A POSSIBLE APPROACH FOR ADDRESSING NEGLECTED HUMAN FACTORS ISSUES OF SYSTEMS ENGINEERING

C.W. Johnson† and C.M. Holloway*,

†Department of Computing Science, University of Glasgow, Glasgow, Scotland, UK, G12 8RZ.
johnson@dcs.gla.ac.uk, http://www.dcs.gla.ac.uk/~johnson

*NASA Langley Research Center, 100 NASA Road, Hampton VA, 23681-2199, USA,
c.m.holloway@nasa.gov

Keywords: human factors, air traffic management, systems engineering, safety-critical systems.

Abstract

The increasing complexity of safety-critical applications has led to the introduction of decision support tools in the transportation and process industries. Automation has also been introduced to support operator intervention in safety-critical applications. These innovations help reduce overall operator workload, and filter application data to maximise the finite cognitive and perceptual resources of system operators. However, these benefits do not come without a cost. Increased computational support for the end-users of safety-critical applications leads to increased reliance on engineers to monitor and maintain automated systems and decision support tools. This paper argues that by focussing on the end-users of complex applications, previous research has tended to neglect the demands that are being placed on systems engineers. The argument is illustrated through discussing three recent accidents. The paper concludes by presenting a possible strategy for building and using highly automated systems based on increased attention by management and regulators, improvements in competency and training for technical staff, sustained support for engineering team resource management, and the development of incident reporting systems for infrastructure failures. This paper represents preliminary work, about which we seek comments and suggestions.

1 Introduction

Automation has been widely proposed as a means of reducing the impact of operator error. In consequence, 'intelligent' control systems have been introduced in many industries (such as energy distribution, healthcare, transportation, and financial services). For example, agent based techniques have been developed that identify and respond to changing patterns in an operating environment without the need for human intervention. Also, inferential reasoning, including Bayesian analysis, has been used to automated decision making under uncertainty [5]. Techniques such as these have enabled increasing levels of integration both between and within individual systems and collections of systems [11].

These advances have come at a cost [12]. Increased automation has reduced the scope for operator intervention. Direct modes of control have been replaced by supervisory functions where, for example, railway signalers only intervene in response to abnormal conditions including the failure of system components. A host of human factors concerns arise as a result [2]. It is difficult to create and sustain high levels of situation awareness when operators are not continually involved in application processes. Increasing levels of integration also jeopardize an individual's ability to understand the diverse subsystems that support complex infrastructures [13]. There may also be problems at shift handover when operators have to brief their colleagues on the state of automated systems that continually evolve over time.

This paper looks beyond the operator and focuses on the teams responsible for installing, maintaining, and decommissioning safety-critical systems. Increasing levels of automation have increased the impact of administration and configuration errors on complex systems [6]. With less scope for direct operator intervention, many industries rely on systems engineers to correctly set the parameters that govern automated control systems. However, these teams often have relatively limited support. They typically rely on command line interfaces with few tools for error detection and correction. Unlike operators in many industries, systems engineers typically do not have access to the sophisticated simulation facilities that enable them to make mistakes and learn to administer application processes before they make changes to primary or secondary systems. Even though there are relatively few international competency requirements for operators (for instance in the air traffic management or energy distribution industries), there is even less consensus over competency criteria for systems engineers.

The remainder of this paper is organized as follows. Section two illustrates the need for research by three examples from the air traffic management domain. Section three suggests an approach to conducting useful research, and section four presents brief concluding remarks.

2. Illustrating The Need

This section motivates the need for research by analyzing three adverse events. Two of these examples represent

serious recent accidents in air traffic management in European air space. The third example did not result in any injuries or damage but helps to typify a class of failures that have relatively simple causes, but which are increasingly difficult to diagnose across multiple interacting sub-systems.

2.1 Überlingen

The Überlingen accident occurred on the 1 July 2002 when a Boeing 767-200 collided in mid-air with a Tupolov TU164M [3]. 71 people were killed. The identified direct causes of the accident centred on the Air Traffic Control Officer's (ATCO) instruction to the Tupolov crew, which contradicted the Aircraft Alert/Collision Avoidance System (ACAS), and ordered them to descend into the Boeing 767, which was also responding to an ACAS warning to avoid the other aircraft. In constrast, we focus on the maintenance procedures at Zurich Air Traffic Control Centre (ACC) that created preconditions where the ATCO was likely to make a mistake.

