2020 Foresight - A systems-engineering approach to assessing the safety of the SESAR Operational Concept

Derek Fowler, Eric Perrin, Ron Pierce
EUROCONTROL
Brétigny-sur-Orge, France
derek.fowler.ext@eurocontrol.int, eric.perrin@eurocontrol.int, ron.pierce.ext@eurocontrol.int

Abstract: The paper explains why a new approach, both broader and more rigorous than that traditionally followed in ATM, is needed for the safety assessment of the major operational and technology changes that are planned for introduction into European ATM over the period up to 2020 and beyond. It presents the theoretical basis for what is a "systems-engineering approach" and describes how that is being applied to the preliminary work on the safety assessment of the SESAR Operational Concept.

Keywords: safety, assessment, safety-case, assurance, SESAR

1. INTRODUCTION

European airspace is fragmented and will become increasingly congested as traffic is forecast to grow steadily over the next 15 years or so. ATM services and systems are not sufficiently integrated and are based on overstretched technologies. Therefore, to meet future air traffic needs, the European ATM services must undergo a massive operational change, supported by innovative technologies.

SESAR - the Single European Sky ATM Research Programme[1] - is the means of defining, designing and delivering the operational and technological changes necessary to achieve a more efficient, better integrated, more cost-effective, safer and more environmentally sustainable European ATM infrastructure by the year 2020.

During the SESAR Definition Phase, the European Commission initiated Episode 3 (EP3), a three-year project to undertake a first assessment of the SESAR Concept of Operations. Closely related to EP3 is an *a priori* safety assessment of the SESAR Concept, to assess as far as practicable that the Concept has been specified to be acceptably safe - this work is based at EUROCONTROL's Brétigny site.

This work is a preliminary safety assessment, laying the foundations of the process and methods, and gathering initial results, that will then feed into the main SESAR programme.

The specific requirements that the safety assessment has to satisfy are as follows:

- it must be soundly based from a theoretical perspective
- it should be pragmatic and of maximum benefit to SESAR Stakeholders
- it should make maximum use of, and contribution to, the work being undertaken on EP3
- it must preserve the integrity required of the safety-assessment process itself.

Reference [1] explained why the traditional, failure-based approach to safety assessment in European ATM was insufficient for the assessment of new operational concepts, and proposed a "broader approach to safety assessment".

Reference [2] presented an Integrated Risk Picture (IRP) of the causes of ATM-related accidents, based on analysis of accidents and incidents up to year 2005, and showed how it could be used to predict the effect of future changes to the ATM system on the risk of an accident.

This paper builds on, and integrates the approaches proposed in [1] and [2] and shows how what has become the "systems-engineering approach to safety assessment" is starting to be applied to the SESAR Operational Concept circa 2020.

2. THEORETICAL PERSPECTIVES

2.1 Risk Basics

Reference [1] uses the simple example of a car airbag to explain why a safety assessment must consider the positive (risk-reducing) properties of a system as well as its negative (risk-inducing) properties. Clearly, we would want an airbag to be reliable - i.e. to operate when it is needed - and to have high integrity – i.e. not to operate when it is not needed. However, above all, we would want it to be effective (in preventing death / serious injury) when it does operate; this would depend on its size, shape, construction and speed of deployment etc – i.e. on its functional / physical and performance properties.

[1] Equivalent to the US NextGen Programme

This is illustrated in Figure 1 which shows the risk (to the driver) with and without the airbag – i.e. R_U and R_A respectively. The safety case for the airbag depends on its saving far more lives / preventing serious injury, when operating as intended (the green, right-to-left arrow) than any deaths / serious injury that might be caused in the event of its failure or spurious operation (the red, left-to-right arrow).

There are a number of very important points to note about this diagram:

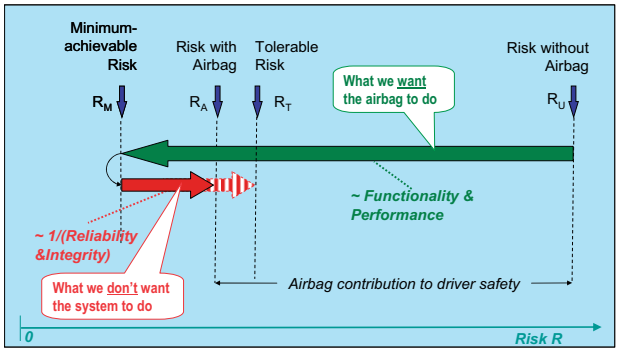

Fig 1: Risk Graph for a Car Driver's Airbag

- R_U has nothing to do with the airbag – for this reason we call it *pre-existing* risk
- R_M is the theoretical minimum risk that would exist in the complete absence of failure of the airbag – it is not zero, because there are some accident scenarios that an airbag cannot mitigate against
- the risk increase R_A - R_M is caused entirely by failure of the airbag - thus we call it *system-generated* risk
- the safety case must show at least qualitatively that $R_A << R_U$
- if we now introduce R_T (the maximum tolerable level of risk) then a most interesting conclusion emerges: the maximum tolerable failure rate of the airbag, the length of the red arrow (R_T - R_M), depends on the length of the green arrow (R_U-R_M) - i.e. on how successful the airbag is in reducing the pre-existing risk
- if, as we desire, (R_T - R_M) << (R_U - R_M) then the overall risk actually achieved (i.e. R_A) is much more sensitive to changes in the length of the green arrow (i.e. to changes in functionality and performance) than to proportionate changes in the length of the red arrow (i.e. to changes in reliability and integrity)[2].

The above points also raise some very important questions regarding the origins and use of traditional risk-classification schemes. It is why the above safety assessment has adopted a more considered approach, based on IRP, as described later.

[2] For ATM, R_A is typically 6 to 7 orders of magnitude less than R_U!

2.2 Application to ATM Risk

ATM is somewhat wider in scope and complexity than a car airbag but the same, fundamental principle holds good – i.e. its primary purpose is to mitigate pre-existing (aviation) risk.

This can be illustrated by expressing the three layers of ATM, described in the ICAO Global ATM Concept [3], in the form of a Barrier Model[3] as shown in Figure 2.

It is self evident that aviation (like driving) is inherently risky! Even for a single aircraft, there are risks of uncontrolled and controlled flight in terrain (UFIT and CFIT). For multiple aircraft in the airspace, there are additional risks of mid-air collision (MAC) and collision between aircraft on the ground.

