SAFETY IN NUMBERS?
(LESSONS LEARNED FROM AVIATION SAFETY ASSESSMENT TECHNIQUES)

Stuart Baskcomb [1]**, Driss Ouedghiri** [2]

[1] Delta System Solutions GmbH, Packenreiterstr. 40, 81247 Munich, Germany.
Email: stuart.baskcomb@delta-system-solutions.com
[2] CASES, 16 Rue Jilali El Oraibi, Casablanca 20000, Morocco.
Email: drissoued@hotmail.co.uk

ABSTRACT

Although well-established and used countless times successfully in support of certifying many different aircraft, there are still aspects of the safety assessment process commonly used in aviation which can be improved.

A lot of the methodology and the techniques used in aviation safety are applicable to other industries, including space. This paper highlights the good, the bad and the ugly of aviation safety based on the authors' experience and makes proposals for lessons learned, the principles of which, can be read across to any domain. The complexity of today's aviation programs is increasing, with the greater reliance on software and complex electronics and the greater number of work-sharing partners.

This means it is more important than ever before to take a pro-active approach and review the now-traditional safety assessment techniques / methods in order to maximize confidence, effectiveness and efficiency.

1 THE CLASSIC AVIATION SAFETY PROCESS

As an introduction and to provide context for the rest of the paper, this section provides a brief overview of the benchmark safety assessment process used in the development of civil aviation products and where it sits in the overall accident prevention model.

1.1 "Holey cheese Batman, they've scored!"

It is widely-accepted that accidents occur due to multiple causes, all coinciding to allow a chain of events to lead to that undesired end effect. This is reflected in two very good analogies.

Firstly, the "Swiss Cheese" model was proposed by James Reason in 1990 and which has since been adopted by a large number of analysts and industry bodies, as well as being the subject of many critical reviews [1]. The basic premise is that the slices of Swiss Cheese represent barriers or layers of protection and the holes must all line up for an accident to occur – see Figure 1.

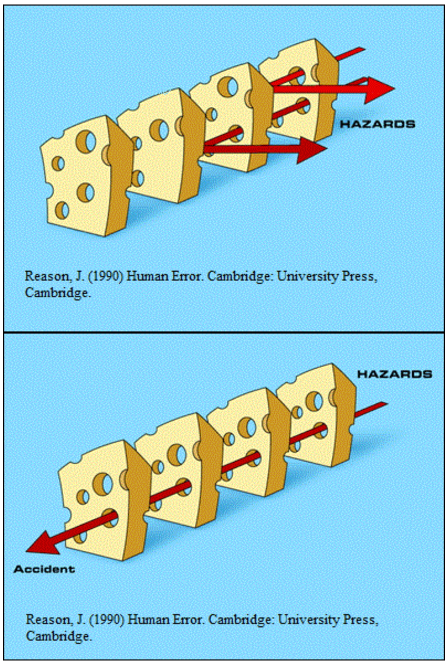

Reason, J. (1990) Human Error. Cambridge: University Press, Cambridge.

Reason, J. (1990) Human Error. Cambridge: University Press, Cambridge.

Figure 1. Reason's Swiss Cheese Model

Secondly, Captain Samir (Sam) Kohli in his award-winning book [2], considers an accident as a goal being conceded in a football match, where the pilot is the goalkeeper, and argues very well that it is unfair and unproductive to blame solely the goalkeeper for conceding a goal. The key contributing roles can be defined as:

- *Goalkeeper*: Pilot / Flight Crew
- *Outfield Players*: Aircraft Operator
- *Captain*: Aviation Authority
- *Coach*: Accident Investigator
- *Manager*: Government
- *Groundsman*: Aircraft Maintainer
- *Stadium & Pitch Management*: Airport Owner / Operator
- *Stadium & Pitch Architect / Builder, Kit & Ball Supplier*: Aircraft / System Manufacturer.

