SAFETY ANALYSIS FOR ADVANCED SEPARATION CONCEPTS

John W. Andrews and Jerry D. Welch

M.I.T. Lincoln Laboratory, Lexington, MA 02420-9185

Heinz Erzberger

NASA Ames Research Center, Moffett Field, CA 90435

Abstract

Aviation planners have called for increasing the capacity of the air transportation system by factors of two or three over the next 20 years. The inherent spatial capacity of en route airspace appears able to accommodate such traffic densities. But controller workload presents a formidable obstacle to achieving such goals. New approaches to providing separation assurance are being investigated to overcome workload limitations and allow airspace capacity to be fully utilized. One approach is to employ computer automation as the basis for separation-assurance task. This would permit traffic densities that exceed the level at which human cognition and decision-making can assure separation. One of the challenges that must be faced involves the ability of such highly automated systems to maintain safety in the presence of inevitable subsystem faults, including the complete failure of the supporting computer system. Traffic density and flow complexity will make it impossible for human service providers to safely reinitiate manual control in the event of computer failure, so the automated system must have inherent fail-soft features.

This paper presents a preliminary analysis of the ability of a highly automated separation assurance system to tolerate general types of faults such as non-conformance and computer outages. Safety-related design features are defined using the Advanced Airspace Concept (AAC) as the base architecture. Special attention is given to the impact of a severe failure in which all computer support is terminated within a defined region. The growth and decay of risk during an outage is evaluated using fault tree methods that integrate risk over time. It is shown that when a conflict free plan covers the region of the outage, this plan can be used to safely transition aircraft to regions where service can still be provided.

Introduction

In the United States, the Joint Planning and Development Office (JPDO) has called for the capacity of the air transportation system (ATS) to increase by as much as a factor of three over the next 20 years while at the same time increasing security, safety, and efficiency [Ref. 1]. In addition to removing barriers to long-term economic growth and to the introduction of new vehicle types, this extra capacity would alleviate bottlenecks experienced today when hazardous weather closes busy routes.

The JPDO acknowledges that achieving significant capacity increases is unlikely unless new approaches are found that allow one to go beyond the limitations of the current ATS paradigm.

The first question that we address is whether the airspace is inherently capable of handling traffic at densities that are severalfold greater than experienced today. We describe the results of a simulation experiment that seeks to verify the existence of assignable four-dimensional (4D) trajectories that accommodate a several fold increase in traffic density.

The second question is how to exploit this airspace capacity. One of the transformational concepts being investigated is to allow computers to assume responsibility for the time-critical decisions that ensure safe separation between aircraft. If successful, this would relax current human workload limitations and allow fuller utilization of the inherent capacity of the airspace. A key concern is whether such computer-based separation can achieve an acceptable level of safety in view of the subsystem faults that might arise. This paper is a preliminary analysis of these safety issues. It is intended to provide an early prioritization of risk factors and to indicate where more detailed modeling or design work is advisable.

To provide a specific concept for analysis, we begin with the description of the Advanced Airspace Concept (AAC) that has been developed by the

National Aeronautics and Space Administration (NASA) [Ref. 2,3]. We then postulate additional safety-related design features that seem appropriate for such a concept.

Any such concept must address several design challenges. Some essential tasks require human judgment and reasoning, and cannot be automated. Human operators must continue to make many decisions, and the system design must compatibly merge human and computer decision-making. Consequently the goal is to automate time-critical separation assurance tasks, not to replace a human sector controller with a computer. For example, setting a proper trade-off between overall system efficiency, user preferences, weather avoidance, and system load balancing requires consultation and judgment that is not readily consigned to computer algorithms. Providing expert consultation to aircrews requires voice contact and human understanding of the questions being asked. Reconfiguring the system to handle poorly defined or unexpected events may require direction by a level of general intelligence that has never been demonstrated in computers. On the other hand, computers excel at computationally intensive tasks, such as simultaneously revising multiple 4D trajectories to provide efficient, conflict free routing. By addressing this more limited problem, computers may provide the required capacity breakthrough without having to master the full range of skills needed to manage the air traffic system.

The evolution pathway for the AAC concept will involve intermediate stages that use AAC building blocks primarily for decision support or for handling traffic subsets. However, this paper focuses on a mature AAC operating in airspace in which all traffic participates in the system. This allows us to define an end state that is both feasible and safe - the first step that the ATM research community must take to provide a basis for defining evolutionary steps.

Although we restrict this analysis to en route airspace, the techniques can be applied to AAC-type operations in terminal airspace. As a minimum, model parameters would have to be adjusted to reflect the special characteristics of terminal airspace.

Inherent Capacity of Airspace

Several studies suggest that available airspace does not constrain en route traffic throughput [Ref. 5, 6]. We examined two en route sectors that are currently operating near the controller workload limit (sectors 46 and 48 in the Cleveland en route center). We found evidence of excess airspace volume. Almost half the aircraft traversing these sectors never came closer than 20 NMI to other traffic (four times the radar separation standard). But a more relevant question is whether one can readily find sufficient conflict free 4D trajectories to accommodate traffic growth.