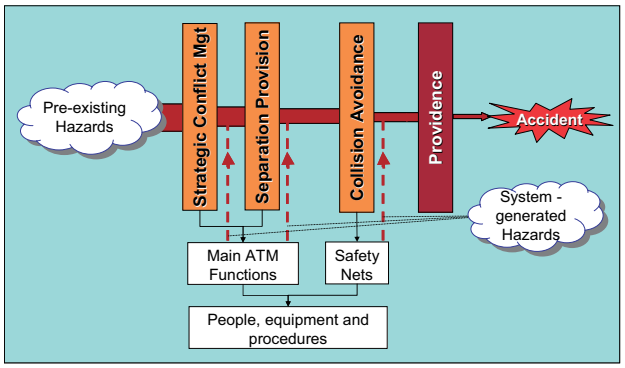

Fig 2: Simple ATM Barrier Model

These risks (or hazards) are inherent in aviation and therefore can be considered as "pre-existing" as far as ATM is concerned - they form the input to the model.

The barriers act in rough sequence from left to right and effectively filter out a proportion of the pre-existing hazards. The final barrier reflects the point that, even when all three layers of ATM have been unable to remove a hazard, there is a (usually high) probability that an actual accident will not result.

As the main barriers are provided by the elements of the ATM system, it is the ATM system functionality and performance that determines the effectiveness of the barriers in removing the pre-existing hazards. Of course, elements of the ATM system can fail or operate spuriously / incorrectly, giving rise to system-generated hazards, as defined above – these are shown in Figure 2 as inputs to the bottom of the model.

To paraphrase SESAR deliverable D4 [4], ATM must:

- "maximize its [positive] contribution to aviation safety", and
- "minimize its [negative] contribution to the risk of an accident".

In [1], these two aspects were referred to respectively as the *success* and *failure* approach; it was also emphasized that traditional ATM safety assessments

[3] Adapted from Prof James Reason's "Swiss Cheese" model – see http://www.bmj.com/cgi/content/full/320/7237/768

had usually assumed the former and focussed almost entirely on the latter.

What is crucial about Figure 2 for SESAR is that, in order to show that ATM achieves a tolerable level of risk overall, we need to understand the relationship between pre-existing risk (R_U), the positive and negative contribution of the three ATM Barriers, and the positive contribution of Providence[4] .

To demonstrate this quantitatively, we have combined the characteristics of the Barrier Model and Risk Graph as a single (slightly unconventional!) Fault Tree, as illustrated in Figure 3.

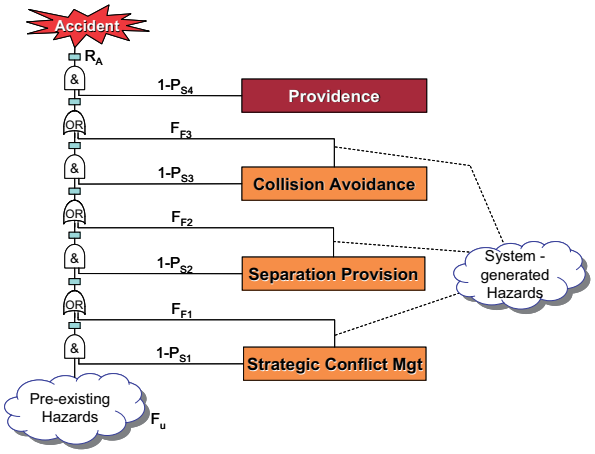

Fig 3: Fault Tree Version of Barrier Model

This Fault Tree allows us to compute the risk of an accident (R_A) from: the pre-existing, aviation hazards (and their frequencies F_U); the probability of success (P_{Sn}) of each barrier in removing those hazards; and the frequency (F_{Fn}) with which failure of each barrier introduces new hazards. Alternatively, of course, if we make the top-level risk our target (R_T) then, given F_U and access to historical accident and incident data, we can make informed judgements about what P_{Sn} and frequency F_{Fn} are required to be in order to satisfy R_T.

This risk model lies at the heart of the first stage in the integration of IRP accident model, being developed under EP3, into the *a priori* safety assessment. In practice, IRP uses a more detailed Barrier Model than the one described above - it exists in both current-ATM and post-2020 versions, as described in section 3.5 of the paper.

2.3 Safety Cases

Safety assessments are often done within the context of a *safety case* [5] which, like a legal case, comprises two main elements:

- a set of *arguments* - i.e. statements which claim that something is true (or false), together with
- supporting *evidence* to show that the argument is valid.

Safety arguments are normally set out hierarchically such that any particular argument statement is valid only if all of the next-level arguments are themselves valid - as shown, using *goal-structuring notation* (GSN), in Figure 4.

GSN is simply a graphical representation of an argument / evidence structure. In safety work it will usually start with the *claim* (Arg 0) that something is (or will be) *safe*; this is then decomposed such that it is true if argument statements Arg 1 to 4 are all true.

The *strategy* text should explain the rationale for that decomposition.

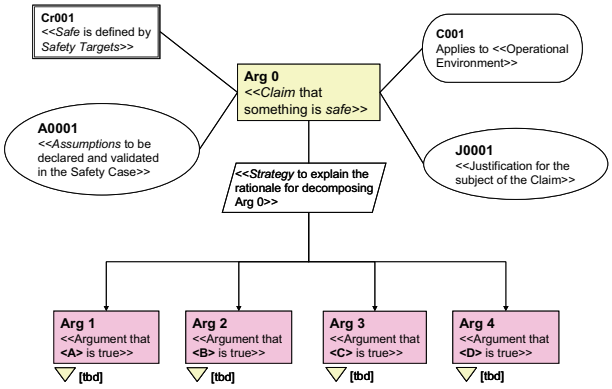

Fig 4: High-level Safety Argument

The claim is supported by vital contextual information:

- what is meant by *safe* is defined by means of *safety targets*, which may be quantitative and / or qualitative
- the *context* for the claim must include a description of the operational environment for which the claim is being made; sub-section 2.5 explains how critical this is to the validity of the claim
- *assumptions* are usually facts on which the claim depends and over which the organization responsible for the safety case has no managerial influence - e.g. traffic will increase by x% per year
- if the claim relates to a major change to a safety-related system, it is good practice to provide a *justification* for that change.