The scope of this paper is the safety assessment carried out in support of a product development cycle. This can be considered as one, or more, slices of Swiss cheese or the supplier of the football kit and architect / builder of the stadium and pitch.

1.2 The Classic Safety “V”

Aerospace Recommended Practice (ARP) 4761 [3] has been the civil aviation safety engineers’ bible for many years. This provides very good guidelines and methods for performing a safety assessment process as an inherent part of an iterative product development process. Civil aviation certification authorities do not insist on its use, however, a correct application of the methods (in whole or part) is recognized as a valid means of compliance to their safety requirements [4].

In the classic development process, the development of a product follows a “V” shape, reflecting time along the horizontal axis and depth of analysis along the vertical axis. It is requirements-based engineering with top-down, requirement-setting carried out on the left-hand side of the “V” and bottom-up, verification of the design on the right-hand side. The associated classic safety assessment process based on the guidelines and methods of [3] can also be summarized in this “V” model. See Figure 2.

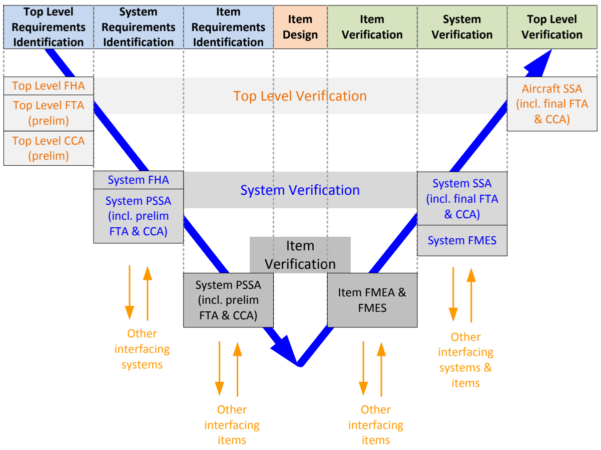
Top Level Requirements Identification
System Requirements Identification
Item Requirements Identification
Item Design
Item Verification
System Verification
Top Level Verification
Top Level FHA
Top Level FTA (prelim)
Top Level CCA (prelim)
Top Level Verification
Aircraft SSA (incl. final FTA & CCA)
System FHA
System PSSA (incl. prelim FTA & CCA)
System Verification
System SSA (incl. final FTA & CCA)
System FMES
Item Verification
System PSSA (incl. prelim FTA & CCA)
Item FMEA & FMES
Other interfacing systems
Other interfacing systems & items
Other interfacing items
Other interfacing items

Figure 2. Classic Safety “V” Development Cycle

How much benefit you gain from your safety assessment process depends on how you apply the guidelines and methods. There are, of course, the certification requirements to comply with. However, these reflect the minimum standard to achieve and they do not specify how much time and money you must invest. Therefore, the remainder of this paper discusses improvements and recommendations on how the safety assessment process is applied in order to optimize your product.

2 SAFETY IN NUMBERS?

Quantifying safety requirements and assessments has the obvious advantage of providing a crystal-clear case for whether compliance is achieved, or not. Project managers and chief engineers love this approach too, because it provides answers at a quick glance. However, it is all too easy to become obsessed with numbers.

2.1 Fault Tree Augmentation System

The Fault Tree Analysis (FTA) of a system can be considered as a link between certification requirements on one side, and reliability / failure data of components on the other side (see Figure 3). Excellent guidelines and training courses exist for FTA methodology. There are also some very good FTA software tools on the market, which provide very accurate calculations, no matter how large the model. Whilst in the main this is a plus point, it can also support the numbers obsession by allowing a monster fault tree (FT) to be built with thousands of events. Such “Frankenstein” FT’s take on a life of their own. They become impractical to verify and can lead to time-consuming debate / re-work, especially when unexpected answers, or answers *thought* to be incorrect, are produced.