To answer this, we captured trajectories of existing traffic to form a baseline set of 4D trajectories that are safe by definition, and that represent the required distribution of cruise altitudes, headings, speeds, etc. To determine whether one can triple the number of such trajectories and still maintain safe separation, we "cloned" four additional 4D trajectories from each actual trajectory. We displaced the clones from the original trajectory as illustrated in Figure 1.

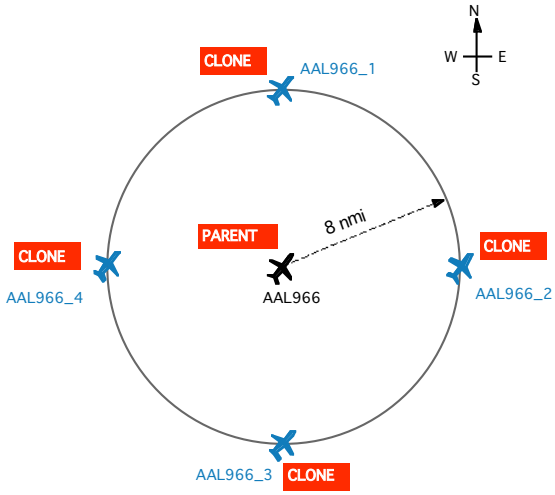

Figure 1 Cloning to produce candidate 4D trajectories.

Using a separation requirement of 8 NMI (the current standard is 5 NMI); we deleted clones that conflicted with a baseline trajectory. Next we deleted one clone from every pair of conflicting clones. In the end, 87% of the cloned trajectories remained. If an automated planner fully utilized these trajectories, the traffic density in the airspace could increase from 763 to 3421, a factor of about 4.5. In an actual implementation, the selection of candidate 4D trajectories would be more versatile than the cloning procedure, and could produce additional candidate trajectories.

These estimates of available capacity using the cloning method have recently been confirmed in simulation using a candidate automated resolution planner. The resolution planner software has been exercised with varying traffic demand in Cleveland

Center airspace. The automated resolution planner was able to perform at three times the current traffic density.

Level of Safety Goals

We seek a rational basis for safety assessment with achievable goals, well informed trade-offs, and efficient resource allocations. Accepted practice is to define a *level of safety (LOS)* that must be achieved in the face of all possible risks. From this requirement, a number of subsidiary requirements can be imposed.

The International Civil Aviation Organization (ICAO) has established Target Level of Safety (TLS) standards for air transportation [Ref. 7]. The overall mission TLS is one fatal accident per 10 million operations or 1.0E-07 per operation. If an operation averages two hours, this corresponds to an allowed rate of 5.0E-08 accidents per flight hour.

The mission TLS is further allocated to particular phases of flight. For en route airspace the allowed rate of fatal accidents per flight hour is 5.0E-09 per direction (vertical, lateral, and longitudinal). Some analysts interpret this as allowing a total accident rate of 1.5E-08 per flight hour [Ref. 8]. The Target Level of Safety (TLS) established for reduction of vertical separations (RVSM) was "2.5 equipment-related fatal accidents in a billion flight hours" or 2.5E-09/flight hour [Ref. 9].

The *baseline risk* is the probability of collision, when a potentially hazardous encounter arises, and all subsystems are operating normally, but none of them succeeds in preventing the collision. It is essential that the factors leading to baseline failure be very rare because the baseline exposure is high (the system is in the baseline state most of

the time). The baseline risk must be considerably lower than the overall target level of safety because deviations from normal operations (faults) will occur that will raise the overall risk of collision above the baseline risk. We choose a design objective of 1.0E-09 accidents per flight hour when no faults are present.

The next section enumerates the characteristics that we believe will make it possible to operate the AAC with this low baseline collision rate.

System Characteristics

Figure 2 shows the basic architecture of the Advanced Airspace Concept as described by Erzberger [Ref. 2]. Data-link equipped aircraft exchange trajectories and planning messages with the ground-based AAC infrastructure. Unequipped aircraft (those without ability to interface via data link) participate through interaction with the human service provider over a voice link. An Automated Trajectory Server (ATS) on the ground analyzes all proposed trajectories and ensures that they are conflict free for a specified period of time.

The ATS may update existing trajectories or add additional conflict free trajectories for aircraft currently within its sector. Its most critical function is to generate conflict free trajectories in response to detected conflicts within the sector. It also checks all submissions of user-requested trajectories for conflicts and for consistency with airspace and flow restrictions. When necessary, minimally modified trajectories that are free of conflicts and restriction violations are negotiated. Furthermore, the controller has tools to interact with the ATS to generate conflict free flight plan amendments for unequipped aircraft or for any equipped aircraft requiring special handling. The trajectory database of currently assigned conflict free trajectories and flight plans is consolidated into a Separation Plan that can be inspected as needed by service providers, traffic flow planners, adjacent facilities, and aircraft.