The arguments would then be further sub-divided until a level is reached at which a piece of documented evidence, of a manageable size, could be produced to show that the corresponding argument is valid. Further guidance on constructing safety arguments is given in [5].

2.4 Safety Assurance

There, however, are two problems with the simple argument / evidence approach.

[4] Providence is unique in that it cannot make a negative contribution – i.e. it cannot introduce new risk

[5] This is consistent with the SESAR Safety Management Plan and European Operational Concept Validation Methodology, (E-OCVM) both of which take a "case-based" approach

The first is that, in itself, it gives no indication how the evidence should be obtained or how rigorous that evidence needs to be. As illustrated in Figure 5, this problem is addressed by bridging the lowest level of decomposition of argument and its supporting evidence with:

- *safety assurance objectives,* which state what has to be done to satisfy the related strand of the argument, and
- *safety assurance activities* which state how the safety assurance objectives will be satisfied – including the tools and techniques etc to be used.

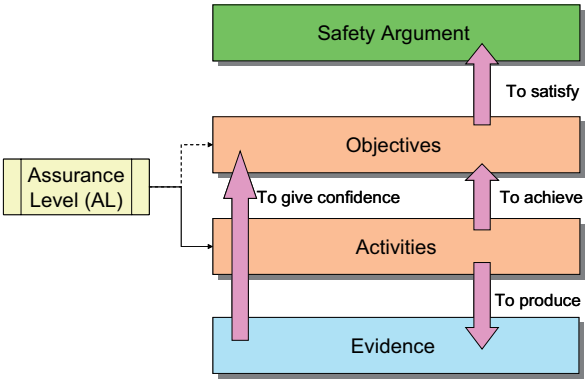

Fig 5: System-level Assurance Structure

The output of the assurance activities is then the evidence that we needed to show in turn that each objective has been met and eventually, therefore, that the safety argument is satisfied.

In many assurance-based approaches, the objectives and activities are, to some degree and extent, determined by an assigned *assurance level* (AL) – these ALs are usually derived by assessing the consequences of failure of the system element under consideration. For the initial SESAR work, we decided to make the objectives independent of the ALs and give only general guidance on the rigour required of the tools, techniques etc used in the safety assessment[6].

There is a second, related problem that safety assurance is often used to address - the fact that the integrity of software functions or human tasks, in particular, is very difficult to show in a direct way - through, for example, analysis of test results - that such safety requirements have been satisfied in implementation.

This is reflected in, for example, airborne software standard DOD 178B [6] and system / software standard IEC 61508 [7] both of which are assurance based. EUROCONTROL itself has adopted such an approach in the safety assessment of the individual software, procedure and (under development) human elements of ATM systems but the application to the overall system, as described herein, is new.

[6] We did not feel that we had the competence or authority to be prescriptive about this – therefore we left it to individual safety assessments / safety cases to justify that the evidence produced is *trustworthy* – see Arg1.4 in section III.

2.5 A Requirements-engineering Model

Capturing a complete and correct set of safety requirements is fundamental to any *a priori* safety assessment.

For the initial SESAR work, we have adopted the simple, but rigorous, requirements-engineering (RE) model shown in Figure 6.

In this model, *systems* exist in the *real world*. The part of the real world that influences the system, and into which the system provides a service, is known as the *application domain*. Users of the service exist in the application domain. The system interacts with the application domain through an *interface (i/f)*.

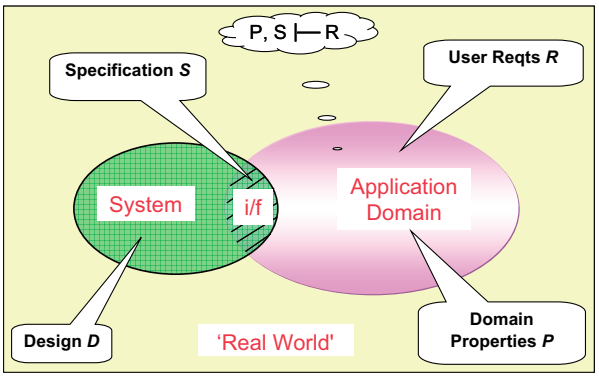

Fig 6: Requirements-engineering Model

User requirements are what we want to make happen in the application domain and are defined in that domain - not in the system.

A *specification* is what the system has to do across the interface in order that the user requirements can be satisfied - i.e. specifications take a "black-box" view of the system.

The formal notation in the "bubble" in Figure 6 defines the key relationship that the specification *S* satisfies the user requirements *R* only for a given set of properties *P* of the application domain; if any one of these three sets of parameters is changed then requirements-satisfaction argument is invalidated until one of the other sets is also changed, in compensation.

Design describes what the system itself is actually like and includes all those characteristics that are not directly required by the users but are implicitly necessary in order for the system to fulfill its specification and thereby satisfy the user requirements. Design is essentially an internal, or "white-box", view of the system.

The distinction, and relationship, between requirements, specifications, domain properties and design are not merely academic niceties but provide the essential foundations for developing systems that do, and can be shown to do, everything required of them. In section 3, it is shown how this is crucial to the construction of a safety argument for the

completeness and correctness of the safety requirements.

3. APPLICATION TO THE SAFETY ASSESSMENT OF THE SESAR OPERATIONAL CONCEPT (CIRCA 2020)

The first point about the SESAR safety assessment is that it is argument-driven – there is a process to be followed but that comprises a series of activities defined as in section 2.4 above.

3.1 High-level Safety Argument

A typical high-level safety argument for SESAR is shown in Figure 7, using the En-route phase of flight as an example.

The top-level *claim* (Arg 0) is that En-route operations for the specified Operational Environment (C001) will be *acceptably safe*, as is defined by the *safety targets* – see sub-section 3.5 below.

