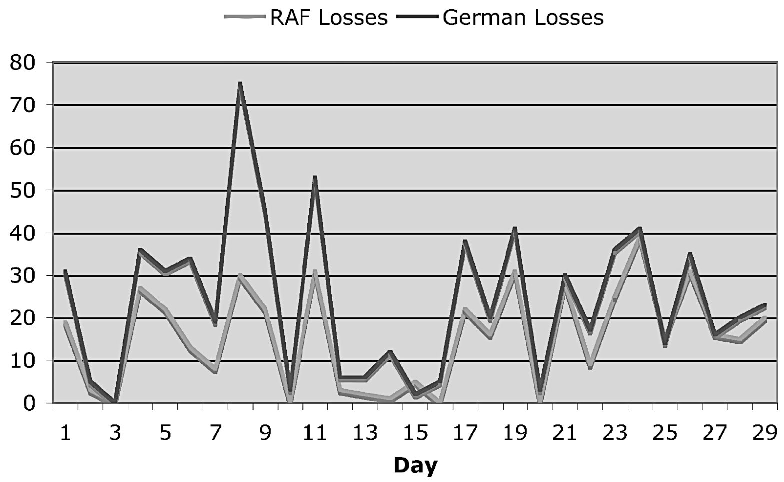

Graph C.1 Battle of Britain: relative aircraft losses, 8 August – 5 September 1940. Losses on each side were broadly in proportion over the 29-day period, with German losses consistently higher than RAF losses. . . .

men was even greater on the German side; this was a cause for great concern to the Luftwaffe, particularly since their aircrews were drawn from the cream of German society.

Moreover, when a Luftwaffe aircraft came down over English soil, any crew who survived would be interned; if the aircraft came down over the channel, the chances were that the aircrew would not survive at all. For the RAF, the situation was different: if a fighter went down, the pilot might bale out, and could be back at his base, ready for operations, in a matter of hours. Not everyone managed to bale out of course. . . , but the advantage, at least in numerical terms was to the RAF. And RAF fighters were not permitted to chase fleeing German aircraft over the Channel, to minimize the risk of losing both aircraft and pilot over the sea.

Graph C.1 is also remarkable for the degree of variability in day-to-day losses. While some of this might have been the result of weather, it is unlikely to have been sufficient to cause such variability. The English summer weather was, and is, notorious in its unpredictability, and the summer of 1940 lived up to its reputation. Moreover, German bombers operated effectively only in clear weather, and could be prevented from operating by mist, low cloud, etc. Poor weather could be a two-edged sword, however, not only preventing raids, but also allowing both sides to repair damage, train up replacement crews, and generally prepare to operate in larger numbers when the weather cleared.

The curious degree of variability in day-to-day downed aircraft statistics is highlighted in Graph C.2. The phase-plane chart is produced from the same figures that gave Chart C.1, but in this case plots the changes in combined aircraft losses from day to day, so that the line on the chart follows that change from day 1 to day 31, in order. Had the statistics varied, say, in a simple wave fashion, then the phase-plane chart would have looked like a circle, or oval. Had the statistics varied chaotically, then the phase plane chart might have looked like the butterfly — see Lepidoptera Lorenzii? on page 35. As it is, Chart C.2 looks nothing like an oval, and can be regarded as vaguely butterfly-shaped only by a major stretch of the imagination.

Did the statistics indicate some fractal characteristic, perhaps? (See Fractals on page 39.) Graph C.3 was formed by sampling the combined loss statistics at different intervals: every day,

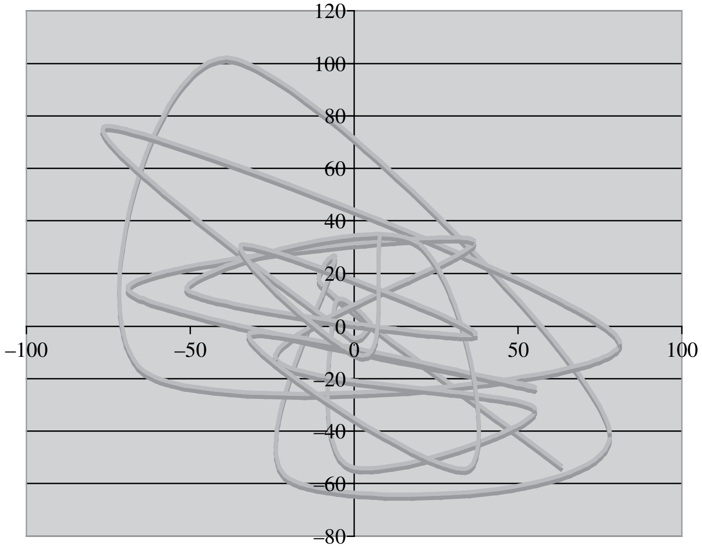

Graph C.2 Luftwaffe vs RAF losses — phase-plane chart, taken from the same historical statistics that generated Graph C.1.

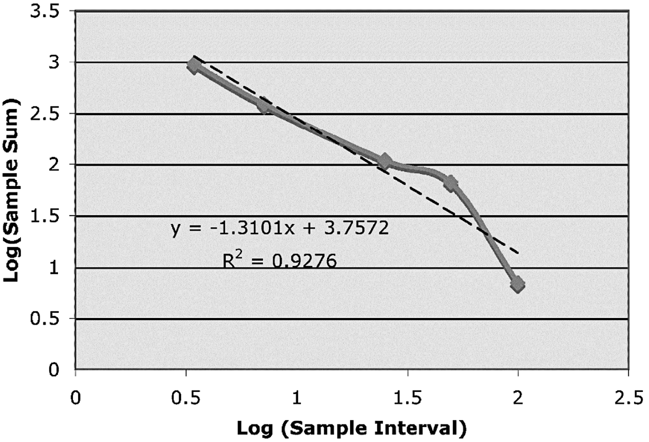

Graph C.3 Analysis of the combined RAF and Luftwaffe loss statistics of Graph C.1, showing linear trend line. See text.

every two days, every four days, and so on, and summing the each set of samples to produce the log–log graph. (This approach highlights the 'bumpiness' of a graph.) Although not conclusive, the results suggest that the pattern of losses may, indeed have been 'weakly chaotic,' or fractal, conforming reasonably well ($R^2 = 0.93$, where 1.0 would be precise conformance) to a power law distribution with a 'fractal' index of 1.3.