The quantitative certification requirements for systems are based on historical data from the 50’s and 60’s and the overall (conservative) aim of new system designs being at least as reliable as those existing at the time [5], [6]. As well as a question mark over the relevance of such old data for today’s modern complex systems, there are other debatable aspects:

- What % of accidents are due to system faults? (10% assumed for requirement-setting)
- How many potentially Catastrophic failure conditions per airplane? (100 assumed for requirement-setting).

Component reliability data from service experience is like gold dust. But it is fool’s gold if the service experience and design does not compare close enough with your new system and/or component. What other sources of data can we use? Reliability predictions from previous projects are often re-used. But these are exactly what they say they are; predictions. Again, service experience is needed to close the loop. Unfortunately, the length of service experience is *inversely* proportional to the chance of re-using the same component in the same system and/or environment. Another source is standard, “text book” published data. This should be a last resort because it is generally overly-pessimistic and it is very difficult to compare the component and/or the service experience with that of the new product. However, it is often a first resort because it is thought to be an easy route to “compliance”.

In practice, this leads to manufacturers using their own, in-house derived factors to increase the reliability on paper. *Real* improvements in reliability can be made by implementing a Design for Reliability process [7].

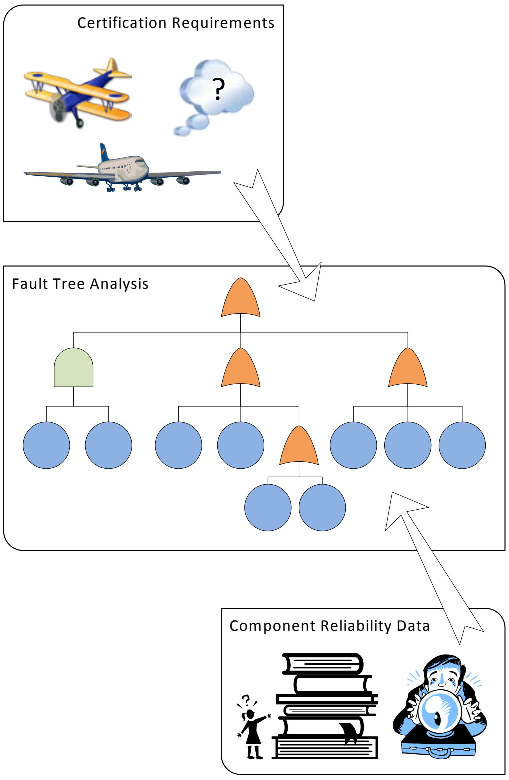

Figure 3. Simplified Quantitative Assessment Chain

In other words, the highly accurate and powerful FTA is surrounded by assumptions, approximations and opportunities for inaccuracies. In aviation, the requirement for the probability of occurrence of a Catastrophic failure condition is 1.00E-09 per flight hour (/FH) [8], [9]. As an example, consider a system FTA which has produced a result of 1.05E-09/FH. Does it make sense to spend significant effort arguing whether it is acceptable and/or re-work the FTA to reduce the number? In the end, this can just become a numbers game and the assessment is massaged until the "correct" answer is obtained.

2.2 The Alternative

FTA should be carried out to maximize the benefits of the safety assessment, rather than to maximize the use of the FT tool. It is easy to not see the woods for the trees when the creation and care of your Frankenstein is taking up all your time and effort.

FTA is an excellent assessment technique and can concisely illustrate safety arguments. The following key recommendations are made for applying FTA to any system in any domain:

- Do not blindly assess down to the lowest component level. Regularly take a step back to review and question what level is necessary *and* practical from a validation aspect.
- Focus more on the order of magnitude of failure probabilities, rather than the 20th decimal place.
- Assess contributing events in each fault tree relatively and review top contributors to see if / how they can be reduced. This can improve, not only the safety, but also the reliability of the system / product [7].
- Identify risks & hazards from service experience as candidates for a more detailed assessment.