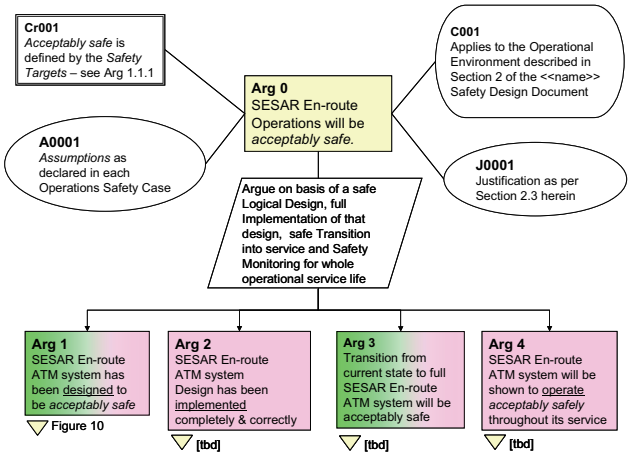

Fig 7: High-level Safety Argument – SESAR En-route Operations

The key *assumption* at this stage is that SESAR will deliver by 2020 a 1.7-fold increase in capacity [8] and that this will be fully taken up by a corresponding increase in traffic levels[7].

The *justification* for SESAR stems from its benefits to the airspace users, including improvements in the capacity, cost- effectiveness, efficiency, environmental sustainability, and flexibility of the overall ATM service.

The claim is then decomposed into the four *arguments*.

Arguments 2 to 4 reflect normal ATM safety practice and are the responsibility mainly of the SESAR stakeholders involved in the implementation of the SESAR Concept (Arg 2) and subsequent SESAR-based operations (Arg 3 and 4). However, it is important to note that Argument 1 applies to the whole SESAR Concept as applicable circa 2020; therefore, because the SESAR Concept is being implemented in stages, the term *transition* in Argument 3 includes the safety of each stage of this phased deployment of the end system, taking account also of the fact that developments in adjacent airspace may be being deployed in a different sequence and/or to different timescales – it is part of the current SESAR work to consider how to address that problem.

The main focus of the current work, however, is Arg 1.

3.2 Decomposing Arg 1

In order to decide how best to decompose Arg 1, we first needed a suitable interpretation of the RE model of Figure 6.

This interpretation is shown in Figure 8. As a (literally) logical representation, the RE model lends itself well to being expressed as a safety argument.

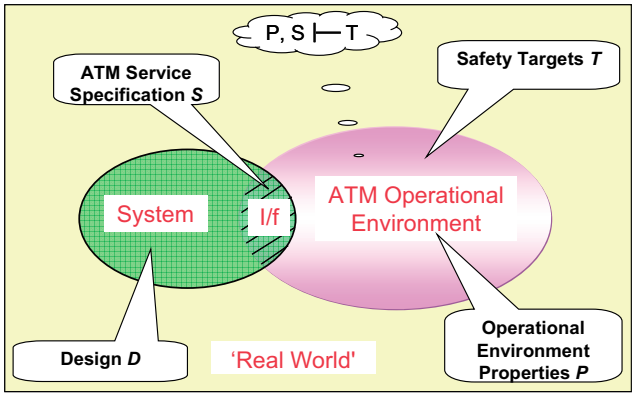

Fig 8: ATM Requirements-engineering Model

Our strategy for developing the argument was as follows:

- firstly to ensure that the properties *P* of the operational environment was properly described. Fortunately, most of the necessary information was readily available from detailed operational descriptions (DODs) produced by EP3 operational experts – it included the statement that the ATC separation minima would remain unchanged
- next to make an argument that the safety targets *T* were appropriate and correct for that environment
- then to make an argument that the ATM service specification *S* (to be produced as part of the safety assessment) would satisfy the safety targets *T* given the operational environment properties *P*.

Thus we could argue, at this stage, that the ATM service had been specified to be acceptably safe. The form of that *specification* is discussed in sub-section 3.5 below.

The next key step was to argue that the ATM system had been designed to satisfy the ATM service specification. It was clear that at this stage it would impracticable for us to attempt a physical design since that would more appropriately be left to

[7] This is the worst case because increasing traffic has an inherent linear or square-law negative affect on safety (depending on the type of accident being considered) for which improvements in the ATM system must compensate [9]

implementation (see Arg 2 above). Thus we needed find a more abstract representation of the system – which we called a *logical* design – as described in sub-section 3.6 below.

Two more issues needed to be addressed in order to complete a satisfactory argument:

- to show that the logical design was realistic – i.e. would be capable of being implemented in a physical system, comprising people, equipment and procedures
- to show that all the evidence under Arg 1 was trustworthy – see the discussion on safety assurance in section 2.4 above.

This is all summarized in GSN form in Figure 9 below.

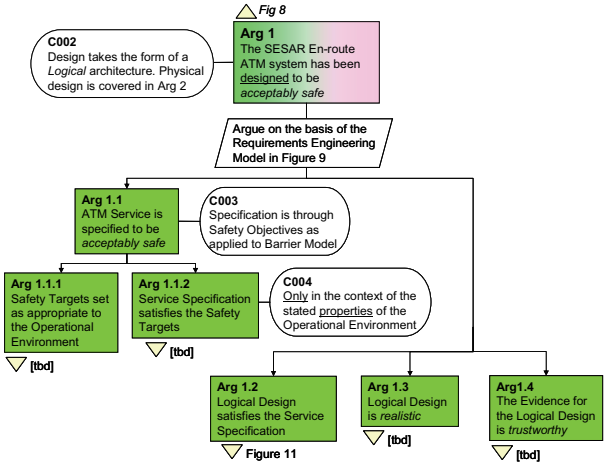

Fig 9: Initial decomposition of Arg 1

3.3 Decomposing Arg1.2

Making an argument for logical design is not simply a matter of showing traceability of the individual safety requirements (that form part of the design) back to the specification. This would ignore the possibility that the design as a whole was in someway functionally incomplete or internally incoherent or that new failure properties would emerge at the design level that were not apparent at the ATM-service level.