ACC Zurich upper airspace was divided both vertically and horizontally. The particular vertical division above flight level 235 into two or three sector operations created particular problems for the operation of Revised Vertical Separation Minima (RVSM). RVSM was a European initiative to increase capacity by relying on new generations of avionics to reduce the required vertical separation between aircraft. ACC Zurich staff, therefore, developed a six hour plan to modify the flight plan processing system to simplify the upper airspace and support the implementation of RVSM. This plan affected a number of different systems: the radar data application; the multi-radar computer system; the flight plan data processing system for tower and approach control; the landing sequence computer; the departures and arrivals traffic system and the ground to ground phone system with neighbouring centres. A further consequence of these effects was that management began to prepare for the upgrade by issuing official instructions to describe the work. An additional memorandum also documented the impact that the work would have in requiring controllers to work in fallback mode without a visual Short Term Conflict Alert (STCA).

ACC Zurich's normal configuration for night operations was based around two controllers supported by two assistants. It was also usual for one of the ATCOs to leave the control room and rest in the lounge as soon as the amount of traffic decreased. Management knew about this practice and there was no apparent pressure to stop it hence there was an assumption of at least implicit acceptance. Even though there were additional systems managers to support the RVSM upgrade, the Chief controller did not brief his colleagues about the additional staff. In consequence, a single controller was placed in a situation where they believed they were responsible for the tasks associated with radar planning, radar execution, shift supervisor, and systems manager at a time when profound changes were being made to the technical infrastructure.

The Zurich radar data processing system consisted of three main Thomson MV9800 computers. The first was used for primary operation, the second was held as a "hot standby", and the third was used for test purposes and software development. The system has a visual and acoustical STCA (Short Term Conflict Alert). If the connection between the MV9800 and the controller workstation system is interrupted, as it was on the night of the accident, then the correlated radar image is lost. Controllers must use the fallback radar computer (fbRDPS). This means that the controller must manually correlate radar targets with flight plans. The maintenance work that led to the loss of the MV9800-ICWS link also deprived the controller of the visual Short Term Conflict Alert, although an audible alarm was available. By forcing the manual correlation of radar targets and flight plans and by removing the prompt visual STCA warnings, the controller was placed in a vulnerable position. Hence, the interactions between staffing and the degraded technical infrastructure created the context in which the ATCO contradicted the ACAS advisory on the Tupolov.

2.2 Linate

On the 8 October 2001, an MD-87 collided with a Cessna 525-A that had taxied onto runway 36R at Milan's Linate Airport [1]. One hundred and fourteen people were killed on the aircraft along with four ground staff who were working in a baggage handling building, which was struck by the MD-87 after the runway collision. The official accident report concluded that the Cessna's crew had mistakenly crossed the active runway under low visibility conditions. Again, however, the engineering of ground based systems played a significant role in the causes of this accident.

In the past, the Linate air traffic control officers had been provided with an analogue Aerodrome Surface Movement Indicator (ASMI) radar system. Traffic increases had exposed the reliability and low definition of this system to a point at which ATM personnel began to look for an alternative. There was a plan to introduce a new Surface Movement Guidance and Control System (SMGCS) using video camera technology. The old AMSI system was, therefore, taken out of service three years before the accident.

Plans to install the new system were jeopardized when the predecessor of ENAC (Italian Civil Aviation Authority) objected to the antenna location. They argued that the planned location would involve additional expense by constructing a temporary structure that would then be moved once a new tower was built. It was also argued that the proposed structure might hinder visibility and that there were few reported problems in handling ground traffic at Linate. There was also concern that the new system would not harmonize with other European initiatives.

In July 2000, ENAV assumed responsibility for air traffic management operations at Linate. One side effect of this hand-over was that approval was finally granted for the development of the new SMGCS. The antenna was to be located in the same position as the previous ASMI radar.

However, the accident occurred while the upgrade project was further stalled. Mothballed hardware had to be re-serviced before the new system could be delivered. Further problems arose because the runway incursion sensors had already been deactivated on taxiway R6. In consequence the ANSV argued that there was "no possibility" to confirm the positions of the various aircraft using technical aids on the morning of the collision.

2.3 Failure of a Network Card

Linate and Überlingen illustrate the role that systems engineering played in Europe's two most serious ATM-related accidents. In contrast, the final case study considers an incident that did not result in any injuries or damage to aircraft. It illustrates the increasing challenge faced by infrastructure teams who must diagnose the causes of complex interactions between multiple systems.

This incident began when the backup mode was activated for the flight data processing system (FDPS) local area network at a major European airport. Aircraft already in the system were unaffected. However, some of the flight data information was not displayed for aircraft entering the system. The ANSPs engineers worked to restore system but could not determine the root cause of the malfunction. The airport continued to operate but with capacity restrictions. A multi-party team was formed among the engineers, the system suppliers and their sub-contractors.