Optimizing the efficiency of the FTA in line with the resolution of the inputs also has the advantage of justifying more time and effort being spent on different aspects of safety assessments that offer a better chance of improving (i) risk management and, (ii) design optimization. This can include making more out of the FTA for other, important aspects, some of which are qualitative and *all* of which are usually either completely excluded or rushed at the end as an "add-on" / "tick-in-the-box" exercise.

The rest of this paper discusses some of the areas where the most significant gains can be made and the lessons learned from the experience of the authors.

3 RECOMMENDATIONS

3.1 Software Influence

Software (SW) plays an ever-increasing role in aviation systems and, therefore, in the safety of these systems. SW does not fail in a random way like hardware. Errors in SW arise due to systematic faults generated during the development process. If SW cannot fail, it "only" needs to be proven that it is fault-free before it goes into service. Unfortunately, the typical amount of SW code contains too many scenarios for them all to be verified by test. Therefore, a certain level of robustness must be assured by an assessment of the SW development process.

Mature and widely-recognized methods and guidelines exist for how to allocate and decompose the SW development assurance level (DAL), as well as how to demonstrate a sufficient level of assurance has been achieved [10], [11]. Basically, the level of assurance required is commensurate with the criticality level of the SW function.

In the requirement-setting phase of development (left-hand side of "V" – see Figure 2), the preliminary system safety assessment (PSSA), including FTA, is used to derive the assurance level required for the SW. This is the starting point for the SW DAL strategy.

However, it is still common practice for safety assessments to basically exclude software from the verification development phase (right-hand side of "V" – see Figure 2). Verification of the SW development strategy is typically the responsibility of the quality assurance department and there is usually very little interaction between quality and safety.

Although it is not possible to quantify SW risk, FTA can still be used to improve how it is assessed. The inclusion of SW errors in the FTA would show where and how SW contributes to a hazard for further assessment (for example, extra attention in the assurance strategy, investigate alternatives / improvements in design) and for highlighting to the higher-level system safety assessment. It would also validate the earlier DAL allocation, help to define an assurance strategy and support any independence assessment.

3.2 H.G. Wells and Superman

Latent faults are key players in the chain of events required before an accident occurs. The provision of detection means (i.e. the avoidance of latent faults) is one of the fail-safe design principles included in the acceptable means of compliance from aviation authorities [12]. They also request special attention to be paid to faults that could be dormant, particularly "significant" ones (contributing to Hazardous or Catastrophic failure conditions), which should be avoided wherever practical [13].

However, it can be human nature to accept all too quickly an answer that we like. The safety engineer may not challenge the mitigation strongly enough and the focus is mainly on the latent faults, rather than the detectable ones. This means there is a danger of implicit *and* optimistic assumptions in the analysis, particularly regarding the timing of the detection and the capabilities of the human when (or if) they receive a fault indication. In extreme cases, time travel and/or super-human powers are required.

FTA can easily indicate latent events by the use of a different shape for the basic event, for example. However, the fault detection, indication and crew response for the detectable events are not so easy to show. In addition, it is usually not possible to indicate with FTA whether the intermediate events can be latent or not.

So, unless H.G. Wells designed your system and/or it is operated by Superman, a concerted effort and specific focus is needed to validate and verify the claims of detection, indication and crew response. It is important to assess the risk associated with this chain itself failing, including a consideration of the timing and the capabilities of the human.

3.3 Human Error

Like latent faults, Human Error (HE) is a key element in the chain of events leading up to an accident. In fact, HE has a hand in all accidents. It can come from the pilot, from the manufacturer and / or from the maintainer.

Assessing HE is a topic that has filled, and continues to fill, many a text book. It is a niche-within-a-niche skill and beyond the scope of this paper. However, it is possible to state here that HE is often not given the attention it deserves in aviation safety assessments.