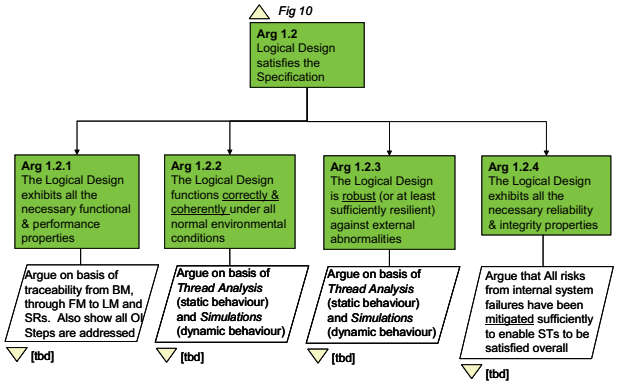

Fig 10: Decomposition of Arg 1.2

Thus we needed to show, as indicated in GSN form in Figure 10, that:

- The design has the functionality and performance attributes that are necessary to satisfy the ATM service-level specification
- The design will deliver that functionality and performance under all normal conditions of the operation environment that the system is expected to encounter in day-to-day operations
- The design is robust against (i.e. work through), or at least resilient to (i.e. recover easily from), any abnormal conditions of the operation environment that the system may exceptionally encounter
- The design has the reliability and integrity attributes that are necessary to satisfy the ATM service-level specification

3.4 The Safety Lifecycle

Albeit very much argument-driven, the safety-assessment approach has to end up with a process that is to be followed through the project lifecycle.

This is illustrated at the highest level in Figure 11, and shows that each safety-lifecycle stage comprises safety assurance *activities* which are determined by the safety *argument* and which produce *evidence* that the *argument* has been satisfied – the SESAR Safety Management Plan maps these on to the SESAR Project and E-OCVM lifecycle stages.

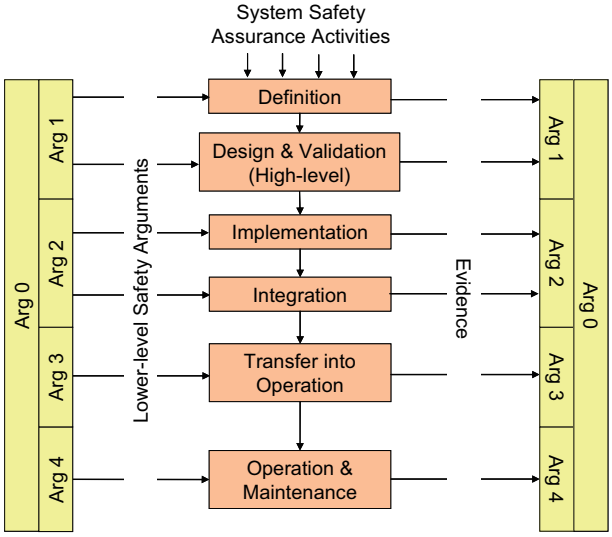

Fig 11: Overall Safety Lifecycle Process

It may be noticed that there is no reference to safety assurance *objectives* in Figure 11. This is because, when safety assurance is put into a safety argument framework, the safety assurance objectives become simply the lowest level of decomposition of the safety argument.

We can now apply the same general model to the Definition and Design & Validation phases of the lifecycle, as described in the next two sub-sections.

3.5 Definition Phase

Figure 12 provides an overview of the safety assurance process for the Definition phase of the safety lifecycle.

Each of the three steps consists of a number of assurance activities necessary to satisfy the associated safety argument (or, in the case of C001, provide vital contextual information to support the argument).

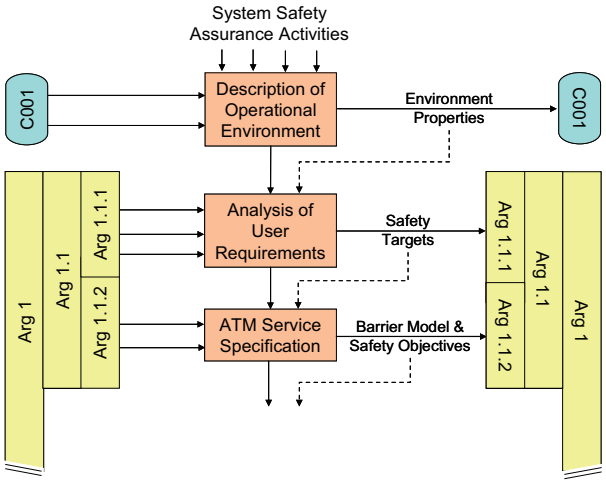

Fig 12: Safety Assurance in Definition Phase

It is impracticable to present the full scope of these activities within this paper – as an example however, the description of the operational environment for SESAR En-route operations would include:

- airspace structure and boundaries
- types of airspace / ICAO classifications
- route structures (as applicable) and any restricted airspace (temporary or otherwise)
- traffic characteristics and complexity
- aircraft ATM capabilities
- air traffic services to be provided, and associated separation standards

It would also need to identify those properties of the environment that are crucial to the safety assessment (C001).

The needs of the airspace users are analyzed from a safety perspective. From this analysis, safety targets are derived so as to satisfy those user needs. For SESAR, we have (provisionally) identified three types of safety target, for each of the four main phase of flight:

#1 the risk of an ATM-related accident (per annum) shall be no higher than for the pre-SESAR situation

#2 the risk of an ATM-related accident shall not exceed [tbd][8] per flight hour

#3 the risk of an ATM-related accident shall be reduced as far as reasonably practicable

The specification of the ATM service – see sub-section 3.2 above – is based on the *barrier model*[9] shown in Figure 13.

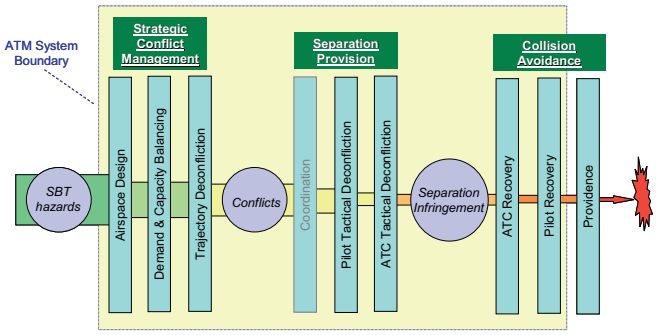

Fig 13: En-route / TMA Barrier Model

The inputs to the model are the pre-existing hazards of conflicts between, what are known on SESAR as, the *shared business trajectories* – in effect, these are the ideal trajectories that the each user would like to fly, unconstrained by any other considerations.

The ATM service specification then comprises:

- a functional description of the operation of each barrier and, qualitatively, how barrier contributes to the removal of the pre-existing, SBT hazards
- *safety objectives* which specify, quantitatively, both the minimum probability of success, and the maximum rate of failure, of each barrier such that the residual accident rate is within the safety targets.