AVM Keith Parks' tactics

One possible explanation for the variable statistics might be the actions of AVM Keith Parks. He adopted what, in retrospect, may be seen as classic defense strategy for a besieged force. In particular, he took great care to minimize losses. Not only did Parks insist that No. 11 Group fighters did not pursue enemy aircraft over the Channel, but he also prevented them from rising to the bait when flights of Luftwaffe Me109 fighters came over without their bombers in tow, evidently spoiling for a fight. Parks saw his job as protecting RAF fighter bases and factories from Luftwaffe bombers, and not taking on marauders who, by virtue of their greater resources could better afford to lose fighters in air-to-air combat.

Parks also rotated the squadrons so that the brunt of the defensive work was shared out; a particular squadron might be working all out for several days, then have a rest for a week or more, while crews and planes recovered. In this way, he always had fresh squadrons to hand, should an unexpected peak in activity arise. The number of squadrons on 'immediate operations' might vary, too, with several squadrons sometimes taking on a relatively small raid, while at other times, perhaps only one or two squadrons would face up to a large raid. Overall, Parks' tactics, although questioned at the time, turned out to be brilliant, making the best economic use of his meager resources so that he would be able to meet an impending major assault, while at the same time presenting an entirely unpredictable face to the enemy, who, as it later turned out, were confused about how large, or small, No. 11 Group's resources were.

Battle of Britain Simulation

The simulation described below was developed for a Granada TV program for the History Channel as part of their 'Battlefield Detectives' Series. The Battle of Britain was particularly brief because the Germans, having failed to suppress the RAF in the few short days they had expected, switched instead to a quite different campaign of bombing centers of population such as London.

The TV program set out to investigate three issues:

1. How long could Fighter Command have lasted if the Luftwaffe had continued to attack the RAF bases and radars, instead of switching to bombing London and other cities?
2. Can the Battle of Britain be reasonably classified as a 'win' for the UK, or not?
3. Was the strategy employed by Dowding and Parks, of conserving RAF fighter aircraft resources in anticipation of prolonged activities better, or worse, than the Big Wing concept — which would have contributed most in the long run. . . ?

This last question was occasioned by critics of Dowding and Parks, notably AVM Trafford Leigh-Mallory of No. 12 Group, to the north of London, supported by Douglas Bader, the famous fighter ace who flew throughout the war with 'tin legs,' from losing both legs in a prewar flying accident. Both critics would have preferred to amass their fighter resources into a so-called Big Wing, and to take on the enemy in a decisive major air battle.

The phase-plane chart of performance, Graph C.2 is, in effect, a behavioral signature of a system: a characteristic indicator of emergent behavior. So, what was going on during the intense Battle of Britain to cause this unusual signature? One way to explore the problem space is to simulate the Battle of Britain, and to vary parameters within the simulation until its signature matches that of the real world. Observing the parameter changes that led to this match may shed light on the then situation. . . .

The simulation is presented as a so-called learning laboratory: the user can experiment with many different facets of the Battle to see what the likely outcome might have been. . . . The simulation has three principal parts: representations of:

- the Luftwaffe, a mixed force of fighters and bombers, based in Northen France;
- Dowding's Command and Control System, including Chain Home radars for early warning, ROC filter stations, Sector Operations Centers, and the communications infrastructure that coupled them into a single system
- Number 11 Group, under AVM Parks, with its sectors in southeast England, and its sector operations stations at Tangmere, Kenley, Biggin Hill, Hornchurch, North Weald, Debden and Northolt.

The various parts (modules) of the simulation were built and tested separately, then brought together after the fashion of Figure C.2, and the whole was tested before being used as a learning laboratory.

No.11 Group had some 20 fighter squadrons, mostly Hurricanes with some Spitfires, divided between the sectors and under the control of the sector stations. The 'engine room' of the simulation represents the squadrons of No. 11 Group, accounting for the numbers of aircraft and pilots, the casualty rates, hospitalization and recovery rates, aircraft damage and repair times, aircraft damaged on the ground by Luftwaffe action, etc. Luftwaffe operations are represented in similar fashion, but without any attacks on Luftwaffe bases.

The top level of the multi-layer simulation presents an Operations Map Room in the form of a Group Operations Plotting Table — a map of No. 11 Group. and southeast England. Surrounding the table are controls to alert the fighter squadrons, plus estimates of raid size — From this position, users can play 'Park's Game,' choosing — as AVM Keith Parks may have done — which squadrons, and how many squadrons, to send up in response to a raid. There is also an intermediate layer Control Panel, where 'players' can change the rules — starting conditions, kill probabilities, etc. They may also change the simulated weather, too, making flying more or less likely on a four-hourly basis.

The simulation runs for 1440 simulated minutes, or one day. In real time, simulating one day takes about 3 minutes. Using a sensitivity mode, it is possible to set the simulation on 'auto,' at which time it will run uninterrupted for some 31 equivalent days. The results of each run are accumulated on tables in the Control Center, and then compared with the 'real world' charts of Graphs C.1 and C.2.