The Tactical Separation Assured Flight Environment (TSAFE) serves as a safety back-up to the primary ATS logic, providing short-term separation assurance as necessary. TSAFE activates when a predicted near-term loss of separation (predicted to occur at some time less than 2 minutes into the future) justifies immediate action in which separation assurance overrides all other considerations. Studies are in progress to define the most effective logic for triggering this transition.

In addition to these basic features, the mature AAC system analyzed in his paper is assumed to have the following characteristics when operating under nominal conditions (without any subsystem failures):

1) Airspace is divided into control regions that are more than four times larger than today's conventional sectors. These regions encounter peak traffic densities three times greater than can be accommodated in today's system.

2) A reliable digital data link allows aircraft to receive 4D trajectories from the AAC system and transfer these trajectories into the flight management system. (This capability builds upon the standards and procedures for 4D trajectory protocols now being developed by RTCA Task Force 3 [Ref. 10]).

3) Every aircraft maintains voice contact with a designated control facility at all times. (Such contact ensures prompt assistance and helps maintain safety when faults occur.)

4) Every aircraft in the controlled region remains reliably within surveillance coverage from the time of entry to the time of exit.

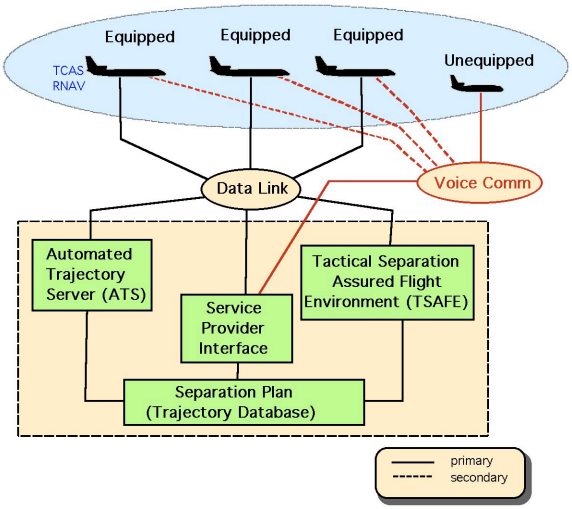

Figure 2 Basic Architecture of the Advanced Airspace Concept.

5) All approved trajectories, together with associated system configuration information, reside in a single published Separation Plan, which is a database of currently assigned conflict free 4D flight plans.

6) Trajectories in the Separation Plan are conflict free at least as far as the exit point from the AAC sector. Trajectory segments further than 20 minutes into the future are unlikely to be executed due to the evolution of the traffic situation or to reconformance. (Re-conformance is the process of *preventing* conformance errors by revising the planned trajectory to start at the current position of the aircraft when it appears to be drifting out of conformance.[Ref. 11])

However, generating such segments for possible use can play an important role in fault management. It also supports the formulation of flow management strategies.

7) In generating trajectories, multiple control objectives are simultaneously considered with proper balancing of competing objectives. Among the control objectives are aircraft separation, user preferences, severe weather avoidance, flow control constraints, fuel efficiency, and airspace loading. When subsystem failures occur, the ATS can sacrifice efficiency to maintain safety.

8) Human service providers and aircrews can request the AAC to revise trajectories. Revisions will be approved only after safety validation. This assures the completeness and integrity of the Separation Plan.

9) Human service providers are not responsible for generating detailed trajectories, nor for approving generated trajectories, nor for monitoring conformance to the Separation Plan. (This is essential for both safety and capacity benefits, and it permits a practical workload allocation.)

10) The Separation Plan is available for inspection at multiple ground sites. The ATS also transmits relevant portions of the plan to aircraft to enhance situation awareness and facilitate trajectory negotiations.

Integrity of the Separation Plan

The prime operating principal of the postulated AAC design is that *no update process is allowed to compromise the safety of the Separation Plan*. The trajectories comprising the Separation Plan provide adequate separation between aircraft for the near future, and do not result in unsafe proximity to terrain or known or predicted hazardous weather. All interactions that can alter the plan are subject to safety validation. User requests may result in a sub-optimum plan with respect to fuel burn or load balancing, but they are never allowed to create a plan that is unsafe if flown. Under these conditions, conformance ensures safety. Risks are associated almost entirely with non-conformance events.

The second important principle is that *the Separation Plan is rapidly and easily modified* to accommodate user requests or to optimize itself in response to new data. Revisions can occur for a number of reasons such as aircrew request, weather forecast updates, unexpected winds, changes in the traffic environment, or re-conformance. Multiple parties (service providers, traffic managers, and aircrews) can negotiate trajectory changes without risk of introducing hazardous incompatibilities. Because no human is required to approve changes or analyze trajectory interactions, user requests can normally be granted in seconds. When incompatible user requests are received, the system suggests acceptable alternatives.

A third important principle is that *the Separation Plan is available for inspection*. Under traditional ATM concepts, portions of the strategic

plan exist only in the minds of individual controllers and cannot be validated or fully communicated to others. Under the AAC concept, the plan is subject to scrutiny by independent safety monitors using different algorithms to avoid common coding errors. Multiple copies are maintained to facilitate transfer of responsibilities in the event of outages. Its accurate transmission to aircraft is ensured by secure digital techniques to preclude the errors in voice transmission that plague conventional control. Conformance monitoring allows early intervention when aircraft start to deviate from their planned trajectories.