3.6 Design & Validation Phase

Figure 14 provides an overview of the safety assurance process for the main part of the Design & Validation phase of the safety lifecycle - activities related to Arg1.3 and 1.4 have been omitted from the diagram for the sake of clarity.

3.6.1 Functional Design

Even though Arg 1.2 is made in the context of *logical* design the first step in the process is development of a *functional* model of the ATM system. This is because:

- we found that to get sufficient assurance of the completeness of the logical design of the ATM system, with respect to the barrier model of the ATM service, it was necessary to bridge the two with a functional representation of the system, and
- it was considered to be good system-engineering practice for deriving the requirements of a functionally rich system like ATM.

A functional model (FM), in this context, is a high-level, abstract representation of the system that is

[8] A figure for each phase of flight is being obtained from the IRP model described earlier in the paper. Each figure will take account of the affect that increasing traffic will have on risk and will be set such that targets #1 and #2 are consistent .

[9] The version of the model shown applies to En-route and Terminal Area operations only – a slightly different Barrier Model has been developed for Airport operations

entirely independent of the logical design and of the eventual physical implementation of the system.

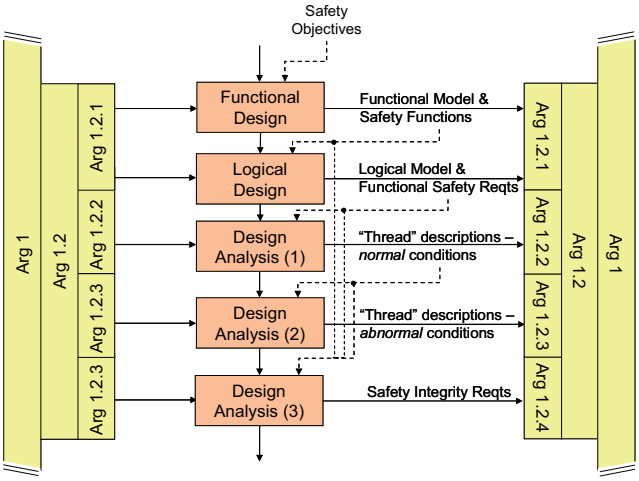

Fig 14.: Safety Assurance in Design & Validation Phase

The FM describes what safety-related functions are performed and the data that is used by, and produced by, those *safety functions* – it does not show who or what performs the safety functions.

It is not practicable to describe a typical FM in this paper but to illustrate the level and structure involved; however, to give some indication of its scope and complexity, Figure 15 shows the graphical representation of the SESAR FM for Terminal Area operations.

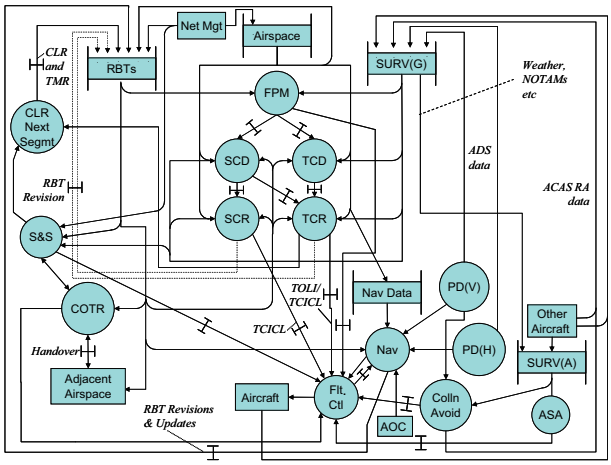

Fig 15: Typical SESAR Functional Model

Safety functions describe in detail what each element of the FM does and, where necessary, what level of performance is required of it.

A typical ATM safety function is *strategic conflict detection* (SCD). It is effectively an abstraction of one of the main role of the multi-sector planner controller / planning tools. It is normally triggered by *flight progress monitoring* (FPM) or directly from airspace / trajectory information, and provides a warning of conflicts between trajectories and between a trajectory and prohibited airspace. SCD needs to: be able to handle a mix of trajectory types, times, aircraft capabilities etc; be able to operate to full effectiveness for trajectories that are based on pre-defined RNAV routes or user-preferred routes; be able to operate to full effectiveness in a mixed traffic environment; to support continuous descents and climbs in Terminal Areas; and take account of the separation mode for each aircraft.

3.6.2 Logical Design

A logical model (LM) is a high-level, architectural representation of the system design that it is entirely independent of the eventual physical implementation of that design. The LM describes the main human tasks, machine-based functions and airspace structures and explains what each of those "actors" provides in terms of functionality and performance. The LM normally does not show elements of the physical design, such as hardware, software, procedures, training etc.

Figure 16 shows the graphical representation of the SESAR LM for Terminal Area operations.

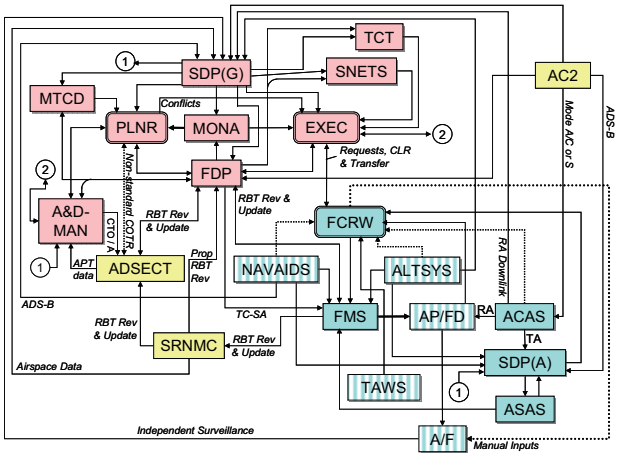

Fig 16: Typical SESAR Logical Model

Functional safety requirements (FSRs) describe in detail what each element of the LM must do from a safety perspective and, where necessary, what level of performance is required of it. As an example, the following are two of the 21 FSRs provisionally specified for the Arrival & Departure Manager (A&DMAN) and two of the 29 FSRs provisionally specified for the EXEC controller:

1. the AMAN sub-function shall compute a Controlled Time of Overfly (CTO) for waypoints extending out well into En-route Airspace (typically as far as 200 NM) and down to a CTA at the Final Approach Fix or at a final merge point
2. the AMAN sub-function shall generate speed advisories for Aircraft without an RTA capability
3. the EXEC shall resolve any conflicts, as follows:
   a). where the situation is time-critical, issue an "open-loop" clearance to one or both Aircraft involved, or
   b). where possible, and the situation is less time-critical, issue a trajectory change to

resolve the conflict but return the Aircraft to its original route, or

c). where proposed by the PLNR and judged appropriate, for crossing / passing traffic, delegate separation responsibility to the FCRW according to the agreed and authorized RBT