A similar, intermittent malfunction occurred two days after the first problems. However, normal levels of operation were resumed after some suspected network components had been replaced and the system seemed to have stabilised. As before, engineers continued to work on identifying the root causes. This task was complicated because it proved to be very difficult to replicate the observed symptoms. After 28 days without any subsequent problems, the malfunction occurred once again. The subsequent report into the failure noted that "the determination of capacity was based on continuous risk assessments of the system performance and the technical and operational mitigations put in place to ensure safe operations". As before, capacity was gradually increased once the system seemed to have stabilised. Additional personnel joined the investigation team but they still faced significant problems identifying the causes of the problem given that the FDPS local area network was still in operational use. Monitoring systems were deployed and operational changes were introduced to ensure that aircraft were not in holding patterns during any potential future malfunction.

A subsequent malfunction again forced the system into a backup mode and flow restrictions were imposed. ATCOs were by now beginning to lose confidence in the systems infrastructure and the decision was taken to halt the use of the FDPS. This time when the system went into a backup mode, ATCOs observed further problems. Some working positions lost flight plan coupling on all aircraft while others only lost data for certain flights. The additional monitoring tools enabled the joint sub-contractor and ANSP engineering teams to diagnose an intermittent hardware problem associated with a network interface card. They were then able to replicate the fault on a standby platform. Capacity was gradually restored, together with ATCO confidence in the state of the underlying systems.

2.4 Common Themes

A number of common themes can be identified across the three incidents. In particular, it is clear that systems engineering played a significant role in the causes and contributory factors leading to failure. It is also clear that the complexity of safety-critical systems in Air Traffic Management are imposing significant demands on the skills and expertise of engineering teams as they work to configure and maintain multiple, interactive systems. The previous incidents also reiterate the communications challenges that remain to be addressed across the industry.

Communication with Operations: The Überlingen accident illustrates some of the hazards that can arise when operational staff do not understand the impact that infrastructure changes can have upon their everyday tasks. In this accident, the information was available but ATCOs failed to appreciate the significance of the documentation and management did not provide clear information about additional engineering support. Looking beyond this specific accident, a more general problem is that the systems engineers often provide information about their activities in a language that cannot easily be understood by systems engineers. Although this does not excuse situations in which ATCOs fail to read maintenance updates, there are often good reasons why operational staff do not follow dozens of pages of technical guidance at a level of detail that does not reflect their understanding of the underlying infrastructures.

Our previous work has revealed many other incidents that have occurred when systems and operational teams do not communicate the extent of maintenance activities during shift handovers. The third incident also illustrates how a series of malfunctions can undermine the confidence that operational staff have in their underlying systems. This can, in turn, create further communications problems where end users continually interrupt systems teams to gain further information about the diagnosis of a fault. Such concerns are understandable but they can erode the finite resources that are available to respond to infrastructure failures.

Communication with Sub-Contractors: The third case study illustrated the impact that sub-contractors can have upon the diagnosis of systems faults. The role of external, technical support is likely to become more and more important with the increasing sophistication of future infrastructures, such as those envisaged within the NextGen and SESAR programmes. Few service providers have the technical resources to develop, implement and maintain the many different automated and decision support tools that are

proposed within these initiatives. In Europe there are further plans under the Single European Skies initiative for neighbouring countries to cooperate in the deployment of increasingly complex, infrastructures. This will further complicate communications with, for instance, sub-contractors that were initially employed by another state. In the case study, described in previous paragraphs, there were strong communications between the various stakeholders. However, this is not always the case when, for instance, the best engineers are reallocated to subsequent projects once a system has been accepted by an end user. In these circumstances, sub-contractors can struggle to martial the necessary expertise to diagnose faults in the systems that they developed. Further problems occur when malfunctions arise from the interaction between systems that were developed and maintained by different external companies.

Communications with Management: All three case studies illustrate the communications problems that can arise between management and systems engineers. At Überlingen, there were precursor incidents under SMOP that were not acted upon. The subsequent BFU report argued that this was indicative of failures in safety management. Hence technical support staff could not easily anticipate the impact that their actions might have during this mode of operation. At Linate, higher levels of management did not appreciate the technical significance of the ground infrastructures that were gradually eroded over time. In the final incident, management were anxious to restore previous levels of service even though the joint engineering teams could not identify the root causes of the FDPS LAN malfunction.