Fault Analysis

Given adequate funding and motivation, there appears to be no fundamental operational barrier that would prevent a mature AAC system from operating in accordance with the proposed new architecture and airspace design. Furthermore, there exists today sufficient technology, management, and implementation expertise to provide the infrastructure characteristics needed to keep the AAC operating nominally with high reliability and availability. If the AAC is operating correctly and a potentially hazardous encounter arises, the system is designed to take the proper action to avoid a collision. It appears reasonable to postulate a design that achieves a baseline collision rate of one accident in 10^9 flight hours when no faults are present to interfere with that design.

However, faults will occur. We design the system to handle faults, and we must examine the system safety in their presence to assure that they do not lower the overall level of safety below the acceptable limit.

In our analysis, mission failures are defined as collisions. Such failures are associated with defined faults. As noted above, a fault is any event that poses a specific risk of collision not otherwise present. A fault may be associated with a particular hardware or software element, or it may result from failure of subsystems or human operators to properly execute a procedure.

We proceed by defining a set of faults that represent all significant risks to safety. Currently, each fault type is broadly defined (for example, all faults that produce non-conformance are grouped into two types.). Because faults occur rarely and the design allows fault conditions to be repaired quickly, the system remains mostly in the no-fault state. Simultaneous occurrence of two or more defined faults need not be considered because such events are much rarer than the single-fault cases and do not significantly elevate risk. Risk arises from non-conformance - which is already present with the first fault to occur.

The analysis is based on identifying the types of faults that can occur and estimating the combined risk. If the collision rate for fault k is F_k, the final collision rate that defines the level of safety is

$$F_{LOS} = \sum_k F_k \qquad (1)$$

We estimate the risk for each fault by using a fault tree analysis. A fault tree is a simplification of the real situation that helps us identify weak elements of the design and leads to strategies for addressing critical safety problems. For most faults, the tree explicitly defines the chain of events that must occur to produce a collision.

Figure 3 illustrates a generic fault tree that postulates a chain of events leading to a collision. It assumes that when a potentially hazardous encounter arises, resolution is attempted in succession by the primary AAC system, TSAFE, and the Traffic Alert Collision Avoidance System (TCAS). If TCAS fails, then there is still a chance for the conflict to be resolved with visual avoidance. Finally, there is a probability that—simply by chance—the aircraft do not collide (although they will still experience a near mid air collision (NMAC)).

Inspection of Figure 3 leads to the following basic mathematical expression for the rate of collisions from a type k fault:

$$F_k = \lambda_k \lambda_{NMAC} \tau_k \alpha_k \beta_k \gamma_k \psi_k P_{coll} \qquad (2)$$

where

λ_k = mean rate at which type k faults occur

λ_{NMAC} = mean rate at which near mid air collisions occur in absence of any separation process

τ_k = mean duration of type k faults

α_k = probability that the basic AAC separation processes fail to resolve a type k fault

β_k = probability that TSAFE fails to resolve a type k fault

γ_k = probability that TCAS fails to resolve a type k fault

ψ_k = probability that visual avoidance fails to resolve a type k fault

P_{coll} = probability of collision for an unresolved near mid air collision.

Fault Taxonomy

For this preliminary analysis, we group all possible faults into four general fault types. Future analysis can add detail by subdividing the types or, if necessary, defining new types. The four general fault types are:

Type 1: Nominal conditions. Aircraft are in conformance with the Separation Plan and none of the other defined faults are present. This is the most common condition for an encounter. Even though conditions are very favorable for separation assurance, there is still a baseline risk that is non-zero.

Type 2: Information fault non-conformance. Non-conformance results from problems in transmission or interpretation of information. The aircraft is controllable and responsive, but may be acting on incorrect information. Such faults can usually be quickly corrected once recognized.

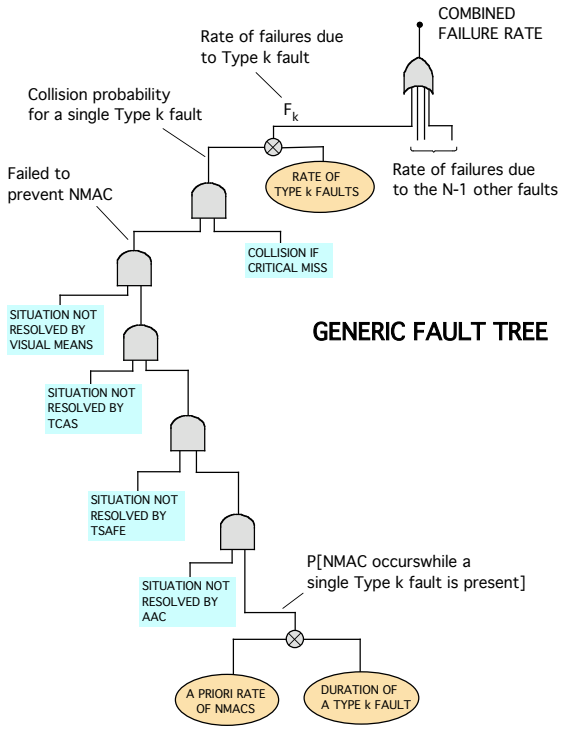

Figure 3. Generic fault tree for assessing level of safety.