4. Whenever EXEC delegates separation responsibility to FCRW, he/she shall:

a). request the FCRW to accept responsibility for separation under ASAS procedures

b). pass the identity of the "target aircraft" to the FCRW

c). continue monitoring of these flights for possible unexpected behavior, and correct as necessary - otherwise the EXEC shall NOT provide instructions, advice or assistance to the FCRW unless specifically requested to do so by the FCR

d). retain responsibility for providing separation between all other aircraft and between those aircraft and the aircraft involved in the ASAS maneuver

e). resume separation responsibility for the Aircraft involved in an ASAS maneuver when advised by the FCRW that the maneuver is complete and the Aircraft involved are on diverging paths.

3.6.3 Design Analysis

Having produced a design that appears to have all the functionality and performance attributes that are necessary to satisfy the ATM service-level specification, the three stages of design analysis are intended to:

1. prove the correctness and coherency of the design, under all normal conditions of the operation environment that the system is expected to encounter in day-to-day operations
2. assess the behavior of the design under any abnormal conditions of the operation environment that the system may exceptionally encounter
3. assess the effects of internal failure of the ATM system on the risk of an accident

The only difference between the first two stages are the operational scenarios that define the normal and abnormal environmental conditions, and the requirement that in the first case the system must deliver full functionality and performance whereas in the second case the system may degrade somewhat provided it can be shown that any associated risk is very low because of the short duration and/or infrequency of the abnormal conditions.

Both stages examine the behavior of the system from a static and dynamic perceptive.

Much of the static assessment employs a modified version of UML system sequence diagrams used in *use case* analysis – which we have called *thread* analysis – illustrated in Figure 17.

The example scenario is that an aircraft requests a change of trajectory.

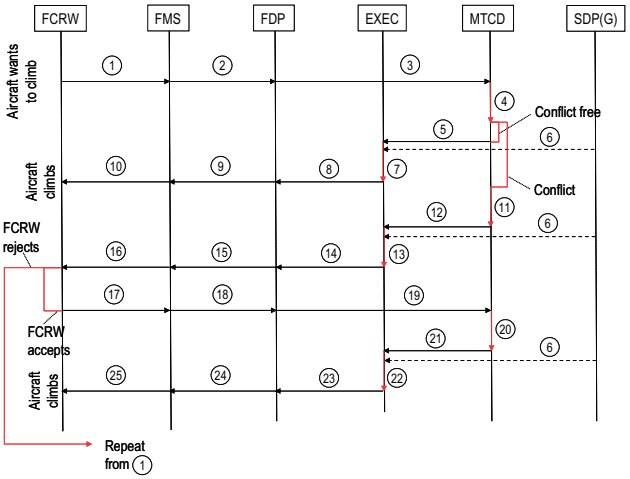

Fig 17: Thread Analysis (Illustrative)

It is left to the reader to work out the details (!) but the key points regarding the technique are as follows:

- the thread starts with an initiating event – "aircraft wants to climb" and/or one or more pre-conditions - e.g. the aircraft has a *level-4* capable FMS (not shown)
- the numbered horizontal arrows denote transactions between the (human and equipment-based) actors shown across the top of the diagram
- the numbered vertical arrows denote functions / tasks performed by an actor
- a dashed horizontal arrow denotes continuous flow of data – e.g. surveillance information (item 6)
- items 4 and 16 both have two possible outcomes, leading to branching of the thread
- each thread is continuous from initiation to conclusion
- each numbered item has an associated written description and a cross-reference to the related Functional Safety Requirement(s).

So far, the use of thread analysis on the SESAR safety assessment has shown the following benefits:

- it has led to a much better understanding of how the SESAR Operational Concept should work in practice – this should be of benefit to the whole EP3 validation program, not just to the safety assessment
- it has helped correct some errors, inefficiencies and inconsistencies in the logical model
- it has proved very effective in identifying missing or incorrect FSRs

Because the threads provide an understanding of the system behavior that cannot be shown solely through the LM and individual FSRs, it follows that the threads themselves should form part of the system design, and of the safety requirements.

Of course, what thread analysis cannot assess are the dynamic aspects of the system behavior – hence the safety assessment needs to make use also of the real-time and fast-time simulation exercises, which will form a very important part of EP3 and SESAR Development Phase. Nevertheless, thread analysis is a very cost-effective way of proving the correctness of the logical design under a wide range of normal and abnormal conditions.

Furthermore, by "breaking" threads, it should be possible to get a better understanding of the effects of **failures** within the system, and identify reversionary modes of operation – i.e. it can be used to enhance the conventional, failure-based safety assessment. Otherwise, Stage 3 of Design Analysis is effectively a conventional, failure-based approach to safety assessment and is not covered further in this paper.

3.7 Documenting the Results

Figure 18 shows the overall SESAR Safety Case structure.

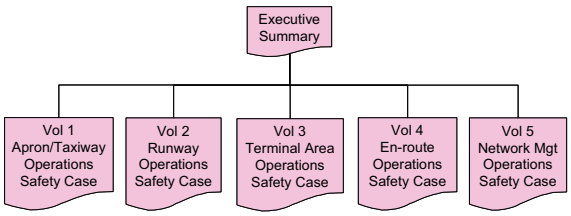

Fig 18: SESAR Safety Case Structure

This structure allows the various volumes of the Safety Case to be developed independently, provided all the interfaces and interdependencies between the phases of flight are dealt with in the appropriate volumes – in general, this proviso is taken care of by means of Safety Requirements placed on one phase of flight by another.

Figure 19 shows the main documentation structure for a typical volume of the SESAR Safety Case.