HE (like SW) is usually left out from analysis, such as FTA mainly due to it being extremely difficult to quantify. Normally, it is considered only as part of the common mode analysis (CMA). The CMA guidance and checklist in Appendix K of [3] is often used in aviation safety. This points the safety engineer to various sources of HE as potential violations of claimed independence. The implicit assumption here is that the human makes no error or, at least, negligible error, which nobody would validate if it were explicitly written. One aspect of HE science which everybody agrees on is that humans make errors, as highlighted by Kohli's book dedication [2].

In a similar way to SW, HE should be included in the FTA to, at least, show where and how it contributes to hazards. This would allow HE contributions to be reviewed, assessed further and/or investigated to see if they can be reduced or eliminated. It would also allow them to be highlighted to the higher-level system safety assessment and help the validation of any independence claims.

Process Failure Modes and Effects Analysis (P-FMEA) should also be carried out as standard practice and with equal attention as the ubiquitous Design equivalent (D-FMEA). The results of which, should be integrated as far as possible into the FTA and rest of the safety case.

3.4 Independence

It is standard practice in aviation safety to assess independence, or rather, faults / events that can violate claimed independence. The aforementioned CMA checklist in Appendix K of [3] is a fundamental part of every good aviation safety toolbox.

FTA explicitly reflects independence claims with the use of AND gates. It relies on ensuring the same basic contributing events are identified with the exact same name. However, this is limited to the one failure condition or hazard being assessed.

Care must be taken when a sub-system FT is integrated into a higher-level system FT. If a "Frankenstein" FT has been created at sub-system level, it is often summarized as a single basic event in the higher-level system FT. Dependent faults, be it via common modes or single point failures, can therefore be missed if multiple sub-system FT's are used in one system FT.

In other words, simplifying the FTA approach will not only allow more time for assessing independence (including SW and HE contributions), but it will also reduce the risk of making false independence claims.

3.5 Interfaces

Interfaces are introduced as part of the system design and development, and usually for very good reasons. More often than not, the safety assessment will follow these same interfaces and/or introduce new ones – again, usually for good reasons. Unfortunately, failures and errors cannot read or listen, so why do we expect them to respect the interfaces and boundaries of the system(s)/sub-system(s) we have assessed?

There are two broad categories of interface which deserve extra attention in order to improve the way they are assessed and to avoid introducing any artificial ones; (i) System Interfaces (physical- or functional-driven) and (ii) Assessment Interfaces.

It has been a trend for many years now for companies to sub-contract bigger pieces of their system. This has led to suppliers taking on more of the safety assessment responsibility and they themselves taking on sub-suppliers. It is also common practice for companies to work together as risk-sharing partners and/or for very large companies to share the work between their different sites. This often means that the safety assessment is carried out by multiple teams, creating artificial assessment interfaces that do not exist physically or functionally and/or underscoring an existing interface, making it more difficult to assess.

Assessment interfaces can also be created between systems / sub-systems being assessed by the same team, even though there is some degree of interaction or function-sharing between the systems.

A process for capturing and managing safety requirements can sometimes be well-defined for interfaces, especially in a mature organization where the interface is with a supplier or a separate system. Whilst this is a good thing, it still relies upon the correct capture, understanding and management of the requirements.

A specific focus is needed on interfaces as part of the safety assessment, in order to be as confident as possible that all credible failure scenarios have been captured. For instance, a sub-system can be defined which overlaps an interface and a specific safety assessment conducted. The aim should be to eliminate as many boundaries as possible and, ideally, reach a point where the assessment is completely seamless.

3.6 The Organizational Balancing Act

It is good practice to have some degree of independence between the safety engineers and the system designers. The difficulty is striking the right balance. At one end of the scale, the safety engineers could become so independent, and so far removed from the system development, that their assessment is overly-conservative or overly-optimistic. At the other end, they could be so involved that they lose the ability to cast a critical eye over the design and development.