Running the BoB simulation

The BoB simulation turned out to be very sensitive: successive runs over a simulated 31-day period could give quite different results, even when variables, such as weather, probability of kill (Pk), Luftwaffe raid size, etc., had not been changed, run to run. This was occasioned, at least in part, by the inclusion of random elements within the simulation. For example, setting a Pk value did not determine that a particular number of aircraft would be downed every time, since the simulation allowed results of engagements to be distributed about mean levels set, in part, by the Pk. Similarly, although the weather could be set to different levels of suitability for flying, the

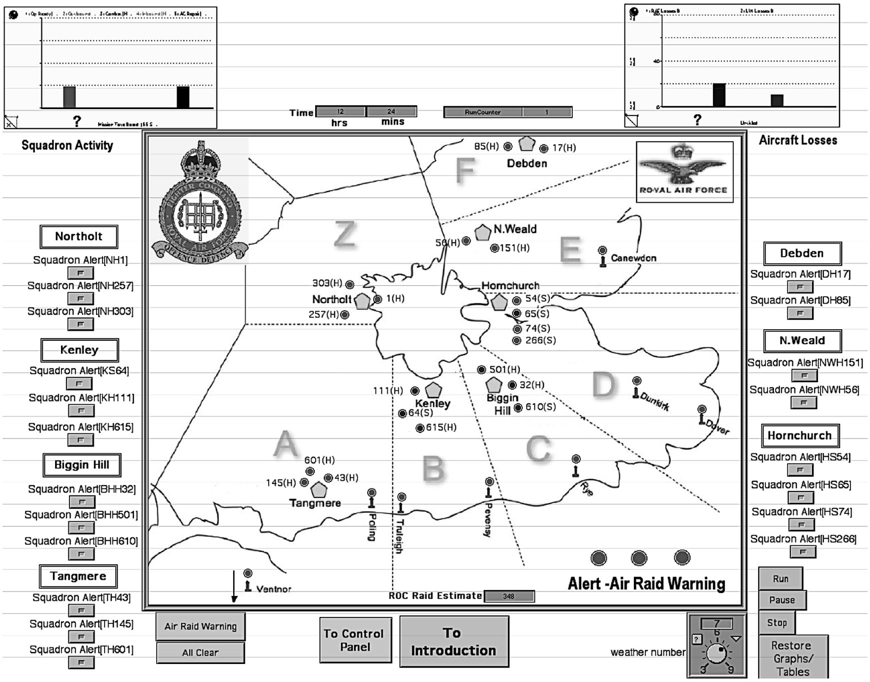

Figure C.4 Battle of Britain — Operations Room simulation. The operations plotting table is shown in the center. Sector Operations Stations are shown left and right, with their respective squadrons. Initials after the squadron number indicate the Sector and the aircraft type: NH is Northolt, Hurricane; KS is Kenley, Spitfire. Illuminated switches show which squadrons are active, and they are also shown as flashing circles on the map. Other squadrons are shown as non-flashing circles. Charts at top left and right are simulated tote boards, and there is a runtime counter at top dead center, showing current time paused at 12 hours 24 minutes. Bottom left are Air Raid Warning switches, to sound the raid alert and the all clear.

periods when a morning or an afternoon were not suitable for flying varied randomly, and so also differed run-to-run.

Sensitivity arose because the variability in results of one engagement determined the number of aircraft returning to their respective bases, to different levels of repair work, to the time available for that repair work before successive missions were due, to the numbers of crew needing medical treatment, and so to the number of aircraft available for the next day's missions. . . which then affected the outcome of the ensuing engagements, and so on. Such variations accumulated over the 31-day simulated period, so that results were quite variable. Moreover, a large number of 31-day runs, all with the same parameter settings, produced a series of results that did not fall into any obvious pattern. . . .

Running the simulation under virtually any set of preconditions failed to produce results that were comparable with the real-world statistics of Graphs C.1 and C.2. Weather could have a serious effect on the combat outcome: in the extreme, of course, poor weather for the whole 31 days would have virtually stopped all operations, since the Luftwaffe could bomb only in clear weather. Perfectly clear weather throughout, on the other hand resulted in a major flurry of losses on both sides, followed by periods of wound licking and repair, before starting again. The pattern of losses tended, in this situation, to become more cyclic, and again unlike the real-world statistics.

Some success in emulating the real world came by combining the effects of moderately good, though unpredictable, weather with the impact of highly variable tactics on the part of the air defense commander, Keith Parks. The simulation made it possible to try out different tactical schemes to see which might produce real-world-like results.

The first possible tactic was to rotate the various squadrons quite quickly, so that each squadron was on active operations for only one or two days, before being 'rested:' at the same time, only a few squadrons were active — three or four at the most. This enhanced the relative kill rate, i.e., the number of Luftwaffe to the number of RAF aircraft downed. The reason was to do with operational availability: each squadron had about 16 fighters; after several days' rest, all of these, including their pilots, would be available for operations. However, once a squadron had been on continuous operations for several days, some of the aircraft would be damaged awaiting repair, lost, etc., and their pilots likewise. By rotating squadrons rapidly, squadrons starting operations would be to full strength.

The second tactic has been mentioned in the first — to deploy only a very few squadrons at a time. The mathematics is simple: if only, say, 48 aircraft (three squadrons out of 20) are deployed, then — at the absolute worst — only 48 could be lost or damaged. Parks did not want to destroy every attacker in one go; instead, his plan was to break up their formations and pick of stragglers — much as lions do when hunting zebra or wildebeest on the Masai Mara.

The third tactic, and that which brought the simulation nearest to real world statistics, was to be quite unpredictable. To be predictable in the situation facing Parks would be to match the size of the defensive force to the size of the incoming raid. So, if a force of, say, 500 aircraft was approaching — as reported by the ROC — then logic might suggest that Parks should put up twice the number of fighters as when 250 aircraft were approaching. If, in the simulation, that simple rationale is NOT followed, then the simulation results start to look more like real world.

In a way, the sense of this seemingly reverse logic is seen in the previous paragraph — if more aircraft are put up against a bigger force, then more aircraft are put at risk. But, that only makes any sense if Parks believed that the Luftwaffe had little chance of hitting their intended airfield targets, and if he believed that even a small defensive force in the air would put the Luftwaffe off their stride. On the other hand, Parks was determined to preserve as much of his air force as he could in anticipation of the ensuing major German offensive, Operation Sea Lion, evidence of which was to be seen on the French coast with a build–up of barges and other shipping that might be used for invasion purposes. Whether wise, or cautious, or both, Parks' tactics won the day. . . and Operation Sea Lion never happened.