Type 3: Control fault non-conformance. The aircraft has the correct information, but cannot be flown to conform to the specified 4D trajectory. Among possible causes could be extreme turbulence, engine failure, or a 4D trajectory outside the performance capabilities of the aircraft. These faults occur less frequently than information faults, but recovery generally takes longer.

Type 4: Service Interruption. All AAC services halt for all aircraft in the sector. This could result from major failures of the computer system or from a deliberate decision to shut down in response to anomalous events.

The examples that follow employ the following common parameters for all fault types: *traffic density* 0.001 aircraft per cubic NMI; *mean relative speed* 560 kt; *NMAC vertical miss distance* 500 ft; *NMAC horizontal miss distance* 100 ft; *collision cross section* 10,000 sq ft.

Other parameters vary with the fault type. Table 1 summarizes their nominal values.

Table 1. Probability Estimates for Fault Types

Parameter	Type 1	Type 2	Type 3	Type 4
Fault Rate (/hr)	1.7	1E-04	1E-05	5.7E-05
Duration (sec)	2144	60	120	1800
p AAC fails	1E-06	0.10	0.20	0.005
p TSAFE fails	0.10	0.10	0.20	1.00
p TCAS fails	0.20	0.20	0.30	0.08
p Visual fails	0.30	0.30	0.40	0.30

For the first three fault types, we compute risk by simply inserting the specified parameters into equation 2. The calculation for Type 4 faults is complicated by the fact that some of the parameters are altered by the system response to the fault. We now discuss the modeling of the Type 4 fault. Then we compute the risks for all four fault types.

Risk Analysis of Outages (Type 4 Fault)

One of the first questions that arise when separation assurance becomes dependent on computers is whether the system safety can be assured despite occasional computer outages. It has been noted [Ref. 12] that if computer outages require that the system revert to fully manual separation, then it may be impossible to increase traffic densities beyond the level that can be handled by manual control

The envisioned AAC resolves this problem with a layered fault management strategy: The core system is made reliable and redundant so that service termination occurs only rarely. When subsystem

faults are detected, the first response is to transition to a back-up configuration in which automated separation is still being provided. The existence of multiple copies of the Separation Plan facilitates the smooth transfer of responsibilities to the back-up configuration. If failures are so severe that back-up subsystems cannot maintain services, the affected airspace shuts down using the procedure described below. Delays and inefficiencies may result from the sector closure, but safety will be maintained. We now analyze the safety implications of the sector closure procedure.

The assumptions regarding the Type 4 fault are:

- At time t=0, all ground-based separation services, including TSAFE, are terminated for all aircraft in the failure airspace.
- TCAS continues to function.
- Navigation services continue normally.
- The time required to restore service is greater than the transit time through the failed sector.

Note that the *recovery* of separation assurance services does not necessarily require repair of the specific problem that caused the fault. Recovery can be achieved by successfully transitioning to a back-up mode. If such recovery occurs quickly, the risk resulting from the service interruption can be greatly reduced. In the calculations here we make the conservative assumption that recovery does not occur soon enough to reduce risk.

We now impose the following AAC design and procedural restrictions:

- Additional aircraft will not be admitted into the failed sector until service is restored.
- The Separation Plan is not corrupted by the abrupt termination of service.
- When service is interrupted, aircraft, to the best of their ability, fly on their previously approved trajectories to the point of sector exit.
- If an aircraft deviates from its approved trajectory, it is unable to regain conformance with that trajectory.

While we consider only failures of a single supersector, the analysis can be applied to multi-sector failures. The main effect of expanding the scope would be to increase S, the time required to fly through the impacted area.

Equation 1 presented the generic formula for the risk generated by a single fault type. For the service interruption fault, this expression is valid, but it must be averaged over time because the traffic density in the failed airspace begins to decrease and the probability of non-conformance begins to increase after the fault occurs. Thus rather than finding F_4 by simply inserting numbers into equation 2, we must find the average value of the expression during the fault period. The model details are described in the following paragraphs.

Duration of a Type 4 Fault

From the viewpoint of a particular aircraft, the duration of the Type 4 fault, τ_4, is the time between service termination and service restoration. Under our assumptions, service is restored only when the aircraft exits the failed sector. Thus from the viewpoint of the aircraft, the duration of the fault, τ_4, is the amount of flight time within the failed sector that remains to be completed at the time the fault appears. If the fault appears at random, and the flight time through the failed airspace is S for an aircraft, the probability distribution for τ_4 will be

$$f_{\tau_4}(\xi) = \frac{1}{S}, \quad 0 \le \xi \le S \tag{3}$$

For simplicity, we use the same value of S for all aircraft, although a more complicated model could define a distribution of transit times that reflects the traffic routing in a particular sector.