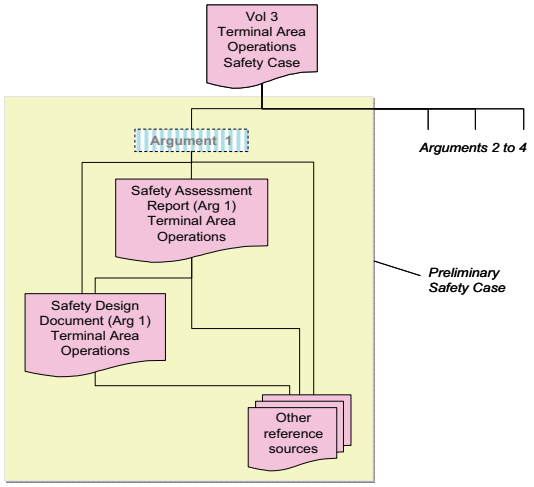

Fig 19: Typical Evidence Structure

The Safety Assessment Report (SAR) records the process, and presents the findings, of the safety assessment within the scope of Argument 1.

As explained above, the safety assessment is based on three models of the ATM service / System – i.e. barrier, functional and logical. Because the information associated with these models, and the description of the operational environment, is quite lengthy and because much of the information could be of significant use in non-safety areas as well, it was decided to place it in separate Safety Design Documents and to confine the SAR to the safety analysis of the three models.

4. CONCLUSIONS

The paper has explained why a broader and more rigorous approach than that traditionally followed in ATM, is needed for the safety assessment of the SESAR Operational Concept.

It has shown that what has become known as the "systems-engineering approach" to safety assessment has a sound theoretical basis.

It has also outlined how the approach is being applied to the major operational and technology changes that are planned for introduction into European ATM over the period up to 2020.

So far, we have validated the approach for the definition phase and the functional and logical stages of the design phase, of the safety lifecycle, for all four phases of flight and are well into developing threads for the initial design analysis for Runway and En-route operations.

Our experience to date has shown that the approach described herein is well able to meet the challenges of what looks to be one of the most wide-ranging ATM safety assessments ever undertaken. Nevertheless, provision has been made in the SESAR Development Phase for further development and refinement of the detailed methods, tools and techniques, within the above framework, as the SESAR safety assessment progresses through its lifecycle.

REFERENCES

[1] D Fowler, G Le Galo, E Perrin and S Thomas, So it's reliable but is it safe?, Proceedings of the 7th US / Europe Seminar on ATM Research & Development, Barcelona, July 2007

[2] E Perrin, B Kirwan and R Stroup, A systemic model of ATM safety: the Integrated Risk Picture, Proceedings of the 7th US / Europe Seminar on ATM Research & Development, Barcelona, July 2007

[3] ICAO Doc 9854, Global ATM operational concept, 1st edition, 2005

[4] SESAR Consortium, the ATM deployment sequence, D4, DLM-0706-001-02-00, January 2008

[5] EUROCONTROL, Safety case development manual, version 2.2, 2006

[6] RTCA, software considerations in airborne systems and equipment certification, DO-178B / ED-12B,

[7] IEC, functional safety of electrical/electronic[etc] safety related systems, IEC 61508, 2000 edition

[8] SESAR Consortium, air transport framework – the performance target, D2, DLM-0607-001-02-00a, December 2006

[9] Episode 3, White paper on the SESAR safety target, D2.4.3-01, 29 September 2008

AUTHOR BIOGRAPHY

Derek Fowler was born in Manchester, UK, in 1945. He was awarded a BSc degree in aeronautical engineering by the Royal Air Force College, Cranwell, UK, in 1968 and an MSc equivalent in aerosystems engineering at the same college 1975.

He served as an engineer officer in the Royal Air Force for 15 years before joining BAe Systems as a consultant engineer, project manager and then Head of the Laser Systems department. In 1990, he moved into the ATM field, with the UK National Air Traffic Services, as a senior project manager and then Deputy Director for Oceanic Systems. His considerable experience in systems engineering and interest in system safety were then combined, in 1998, when he took up successive senior technical positions with two of the UK's leading systems / safety consultancy companies. For the past 5 years he has operated as an independent safety consultant, setting up his own company, JDF Consultancy, in 2005. Working under contract for EUROCONTROL, he has provided safety expertise to more than 30 ATM programmes and, since January 2008, has been leading the initial safety assessment of the SESAR operational concept, at their Brétigny facility. He has many papers on ATM safety issues to his credit, most of them on the development of safety engineering techniques to keep pace with the increasingly rapid changes in ATM technology and operations.

Mr Fowler is a Charted Engineer and a Fellow of the UK Institution of Engineering and Technology

Eric Perrin was born in Saint-Etienne, France in 1969. He was awarded an Engineer degree in Aeronautics and Computer Science from the French Civil Aviation School (ENAC) in Toulouse in 1993.

He has more than 14 years experience of air traffic management, 8 of which have been spent on safety assessment and safety management. He joined EUROCONTROL in 2002 as GPS Ground-Based Augmentation System (GBAS) Manager. Prior to that, he worked as a Project Manager responsible for the design and development of aeronautical mobile communication systems. As EUROCONTROL Safety Assessment and Safety Case Manager, he currently leads a team of safety practitioners at Brétigny, south of Paris, working on a range of short- and medium-term ATM issues. He has made over 50 presentations on aviation technical issues (COM, satellite navigation, safety assessments) to international fora (GNSS, NAVSAT, ESREL, FAA Risk Conference, ATM R&D Seminars, etc.). He currently works on the safety validation of major aviation operational and technical changes and on safety techniques development to keep pace with foreseen air traffic management evolutions, in particular with SESAR.

Ronald H Pierce was born in Glasgow, UK in 1948, and studied at the University of Manchester where he gained his BSc and MSc degrees in computer science, the latter by research.

From 1975 to 1993 he worked for a number of the UK's leading software houses, gaining extensive experience in software engineering topics - compilers, program analysis tools and software engineering methods. Since 1993, he has worked as a Principal Consultant for CSE International Ltd in Flixborough, UK, specializing in software and system safety assessment for industry domains including ATM, railway control and signalling, and automotive, and has been responsible for the development of a number of safety cases for large-scale ATM projects such as new operations rooms and their associated equipment. He is currently working half of his time for EUROCONTROL Brétigny on an initial safety assessment of the SESAR operational concept. He is the secretary of the working group responsible for the maintenance of international functional safety standard IEC 61508 Part 3. He has published a number of papers in software engineering and safety topics and teaches courses in engineering safety management.

Mr Pierce is a Charted Engineer and a Fellow of the British Computer Society.

COPYRIGHT