Graph C.4 shows the simulation results for a 31-day run, with moderate weather and the employment of seemingly irrational tactics, which we may notionally attribute to Parks. This graph should be compared with Graph C.1; clearly, they are not the same, but there are degrees of similarity to be seen in the day-to-day variability, in the consistency with which Luftwaffe losses are higher than RAF losses, etc.

Similarity can be seen, too, at the phase-plane graph, C.3, drawn from the same simulation as Graph C.4: the 'signature' in the phase plane is vaguely similar to that of Graph C.2. As with people, no person's signature is the same twice — it would be suspicious if it were: similarly, one should not expect the simulation signature to be identical to the real world.

Graph C.6 provides a closer, more detailed look at the pattern of aircraft losses shown in Graphs C.3 and C.4. Graph C.6 was formed, as was Graph C.3, by taking samples of the combined aircraft losses at different intervals, and summing the samples in each case — this enables the 'bumpiness' of the Graph C.4 to be assessed. The results, shown in C.6, are encouraging. Conformance of

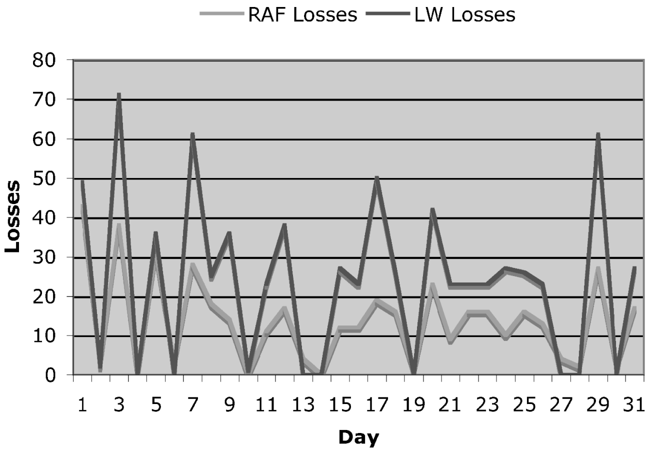

Graph C.4 BoB simulated combat over 31-day period. Weather: moderate. Continual change of defensive tactics.

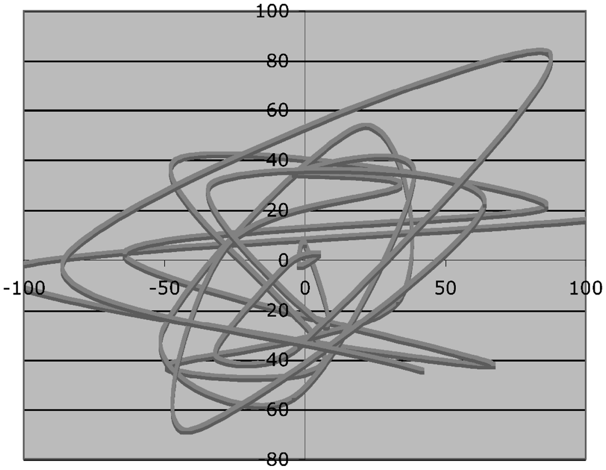

Graph C.5 Phase-plane signature form Chart C.3. Compares with Graph C.2.

the results to a straight line (the R^2 value) is convincing, indicating that the results from the simulation are fractal, while the slope of the line, the fractal index is -1.0312 — different from the real world fractal index of -1.3 — see Graph C.3. Judging by Graph C.6, perhaps Parks really was unpredictable — or perhaps the tactics of the opposing Luftwaffe commander Field Marshall Kesselring, were contributing to the unusual pattern, too. . . .

None of which explains just why the phase-plane diagram — the signature — is the shape that it is — irregular, nonperiodic, nonchaotic. . . . The general form of this signature has occurred before in the book — see Dynamic simulation of phenomena on page 68. In that instance, a phase-plane chart was drawn up to show the variability in population on the River Nile some 5000 years

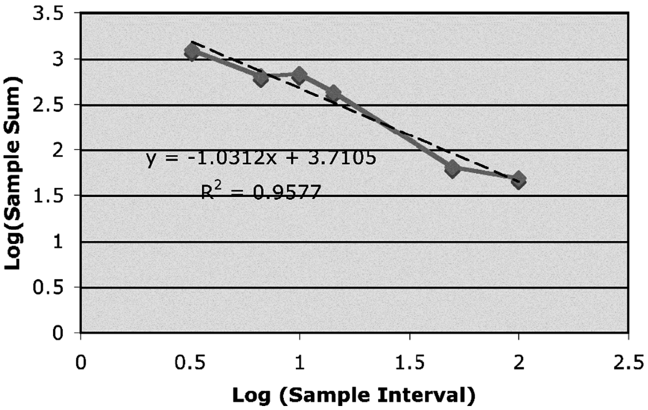

Graph C.6 Analysis of the combined RAF and Luftwaffe loss statistics of Graph C.3, showing good correlation with linear trend line, suggestive of 'weak chaos' and/or self-organized criticality.

ago, caused by variation in the annual Nile inundation triggering feast and famine such that the population rose during feast, but died off during famine.

The variation in Nile population was an instance of self-organized criticality. Graph C.6 is also indicative of self-organized criticality, or 'weak chaos.' And, curious though such a comparison may seem, observing a broadly similar general phase-plane pattern suggests that the interacting systems in the Battle of Britain may also have been in a state of self-organized criticality. In working up ('orchestrating') the various parts of the air defense system into an organismic, cooperative, coordinated unified force, Dowding brought it up to critical performance level — at which level it was, indeed, just able to see off the Luftwaffe.