Rate of NMAC Events

The rate of NMAC events can be computed using the encounter rate formula.

$$\lambda_{NMAC} = 4\,\kappa\, M_h\, M_v\, E[V_{ij}], \tag{4}$$

where κ is the volume density of traffic, M_h and M_v are the horizontal and vertical separations defining a NMAC, and $E[V_{ij}]$ is the expected value of the relative velocity between two aircraft. Because aircraft are forbidden to enter failed airspace, the density in the airspace decreases from its initial value, κ_0, to zero over a time period S. The time-dependent value of the traffic density is thus

$$\kappa(t) = \kappa_0\left(1 - \frac{t}{S}\right), \qquad 0 \le t \le S \tag{5}$$

For this analysis, we assume that using the average traffic density provides an accurate estimation of average risk, even though the density experienced by an individual aircraft may decline either faster or slower than the average. The rate of NMACs at time t after the fault occurs is

$$\lambda_{NMAC}(t) = 4\,\kappa_0\left(1-\frac{t}{S}\right) M_h\; M_v\; E[V_{ij}]$$
$$where \quad 0 \le t \le S \qquad (6)$$

α_4, Probability Encounter is Not Resolved

If two aircraft in an encounter both conform with the trajectories in the Separation Plan, the probability of collision is approximately α_0, the value that existed before the fault appeared. However if one or both aircraft in the encounter have deviated from the Separation Plan we assume that the Plan provides no protection.

Let the rate at which aircraft deviate from their conflict free trajectories be μ (deviations/hour). Because a Type 4 fault prevents re-conformance and desired trajectory updates, the rate of non-conformance is greater than it would be under non-fault conditions. For an initial estimate, we assume μ = 0.02/hr which is about 200 times greater that the value used in the absence of a Type 4 fault. Loss of conformance is still unusual: an aircraft that flies for the maximum 1800 second transit time has only a 1 per cent chance of losing conformance before exiting the failed airspace.

For an encounter that occurs at time t after the start of the fault, the probability that both aircraft involved in the encounter are still in conformance is exp[-2μt]. Thus, the probability of AAC failure for an encounter at time t is

$$\alpha_4 = 1-(1-\alpha_0)\exp[-2\mu t] \qquad (7)$$

Rate of Service Interruptions

The rate at which service interruptions occur, λ_4, depends on the combined reliability of the AAC and its supporting infrastructure. Today, an extended loss of service in an en route sector is an unusual event. For purposes of this analysis we will employ a rate of 0.5 outages/year for any given sector. Because each aircraft is assigned to a single sector at all times, this is the outage rate that will be experienced by system users.

F_4, Final Failure Rate for Type 4 Faults

For a given duration τ_4, the differential contribution to F_4 that is contributed by time interval dt at t is

$$dF_4 = \lambda_4\,\lambda_{NMAC}(t)\alpha_4(t)\beta_4\gamma_4\psi_4 P_{coll}\quad dt,$$
$$where \quad 0 \le t \le \tau_4 \qquad (8)$$

We can now find F_4 by integrating over the duration of the fault and averaging over all values of τ_4:

$$F_4 = \lambda_4\beta_4\gamma_4\psi_4 P_{coll} E_S\left[\int_0^{\tau_4} \lambda_{NMAC}(t)\alpha(t)\,dt\right] \qquad (9)$$

Here the expectation operator E_S averages the value of τ_4 as specified by the distribution of equation 3.

Figure 4 shows how the key variables change with time during the fault period for standard parameters (with S=1800 sec). Both traffic density and the associated rate of NMACs decrease linearly while the failed airspace is being evacuated. The probability of non-conformance increases in a nearly linear manner. The curve labeled "Instantaneous Failure Rate" is the risk experienced by an aircraft still in the failed airspace at the indicated time. There is less risk early during the fault period because aircraft have not had time to deviate from the Separation Plan. There is less risk late in the period because the failed sector has been almost completely cleared of traffic.

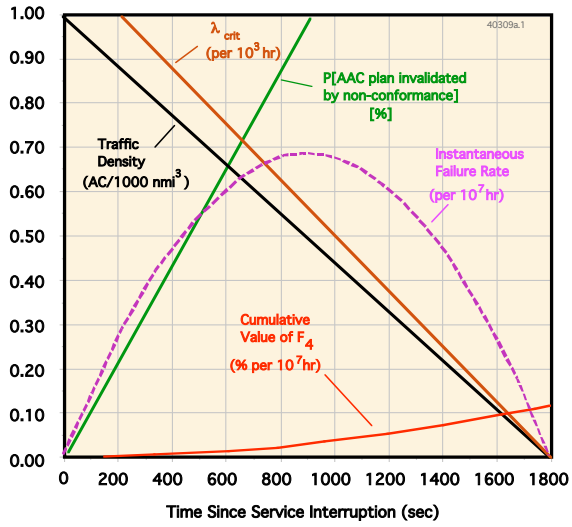

Figure 4. Evolution of risk during a service interruption (Type 4 fault).