Companies usually follow one of two broad routes in their organization:

(i) *Functional-based*: safety engineers are in a separate team to the designers, etc., who are working on the same project.

(ii) *Project-based*: safety engineers are in an integrated project team (IPT) together with designers, etc., who are working on the same project.

Function-based offers a higher degree of independence, whilst project-based offers a higher level of knowledge and understanding to support a more comprehensive assessment, as well as increased potential for influencing the design (e.g. optimization via trade studies). Clearly, there is no right and wrong, however, the project-based approach is recommended.

The key disadvantage of the function-based approach is the risk of a wall being built between the teams. Think of it as a (very slow) team game of Tetris®, where a wall will gradually be built over time, without a continuous and conscious effort to remove the bricks that keep appearing. The designers and developers can fall into the trap of thinking less about safety in their day-to-day tasks because they believe it is somebody else's sole responsibility. Equally, safety engineers can become overly-critical and feel little or no ownership in the system design and development. Safety needs to be part of the thought-process of everybody involved, not just the safety engineers. It requires *all* players to work together to prevent the wall from being built.

A project-based organization can avoid the risk of such a "them and us" attitude developing. All the engineers working on the project are part of the same team, they all sit together and interact on a daily basis. The concern of an IPT organization is the loss of independence of the safety engineer. However, this can be managed via regular reporting / communication lines / reviewing between the IPT safety engineer and a functional safety manager or peer, who is outside of the IPT. It is easier to manage such an interface because the functional safety manager has an explicit and vested interest. Therefore, there should be no risk of a wall developing.

In other words, the improved level of involvement and knowledge of the system offered by a project-based organization far outweighs the effort required to support the sufficiently independent argument.

3.7 Requirements Capture

The principles of systems, or requirements-based, engineering has been used for many years in aviation system development. There are too-many-to-mention textbooks and guidance material available on this topic. Safety can be considered as a specialized branch of systems engineering, again requirements-based and again with a wide consensus on the principles.

However, safety case development does not always include requirement-specific documents, interface documents and requirement management tools that are common-place in system design and development. Instead, safety requirements can be a secondary aspect, with an attempt to capture requirements in the assessment documents and no formal management.

This is fine for the handling of safety requirements from, for example, the certification authorities and their flow down (e.g. the probability of occurrence of a Catastrophic failure condition shall be less than 1.00E-09/FH [8], [9]). However, it runs the risk of implicit assumptions and requirements (regarding, for example, system design and behavior, maintenance, interfacing systems), used in support of the safety assessment remaining hidden – sometimes even from the safety team.

When a dedicated safety requirements suite is used, explicitly documented and integrated into the requirements management process of the system design and development, it can provide the following advantages:

- Increase the confidence in the safety assessment reflecting accurately the system design and, more importantly, its behavior with faults.
- Speed up the production of the requirements, therefore increasing the influence on design (optimization), reducing the risk of late design changes and improving input from suppliers. (Safety requirements cannot be written too early in a project development lifecycle).
- Support the smooth integration of interfacing safety assessments
- Simplify the capture of lessons learned from one project to the next.

4 THE FAMOUS FIVE MISCONCEPTIONS

There are several, well-used phrases related to safety that are often repeated by people at all levels of an organization. These can be heard in a design review meeting or read in a magazine article or even a safety report. As the ICAO Safety Management System training concurs [14], these are popular misconceptions.

Such benign statements as those that follow, do not help develop a safety culture in a business. In fact, they will hinder it and a conscious effort to advise people not to say or write them will go some way to promote safety as a worthwhile, value-adding activity.

4.1 "Safety first"

Do not believe anyone who says this. If this were true, then companies would go bankrupt. By putting safety first and above all other business aims, no airplane or rocket would leave the ground. There has to be a balancing act between production and protection [2], [15] - see Figure 4. Safety should be one of the key management topics / processes in a business, but not the first.