Which leaves open the other questions posed at the outset, for which the simulation was constructed, and continued running of which suggests the following answers:

I. How long could Fighter Command have lasted if the Luftwaffe had continued to attack the RAF bases and radars, instead of switching to bombing London and other cities?

   - Difficult to be precise, but simulating the next few months, allowing for variations in weather, and assuming that neither side changed its tactics, it seems likely that the Luftwaffe would have effectively run short of aircraft and crews after some six months. . . .

II. Can the Battle of Britain be reasonably classified as a 'win' for the UK, or not?

   - In recent years, it has become popular amongst would-be historical analysts to pose such questions. Were a besieged castle to bloody and see-off its attackers, such that they gave up and went to do something different, then there would be little doubt in classifying the attacker as defeated and the defenders as winners. In this case, southeast England was a fortress with a coastline and a channel instead of rampart walls — but the rationale is the same. Of course, it was a victory — certainly, the Germans believed they had been defeated!

III. Was the strategy employed by Dowding and Parks, of conserving RAF fighter aircraft resources in anticipation of prolonged activities better, or worse, than the Big Wing concept — which would have contributed most in the long run. . . ?

- The problem that Dowding and Parks faced included the prospect of a long-drawn-out battle, some of it at least on English soil as the English retreated inland in face of Operation Sea Lion with its landing forces. Even without that prospect to keep in mind, they were grossly outnumbered, and had no real idea of the size of the Luftwaffe reserves.
- Had Dowding and Parks succumbed to the Big Wing concept, then they would have faced two risks: the Big Wing would have been vulnerable on the ground as it mustered fighter aircraft from various stations to assemble in one location; putting the Big Wing up against superior odds was a greater risk, for the obvious reason that the Luftwaffe could afford to trade aircraft with the RAF one-to-one, or even two-to-one, and still turn out the winners, in the style of ‘last man standing.’
- As the fight progressed, and as the Luftwaffe lost more aircraft than the RAF, there might have come a time when an all-out attack by the RAF would have made sense. During the time of the Battle of Britain — August–September 1940 — it would have been, and would have been seen as, a naïve mistake.

Despite that last bullet, and despite having won one of the most famous victories in history, within six months both Dowding and Parks were ‘moved sideways’ (effectively, demoted): winning, it seems, was not enough. . . .

One other point of issue concerns the various systems within the air defense system overall. Would-be experts delight in pointing out the singular ‘key to victory.’ For some it was the Chain Home ground radar. For others, it was the Royal Observer Corps. For many it was the Spitfire — in spite of many more Hurricanes being involved in the actual battle.

Their enthusiasm for their various causes is admirable, but misplaced. The Battle of Britain was won by the total system, with all the parts operating, cooperating and coordinating as a close-coupled, unified whole. To those who think that is not correct, consider which element of the overall system could be *removed*, leaving the rest to operate. The answer is *none*!

Hence, it is the whole system, or nothing, and it is unreasonable to pick out any one element as the key. That is not to decry the incredible bravery of the fighter pilots, nor the technological innovation of Chain Home, nor the tireless efforts of the ROC volunteers, day after day, night after night, nor the ground crews working all hours, in all weathers and while being attacked by enemy aircraft; and we must never forget the sector operations personnel, the radar personnel, the post office engineers keeping the telephone lines working. . . .

But, it was the system as a whole, and — although the term ‘systems engineering’ was yet to be coined, it seems entirely reasonable to describe Dowding’s development of the air defense systems overall performance, in retrospect, as operational systems engineering of the first order.

As Winston Churchill put it: ‘Never in the field of human conflict has so much been owed by so many to so few.’

C(2) The Lightning — Realizing the Total Weapon Systems Concept

Introduction

The end of World War II ushered in the Cold War: the UK became concerned about potential air attacks from the East, and of the potential Soviet use of standoff weapons, which would allow attacking aircraft to launch their attack weapon when still some distance offshore; the weapon would carry on to its destination, while the launch aircraft would return to base unscathed. Soviet aircraft might carry more than one weapon, so might launch at more than one target, and the standoff weapons were smaller and much faster than a manned bomber, so would be much more difficult to intercept. Or so intelligence believed . . .

Faced with such a daunting threat, the UK developed a postwar air defense system tuned towards the particular threat. New, advanced ground radars were established up and down the east coast of the UK, looking primarily eastwards, out to sea. Operations centers were built underground, in fortified bunkers. The whole of this so-called Air Defence Ground Environment (ADGE) was tied together with a communications infrastructure, including digital data links. Airborne early warning facilities were developed, so that their radars could see beyond the horizon, and they were equipped with data links to report what they had seen to the operations center.

Although using newly developed technology, and facing a different enemy, the architecture of this new ADGE was much the same as Dowding's innovative prewar air defense system — indeed, most of the air defense airbases were the same RAF stations Dowding had used: sector operations centers replaced sector operation stations, data links replaced the hard-wired, voice-operated intranet of the Battle of Britain air defense system, but — yes, the architecture was the same. In place of Hurricanes and Spitfires, there came a succession of jet fighters: Hunters, Javelins and the Lightning.

The Lightning

During the late 1940s and 1950s, the UK was very active in aircraft development, with a host of new aircraft types, engines, wing plan forms, etc. The Fairy Delta Mk2, for instance, as the name implies, was a delta-wing aircraft that led eventually to the development of Concorde. Large bombers were developed, capable of carrying nuclear weapons: the so-called V-bombers, Valliant, Victor and Vulcan, all of which saw service in the RAF.

Among the developments was an experimental aircraft, the English Electric P1. It was a 'notched delta' design, giving it a distinctive shape and an awesome rate of climb to high altitude. With rocket-assisted takeoff (RAT), it is said to have reached 90 000 feet from a standing start in 3 minutes. . . .