The curve labeled "Cumulative Value of F_4" is the integrated failure rate considering all aircraft whose planned trajectories exit the failed airspace at or before the indicated time. The final value shown (at t=1800 sec) is F_4, the Type 4 risk defined by equation 9. For this example, the risk contribution of the Type 4 fault is well below the targeted level of safety.

Comparison of Risks for all Fault Types

Using equation 2 and the preceding model for service interruptions, we can compute the risk presented by each of the four defined fault types. Figure 5 provides a graphic depiction of the way in which each of the safety factors contributes to the overall risk for each fault type. In this figure, the vertical axis employs an inverted logarithmic scale in which higher ordinate values correspond to greater levels of safety (lower risk per hour of flight). The contribution of each safety factor can be determined from the segment lengths within each bar.

For this example, the total risk for all four faults combined is 1.8E-12, which is well below the targeted risk level. No one fault dominates the risk. The Type 1 failure is not a fault in the normal sense because the collision occurs when all elements of the system (including the AAC logic, TSAFE, and TCAS) function as designed. The failure rate of the core AAC logic is low by design, but it is not negligible. For purposes of this initial analysis we assumed that it is 1.0E-06 per flight hour, a realizable rate for a mission-critical system. TSAFE, TCAS, and visual avoidance contribute to reduce the overall failure rate by more than two orders of magnitude. The probability of collision resulting from a random critical miss reduces the overall failure rate by another four orders of magnitude.

The overall estimated failure rate for Type 1 encounters is 9.1E-13 per flight hour, about three orders of magnitude better than our design objective for the baseline system. If a margin of this magnitude remained after more precise analysis, it could be budgeted to reduce the safety requirements of other components such as the AAC logic.

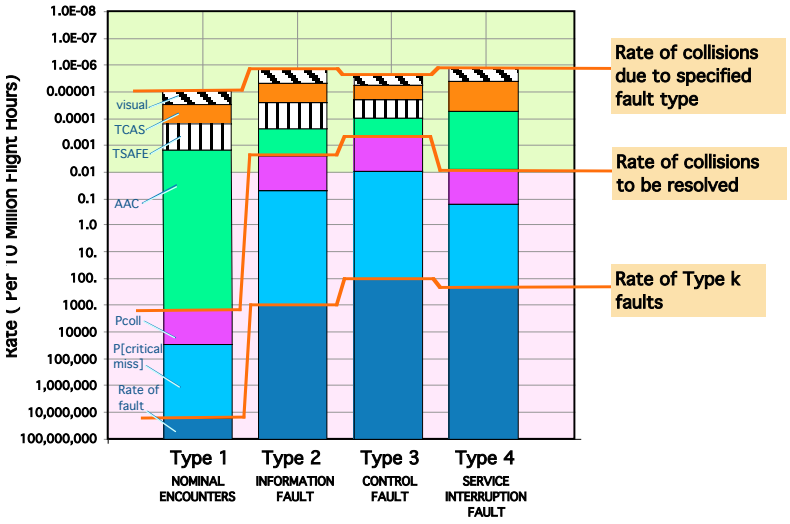

Figure 5. Level of safety calculation for four basic fault types.

Summary and Conclusions

This paper has described a fault-based approach to analyzing the safety of advanced separation systems. The particular concept analyzed, the Advanced Airspace System, has has a number of advantages that allow it to maintain safety despite high traffic densities and possible subsystem failures. Among these advantages are:

- Rapid detection of aircraft that deviate from their planned trajectory.
- Multiple copies of the Separation Plan to enable back-up to begin immediately with full awareness of prior actions.
- Ability of independent safety monitors and aircraft to inspect and validate the Separation Plan.
- Trajectories exchanged digitally rather than by voice to eliminate high error rates associated with voice communications.
- TSAFE as an additional separation assurance mechanism.

A broader safety advantage is associated with the support AAC provides to the safety management process. By comparing actual achieved separation with planned separation, the AAC can detect software and procedural flaws as well as lapses in execution. Monitoring and analysis of such events facilitates timely detection and correction of imperfections that might lead to accidents.

This work has addressed the troubling question of whether the possibility of service outages means that automation cannot be permitted to exceed the traffic densities that are safe to handle by manual control. The analysis suggests that service outages may be tolerable as long as certain safeguards are inherent in the design of the automated system. A key provision is the creation of a Separation Plan that is conflict free for an extended period and remains in effect while the system is reconfigured and traffic is rerouted.

The reason that the risk of service interruption is acceptable under the postulated AAC is that the conflict free planning horizon extends to the point at which an aircraft exits the sector. This is different from conventional control in which the effective conflict free horizon may be only four or five minutes - much less than the sector transit time. In this case, all protection is lost before exiting the sector, and some entity must assume full responsibility for separation assurance very soon after the appearance of the fault.

This preliminary analysis suggests that computer-based separation concepts with carefully designed safety features may be able to achieve higher safety as well as higher capacity and efficiency.