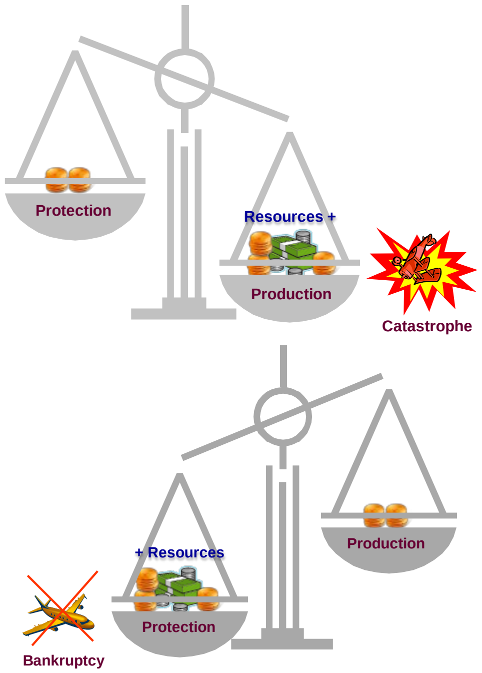

Figure 4. The Management Dilemma [15]

4.2 "Safety is everyone's responsibility"

This can have the opposite effect to that desired and encourage people outside of the safety team to switch-off to safety. Everyone has a level of *involvement* with safety and has a responsibility for safety *awareness*. Safety should be integral to the day-to-day work of any sound engineer / manager and not stand-alone. As Kohli recommends [2], it is better to promote reporting hazards and managing risks as the responsibility of everyone.

4.3 "If it ain't broke, don't fix it"

Why wait until your system is broken? Do you fully understand why it works or why the hazard did not occur? A white paper by Eurocontrol [16] discusses the principle of assessing the more frequent / less severe failure cases and learning from them to help prevent accidents occurring. Kohli's football match analogy [2] and Reason's Swiss Cheese analogy [1] both illustrate nicely the benefits to be gained by assessing the parts that worked as well as those that failed.

4.4 "74% of accidents are due to human error"

Whether it is actually 74%, or not, is irrelevant. The point here is that it is an over-simplification of the root *causes* of an accident. Human error is usually involved somewhere in the chain of events but it is very rarely the sole cause. Eurocontrol [16], Kohli [2] and Reason [1] all argue that accidents arise because of a combination of multiple factors, such as; equipment faults, human error, missing / inadequate procedures, environmental conditions, organizational influences, fatigue, luck, etc.

4.5 "If you believe safety is expensive, try an accident"

Safety should not be viewed as a cost or an overhead, rather it is an investment. Safety assessments carried out properly can lead to design optimization and improved efficiency in a business. This refers back to the earlier point regarding the assessment of things that work, as well as those that fail. An approach supported by Eurocontrol [16], amongst others.

5 A DELTA IN THE SAFETY ASSESSMENT PROCESS

Overall, the safety assessment process can be sometimes quite detached from that of the system development. The problem with this is that a good safety assessment cannot be carried out in isolation by the safety team. The input from the various system specialists is invaluable and it is vitally important to ensure the key people are engaged and feel some degree of ownership / responsibility of the safety assessment.

The safety process can be made more inclusive by overlaying the standard development process with safety workshops. These are designed to involve the system designers in the safety assessment from day one, *and* continuously thereafter.

The typical system development lifecycle involves checkpoints to pass through, reflecting that the design is maturing at the desired rate. As a detailed review of the design, they can also identify key risks to achieving on-time and on-spec delivery of the product. Another benefit of this more inclusive approach is that the safety workshops can be timed to support the preparation of the work necessary to pass through these checkpoints and maintain the influence on the design from the safety team.

Prompt lists and checklists can also be used to support the process and strike a good balance between a systematic approach to maintain consistency and a freedom of thought approach (e.g. brain-storming) to allow new concepts to be proposed.