The P1's potential made it an ideal candidate for inclusion as the intercept element of the UK air defense system — except, that is, for such trifles as a lack of any sensors and weapons. The aircraft had, in effect, to be redesigned as a practical military interceptor: it needed

- a good radar, to see the threat Soviet aircraft;
- phenomenal speed, to fly out in time to meet the threat aircraft before it could launch its standoff weapons;
- even greater speed to overtake any standoff weapons that had already been launched;
- new kinds of air-to-air weapons to combat the new threat, to reach ahead of the interceptor, to catch the standoff weapon, etc.;
- superior climb and turn capability in case the enemy chose to attack at high altitude

Designers set to create a total weapon system aircraft, initially the P1A, then the English Electric Lightning Mark 1A. This aircraft was to have one purpose — to intercept Soviet aircraft before they could launch their weapons. It was to be part of an air defense system also dedicated to the same purpose, so that the whole air defense system, including the Lighting, was to be a coordinated, cooperative, organismic unified whole.

The objectives were clear: the means of making the P1 into an interceptor much less so. To achieve its outstanding performance, it was, in effect, 'two kerosene burners place one above the other, with wings stuck out either side.' In other words, it was an experimental aircraft, not only without any of the equipments needed to become an interceptor, but also with very few places in which to put them.

The P1, or the Lightning as it would become, needed radar — but where to fit it? The ingenious solution was to create a radar 'bullet' that would fit into the shape of an intake center-body. If this could be sensibly achieved, it would have a hopefully limited effect on engine performance and on aircraft drag. The solution was the Ferranti AI23B, an advanced, 3-cm pulse radar with a sophisticated four-beam antenna system that enabled the radar to lock to a target in azimuth and elevation. The radar was fitted with elegant analog computing facilities that enabled the Lightning to fly intercept paths that were suited particularly to the new missile that was also being developed — the de Havilland Firestreak. See Figure C.5, which shows both radar and missile *in situ*.

Firestreak, based on another research program, Blue Jay, had a solid-rocket motor, infrared direction sensing and fuzing, fragmenting warhead and its own navigation and autopilot. It made use, unusually, of magnetic amplifiers as opposed to the more conventional electronic ones, to overcome problems from heat and vibration. Firestreak had a limitation, however — its sensor was able to pick up only infrared signature emanating from the hot turbine disk of an enemy aircraft, and so was able to attack the enemy aircraft only from behind. This meant that the Lightning would have to fly out towards the target aircraft, and perform a looping maneuver to roll out some five miles, or so, behind the target before firing the missile.

The looping maneuver, which pilots came to call a 'butcher's hook' after its shape, took precious time, and so was a disadvantage. To combat this, the AI23B computer was programmed to fly the ideal 'butcher's hook' relative to the target, and it presented suitable signals both to the autopilot

Figure C.5 A Lighting Mk 1A of No 111(F) Squadron undergoing night OTR (operational turn round.) Note the radome in the nose of the aircraft, the Firestreak Missile with protective covers, center, and the refueling probe at the top of the roundel (photograph by author).

and on the pilot's attack sight (PAS), a smart head-up display, so that the pilot could steer manually if so desired. In either event, the objective was to minimize the time taken to reach a firing solution. . . .

Nonetheless, Firestreak was less than ideal, so an improved version, Red Top, was in development, which would permit head-on engagements. Red Top would not be available with the first Lightnings into RAF service, however.

With no place on the P1 to 'embed' missiles, they had to be externally mounted. This was achieved by fitting interchangeable ventral weapons packs, which could carry Firestreak, Red Top, rockets, or cannon if required.

So, the designers of the soon-to-be Lightning found ways to fit sensors and weapons, but there were still serious problems. The flight duration was brief, to say the least; the aircraft could not carry much fuel in its stainless-steel wings: a detachable ventral tank was fitted, extending flight duration significantly, although durations was still short, especially using full engine power and reheat. Fitting two external missiles and the radar center-body increased drag, which made matters worse, so the hoped for speed of Mach 2.0 was hard to realize. And, having only two missiles did not permit one Lighting to take on many opponents — even supposing it had the time.

Optimizing the Design

The design problem was formidable. The objective was to reach the target when it was as far from the UK coast as possible, and certainly before standoff weapon launch distance. The pilot had to be able to see the target at a distance, day or night. The detection range of the radar was dependent, *inter alia*, on scanner dish diameter: the greater the dish diameter, the greater the radar's range. However, increasing the diameter would increase the cross section of the center body, or bullet, which would increase drag and fuel consumption, and reduce speed.

Similarly, mounting Firestreak missiles externally increased drag and fuel consumption. Increasing the cross-section of the Firestreak might have increased its target detection range, and a larger motor might have increased its firing range, but at a cost to overall aircraft performance.

Some of the tradeoffs are presented in the simplistic model of Figure C.6. The model employs Newton's Law relating force, mass and acceleration. In this case, the force is the difference between (engine) thrust and aircraft drag. The resulting acceleration is integrated to give velocity, and integrated again to give (aircraft) distance traveled. Velocity is used to calculate the drag caused by both radar and missile, such that the larger their cross-sectional areas, the greater the drag and the less the performance.

So, the designers were faced with an optimizing problem, for which the solution lay in balancing the various factors to enable the shortest time to kill. The simple model of Figure C.5, with typical results at Graph C.7, shows that increasing the radar cross-sectional area beyond a necessary minimum gave increasingly poor returns — although a fatter radar might see further, the aircraft would take so much longer to reach the target area, that there would be an overall loss in effective operational range. Besides, there was little point in being able to see much further than the range at which the air-to-air missile could be launched.