While these results are encouraging, important caveats remain:

1) The supporting system infrastructure (surveillance, communication, navigation, and computing architecture) must continue to evolve toward greater reliability and robustness. Current roadmaps for infrastructure evolution appear to provide a solid foundation for advanced concepts, and they should be pursued diligently.

2) The software, hardware, and procedures employed by advanced concepts must be designed for safety from the beginning. Certain innovative safety techniques must be the focus of their own dedicated research and development efforts.

3) The current analysis has considered only a mature AAC system. Interim systems may exhibit their own unique faults and should be subjected to their own separate analyses.

References

1. "Next Generation Air Transportation System Integrated Plan", Joint Planning and Development Office, U.S. Department of Transportation, December 12, 2004.

2. Erzberger, Heinz and Paielli, Russel A., "Concept for Next Generation Air Traffic Control System", Air Traffic Control Quarterly, Vol. 10 No. 4, 2002

3. Erzberger, Heinz, "Transforming the NAS: The Next Generation Traffic Control System", ICAS-2004, August 30, 2004, Yokahama, Japan.

4. Donahue, G.L.,"A Simplified Air Transportation system Capacity Model", Journal of ATC, April-June 1999.

5. Hoekstra, J.M. et al, "Free Flight in a Crowded Airspace?" 3rd USA/Europe Air Traffic Management R&D Seminar, Naples, 13-16 June 2000.

6. Andrews, J.W. and Welch, J.D., "Workload Implications of Free Flight Concepts", 1st USA/EUROCONTROL Air Traffic Management R&D Seminar, Saclay, France, June 1997.

7. ICAO Annex 11 - Air Traffic Services, 12th edition, International Civil Aviation Organization, July 1998.

8. Blom, Henk A.P., Klopstra, Margriet B., and Bakker, Bert, "Accident Risk Assessment of Simultaneous Converging Instrument Approaches", 4th USA/Europe Air Traffic Management R&D Seminar, Santa Fe, 3-7 December 2001.

9. Final Rule: Federal Aviation Administration 14 CFR Part 91, Docket No. FAA-1999-5925 [Amendment. No. 91-261], RIN 2120-AG82.

10. "Couluris, G.J., "Detailed Description for CE6 En route Trajectory Negotiation", NAS2-98005 RTO-41, Titan Systems Corporation and Seagull Technology Inc., November 2000.

11. Paielli, Russell A. and Erzberger, Heinz, "Conflict Probability Estimation for Free Flight", NASA Technical Memorandum 110411, October 1996.

12. "Policy Paper: Automation", Guild of Air Traffic Control Officers, September 5, 2003.

Author Biographies

John W. Andrews

John Andrews is a senior staff member in the Air Traffic Control Systems Group at M.I.T. Lincoln Laboratory. He was one of the senior system analysts during the development of the Traffic Alert and Collision Avoidance System (TCAS) and has made contributions to tracking techniques, algorithm analysis, human subject flight testing, and the

modeling of pilot visual acquisition performance. He has served as a consultant to the National Transportation Traffic Board in the investigation of mid-air collisions. Mr. Andrews holds a B.S. in Physics from the Georgia Institute of Technology and a Master's degree in aeronautical engineering from the Massachusetts Institute of Technology.

Heinz Erzberger

Heinz Erzberger joined Ames Research Center in 1965 after receiving his Ph.D. in electrical engineering from Cornell University. During his career at Ames he pioneered several concepts for improving the safety and efficiency of aircraft and air traffic control operations. Early in his career he developed the basic algorithms for four-dimensional guidance and for fuel-optimum flight management. In recent years he designed the Center-TRACON Automation System (CTAS) to help controllers improve the efficiency and safety of air traffic control. The FAA has deployed the CTAS tool for managing arrival traffic, the Traffic Management Advisor, at major airports in the US. He has published more than 90 papers and has received numerous honors and awards for his research. He holds two patents in the design of air traffic automation tools. He is a Fellow of the AIAA and a Fellow of Ames Research Center. He was awarded the ASME's Holley Medal in 2001 and the AIAA's Reed Aeronautical Award in 2004, the AIAA's highest honor for achievements in aeronautical science and engineering.

Jerry D. Welch

Dr. Welch is a senior staff member in the Air Traffic Control Systems Group at M.I.T. Lincoln Laboratory. He was on the team that initiated the development of the FAA's Mode S beacon system, leading the Mode S transponder design and standardization effort. He lead Lincoln Laboratory's program to develop surveillance techniques for TCAS collision avoidance. He supported the FAA and NASA in the development of the Center Terminal Automation System (CTAS). His current research includes aviation system capacity analysis and decision support tool benefits analysis. Dr. Welch received the S.B. and S.M. in Electrical Engineering from M.I.T. and the Ph.D. in Electrical Engineering from Northeastern University.

Acknowledgements

This work is sponsored by NASA Ames under the Air Force Contract #F1968-00-C-002. Opinions, interpretations, recommendations, and conclusions are those of the authors and are not necessarily endorsed by the United States Government.