If the lessons learned and recommendations discussed earlier are also incorporated into this more inclusive approach, a step-change improvement in the safety process can be realized, or a "Safety Delta". More inclusive in this context means a high level of involvement of designers, etc. in the safety process and improved communication lines between the safety team and the other system development teams, *in both directions*.

The knock-on benefits of this "Safety Delta" approach to the system design and development, compared to the traditional approach, include:

- Increased confidence over the safety of the design.
- Reduction in development costs by identifying changes earlier and reducing the risk of late, expensive changes.
- Optimization of the design and the development activities (including expensive testing).
- Improved safety culture throughout the organization.

6 REFERENCES

1. Reason, J., Hollnagel, E., Paries, J. (2006). *Revisiting the "Swiss Cheese" Model of Accidents*. EUROCONTROL Experimental Centre, France. EEC Note No. 13/06.

2. Kohli, Captain S. (2013). *Waiting… To Happen!: The tragedy of Air India Express Flight IX812.* CreateSpace Independent Publishing Platform. ISBN-13: 978-1494762889

3. S-18 Committee (1996). *Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment.* SAE, USA. ARP 4761.

4. EASA (2014). *Certification Specifications and Acceptable Means of Compliance for Large Aeroplanes.* EASA, Germany. CS-25 Amendment 15. Book 2, AMC 25.1309 *System Design and Analysis*, section 9(b)(3), pp 2-F-47.

5. EASA (2014). *Certification Specifications and Acceptable Means of Compliance for Large Aeroplanes.* EASA, Germany. CS-25 Amendment 15. Book 2, AMC 25.1309 *System Design and Analysis*, section 6(a), pp 2-F-41.

6. Gilles, D.L. (1983). *The Effect of Regulation 25.1309 on Aircraft Design and Maintenance.*

SAE, USA. Technical Paper no. 831406, pp.1-2.

7. Ouedghiri, D., Baskcomb, S. (2014). *Design for Reliability in Aviation (A must to improve Life Cycle Cost, Safety and Availability).* IAASS 7th Annual Conference, Germany, CD-ROM.

8. EASA (2014). *Certification Specifications and Acceptable Means of Compliance for Large Aeroplanes.* EASA, Germany. CS-25 Amendment 15. Book 2, AMC 25.1309 *System Design and Analysis*, Figure 2, pp 2-F-45.

9. FAA (1988). *Advisory Circular: System Design and Analysis.* FAA, USA. AC No: 25.1309-1A, section 7.d.(3) pp 7 and section 10.b.(3) pp15.

10. EUROCAE WG-42 / SAE SIRT (1996). *Certification Considerations for Highly-Integrated or Complex Aircraft Systems.* EUROCAE, France / SAE, USA. ED-79 / ARP 4754, chapter 15, pp 15-24.

11. RTCA SC-205 (2011). *Software Considerations in Airborne Systems and Equipment Certification.* RTCA, USA. DO-178C.

12. EASA (2014). *Certification Specifications and Acceptable Means of Compliance for Large Aeroplanes.* EASA, Germany. CS-25 Amendment 15. Book 2, AMC 25.1309 *System Design and Analysis*, section 6.b.(2)(v), pp 2-F-42.

13. EASA (2014). *Certification Specifications and Acceptable Means of Compliance for Large Aeroplanes.* EASA, Germany. CS-25 Amendment 15. Book 2, AMC 25.1309 *System Design and Analysis*, section 9.c.(6), pp 2-F-49.

14. ICAO (2009). *Safety Management System (SMS) Course Notes.* ICAO, Canada. Module 07 *Introduction to SMS,* slide 22.

15. ICAO (2009). *Safety Management System (SMS) Course Notes.* ICAO, Canada. Module 03 *Introduction to Safety Management*, slide 22.

16. Hollnagel, E., Leonhardt, J., Licu, T., Shorrock, S. (2013). *From Safety-I to Safety-II: A White Paper.* Eurocontrol.