The answer had to be to keep the radar and missile cross-sections as small as practicable to minimize drag and maximize aerodynamic performance, while at the same time directing the Lightning on to its target using ground radar, airborne early warning (AEW) or both: these other radars would act as eyes for the Lightning until it was close enough for its interceptor radar to see

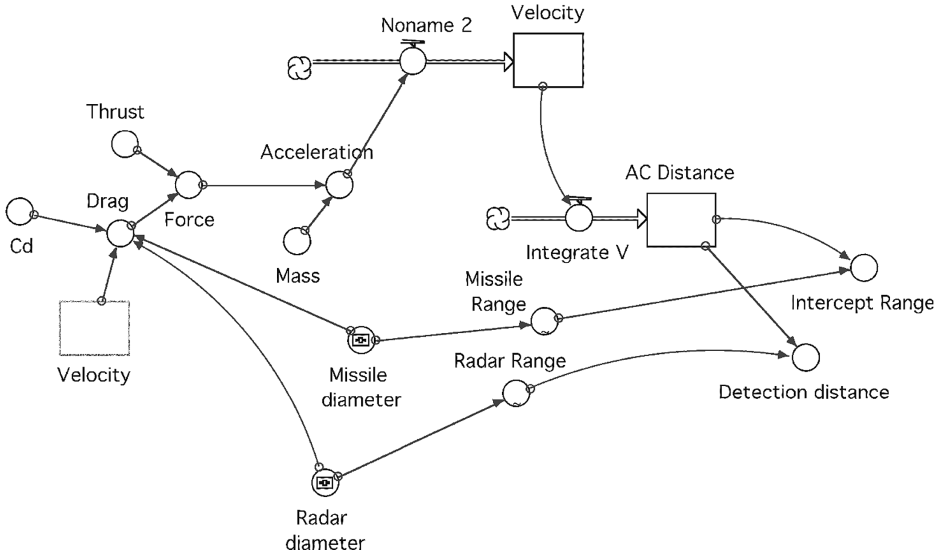

Figure C.6 Simple model to show interaction between radar and missile cross-sectional area and Lightning aircraft operational performance.

for itself. Evidently, it was not sensible to optimize the Lightning on its own — if it was to part of an overall system, then that overall system had to be optimized, and the Lightning would have to be configured as an integral part of the whole weapon system.

Keeping the drag low, and the performance high were essential to deal with another aspect of the supposed threat: Soviet aircraft might try to enter UK airspace at high altitude, above the ceiling of ground-based missile systems. To combat this, the Lightning, with its superb climb rate, would be equipped with a zoom-climb capability. Zoom-climb is a technique for reaching high altitude

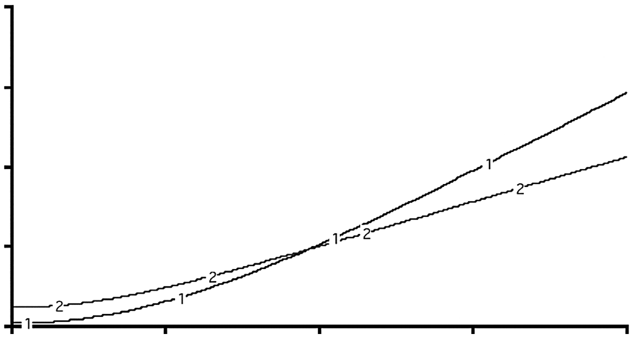

Graph C.7 Lightning potential intercept range (y-axis) against time (x-axis), from the model of Figure C.5. Line 1: small radar and missile cross-section; Line 2: larger cross-section. Beyond the crossover, smaller radar and missile cross-section would give better overall intercept range. . . .

by accelerating at medium altitudes to very high speeds and then climbing steeply, turning the kinetic energy of forward velocity into potential energy of height. So-called programmed zoom-climb profiles were built into the radar and autopilot, such that the pilot could be guided to zoom climb, topping out in a suitable engagement position. This could be a precise operation, since the Lightning would find it difficult to maneuver at very high altitudes, and would in all probability get only one shot at the target.

The Jamming Problem

If the Lightning interceptor was to depend on ground and AEW radars for directions toward a distant target, then what would happen if the enemy decided to jam the radars, or the radio communications from ground radars to the Lightning, or both? Intelligence indicated that the Soviets were investing in airborne jamming capabilities, so their attack aircraft might carry jammers, they might be accompanied by jamming aircraft, or they might use standoff jammers, i.e., jamming aircraft (or ships) that kept back from the line of interception.

Jamming would have put a major spoke in the defensive wheel — what to do about it? The ground radars (which were to become the Type 84 and Type 85 radars of the Linesman–Mediator system, which combined UK air defense on the one hand — Linesman — with the management of civilian air traffic on the other — Mediator) were fitted with comprehensive anti-jamming facilities and features, including passive detection arrangements that permitted triangulation of multiple jamming aircraft, so they, at least, might be located and, potentially, intercepted. But, that left the problem of directing the Lightning aircraft on to the target aircraft — the enemy might be using communications jamming.

Digital Data Links to the Rescue

Digital data links were in their infancy at the time, and they seemed to hold the potential key to solving the problem of communication jamming. Supposing that target data could be sent automatically from the ground radars, the Types 84 and 85 radars, digitized and transmitted to airborne Lightning interceptors, together with flight control instructions for the interceptor pilots to follow, such that they would roll out just behind the target aircraft in the ideal firing solution.

The system designers went one further: suppose all of this data being sent from the ground radars to the airborne interceptor was not just displayed to the pilot: suppose it was used to automatically steer the interceptor and to operate the engine throttles, such that the pilot — having engaged the system — would have had nothing to do other than monitor the situation. Further, using digital data links, if this arrangement could work for one interceptor aircraft, then it would work for many simultaneously. . . .

A research program was set up to construct and trial a ground to air digital data communication system. To be effective, it would have to be high powered, and the transmitters would have to be on the east coast — preferably near the ground radars, also on the coast.

The research program comprised a number of elements:

- a processing system at the ground radar to gather data on airborne targets and airborne Lightning interceptors, to calculate intercept paths, speeds, climb/descend profiles, etc.;
- a processor to fit the data into packets, with headers, addresses and parity, to be sent via data links;

- a multiplexed ground-to-ground data link to pass the data packets over telephone lines to a remote radio transmitter;
- a ground-to-air radio data link transmitter, converting the data stream to modulated radio signals;
- an airborne radio receiver to receive and decode the radio signals, reconvert them into message packets, check parity, and send messages to their respective destinations;
- an airborne interface to use the messages to display target data (range, bearing, altitude, track, velocity, etc.), to control aircraft interceptor direction, altitude, speed, etc., via the aircraft autopilot and autothrottle.

The technology of the day was such that the data link could be unidirectional only: the radio uplink employed the full bandwidth of the airborne UHF transmitter–receiver, a modified version of the conventional set, incorporating a data demodulator. There was insufficient room in the Lightning to fit a second transmitter–receiver: when the pilot selected 'data-link,' the radio received and decoded the data stream only: there could be no voice communication with the ground. To get around this, a special, small voice recorder unit was fitted behind the cockpit with a number of prerecorded voice messages on it; each of these messages could be activated from the ground radar, to say 'Check speed,' 'Check height,' and — most importantly — 'Return to R/T,' indicating that the pilot should revert to voice communications.

Unexpectedly, this recorded voice communication discomforted pilot during early trials. For the sake of clarity, the voice used had been that of a professional woman radio announcer, and pilots complained that the voice messages alarmed them, sounding like there was another person with them in a single-seat cockpit; moreover, the voice was that of a woman! They proposed that the voice messages should be re-recorded using a man, and that artificial 'static' should be included, to render the messages similar to those received over the radio. In retrospect, such trivial complaints were symptomatic of a deeper malaise, and should have forewarned developers of the system. . . .

Developments of prototype equipments took some time, and the Lightning Mk1A had been in service several years before the complete data link system could be set up for trial, in the mid-1960s. A Lightning was fitted with the prototype airborne elements, including a data decoding system that was so large it was referred to as 'the dustbin;' if the trials were successful, a much smaller version would be developed. A ground-to-ground link had been established from the UK midlands, where there was a prototype Type 85 radar, to a more southerly location over several telephone lines. In this more southerly location, a 20 kW UHF transmitter had been set up, with a 20 dB gain antenna, giving an effective radiated power (ERP) of 200 kW; this may be compared with the transmitter power in common use for airborne UHF communications of 25 W.

The trial

The trial was to take place over the English Channel — not far from the coastline that had withstood the worst of the Battle of Britain. The Lightning was to take off and fly over the Channel, be picked up on the prototype Type 85 radar, was then to receive flight control signals from the Type 85, over the ground to ground and ground to air data links, and was to engage the automatic flight controls system (AFCS) for short period of ground-controlled flight. The distances between the ground radar and the ground radio transmitter, and between that transmitter and the aircraft, were representative of the distances supposedly involved in intercepting a Soviet bomber prior to release of its standoff weapon.

Everything was tested on the ground and retested. Came the fateful day, and the Lightning took off as required, transited to the English Channel. . . and later returned. The pilot emerged from the single-seat fighter looking 'unhappy.' It seems that the equipment had worked as it should, but that the pilot was extremely uncomfortable with being flown by some remote agency, and with the manner of the flying.

Unsubstantiated rumor at the time suggested that a scientist or technician at the Type 85 radar site had seen a switch supposedly in the wrong position, and had 'corrected' it, causing the Lighting to perform an instantaneous 180° roll. Be that as it may, and the story may well be apocryphal, the trial proceeded no further, and plans to go to the next stage of technological development were abandoned, apparently without further regard to the Soviet jamming threat.

Conclusions

In spite of the data link fiasco, for that is in effect what it became, the Lightning went on to perform as the principal intercept element in the UK air defense system of the time, but using only voice control — no digital data. Of its type, it was a good example: it had speed, agility and an excellent rate of climb, all features shared by good interceptors. It lacked stamina, and, with only two missiles, it was short of weapon power. A gun had to be fitted retrospectively, when it was realized that it was not a good idea to shoot down every intruder — someone might simply have lost their way: missiles could only kill; a gun was needed, to warn intruders without killing.

In an extension of the iconic BoB image of aircrews relaxing outside their squadron dispersal huts listening to the loudspeaker, Lightnings could be parked on an auxiliary service platform (ASP) at the end of the runway, with their aircraft intercommunications system connected to the sector operations center by 'telebriefing,' through landlines. Pilots could be strapped in, with ground power connected to the aircraft, along with cooled air to feed the pilots' air ventilated suits; one word from the sector operations center was sufficient to scramble instantly, with automatic disconnection of ground power and cooling air and telebrief, followed by takeoff, literally within seconds.

The Lighting had been fitted with a number of automated devices that, in the end, were rarely, if ever, used. Why not? The interceptor radar, Ferranti's AI23B had a computer which would fly the pilot in the ideal 'butcher's hook' to perform a rear hemisphere engagement. Pilots agreed that it worked, but would never use it. When asked why, they had no good reason, other than they wanted to fly the aircraft themselves, in their own way, and they believed that they could do better than any computer — which, incidentally, was not the case. They had a point, though — they were reliable, whereas the AI23B, along with most avionics of the day, was not!

The problem with the data link was, essentially, the same: it was not that the technology did not work, because it did; it was more that the pilots wanted to fly the aircraft themselves, in their own way. They really did not like being sidelined and made to feel useless while some 'faceless boffin' and his equipment flew the aircraft remotely and impersonally from many miles away.

Looking back at Dowding and the way he progressively improved the BoB air defense system, it can be seen that, while he 'tuned' the system by using the adaptability of the human operators, both on the ground and in the air, he did not attempt to change the technology; he trained and organized the people to make better use of their technology. In the Lightning era, it was no longer aviators tuning the adaptive, people-element of the sociotechnical system, but scientists and technicians designing it literally from the ground up; their ambition seems to have been to cut the human out of the loop altogether. It was clear that attempts by scientists and engineers to supplant the human element in sociotechnical systems did not always go down well with human operators. . . .