A number of wider concerns help to explain the general communications barriers that often exist between systems teams and higher levels within their own organisation. In particular, there is an increasing trend to recruit Board level management from either operational or financial backgrounds. There are very few engineers at the top level of European Air Traffic Management organisations; similar caveats can be raised across the safety-critical industries. It should, therefore, not be surprising that in times of economic stringency we see safety budgets under increasing pressure with the consequent risk to infrastructure engineering.

This paper focuses on Air Traffic Management. However, similar concerns can be raised across many other industries. For example, previous studies of information technology failures in healthcare have revealed the same lack of understanding between systems engineers, clinicians, management and politicians [8].

3. A Roadmap for Safer Systems Engineering

Figure 1 provides a roadmap intended to support safer systems engineering. Higher levels of management as well as regulatory organisations often underestimate the contribution that infrastructure support makes to the safety of complex applications. In the past, engineers may only have had an indirect impact on safety because most applications only supported operational decision making. ATCOs retain the right to 'close the skies' or adjust the amount of traffic in response to concerns over underlying systems [9]. However, this position cannot be sustained. The next generation automated systems will stretch the ability of operational staff to directly intervene without decision support tools. Even today it can be difficult for operational staff to accurately judge the capacity restrictions that can be used to offset the risks from malfunctions in the underlying infrastructure.

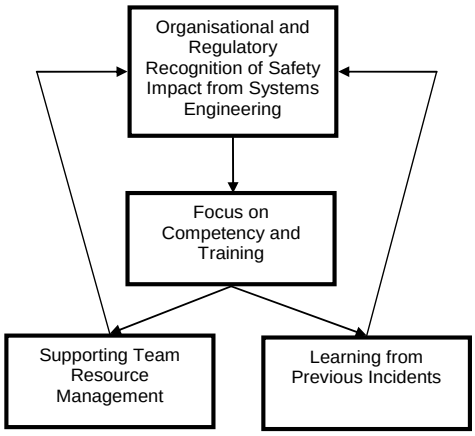

Figure 1: A Roadmap for Safer-Systems Engineering

The argument that management and regulators need to pay greater attention to the safety implications of infrastructure engineering is based around the previous analysis of our three case studies. As mentioned, there is a need to redress the under-representation of technical staff at higher-levels in many industries. If this is not done then there is little prospect of increasing safety margins beyond the absolute minimum necessary to meet regulatory requirements. This raises particular concerns given the lack of regulatory expertise in key technical areas, including human factors and software engineering.

A second waypoint in the roadmap is the need to identify competency and training requirements for systems engineering. These are, typically, far less developed than those available for operational staff [4]. For instance, there are no existing studies into the reliability of competency assessments for engineering teams that can be compared to those for operational staff [10]. There is considerable international disagreement over the training and skills that should be required of engineering teams. This disagreement is crystallised in the controversy over the future role of Air Traffic Safety Electronics Personnel (ATSEP). It is unclear whether increasing the levels of redundancy and automation in air traffic management, identified by the European SESAR and US NextGen programmes, will create a situation in which ATSEPs perform a purely technical function. In this view, their role includes minimal intervention in the diagnosis and correction of potential faults. Alternatively, the complexity of these integrated systems might create the need for more highly trained engineers with the skills and expertise to trace

complex interactions across multiple interacting applications. The three case studies presented in the opening sections of this paper clearly show the need for integrated support across diverse areas, including but not limited to hardware and software systems, human factors and risk assessment as they relate to technical infrastructures.

It is increasingly difficult for systems engineers to fully understand the impact that low level configuration changes might have at an operational level. In air traffic management, engineers are increasingly forced to focus on particular applications. Teams that support Flight Data Processing Systems may have only limited involvement in the engineering of Voice Communications Systems and vice versa. Specialist integration teams then help to support the interactions that arise at the interfaces between these different technical areas. This helps to promote a detailed knowledge of the underlying infrastructures but also creates barriers to engineers developing a global picture of the higher-level operational environment. It is time that regulators and management recognised the impact that these changes will have upon systems engineering; competency requirements must be carefully considered and documented if technical staff are to develop the skills necessary to support increasingly complex safety-related infrastructures.

The competency and training of systems staff also raises issues of selection. In some European states there is a tradition that individuals who fail the initial training for ATCOs are then directed towards a career as ATSEPs. This has a significant impact on the working culture when operational teams interact with support staff. It is also important to review the pay and conditions of systems engineers. In particular, a shortage of operational staff has led some states to increase wages well beyond those of the engineering teams. This undermines communications when two social groups develop. Further concerns also centre on the working conditions of systems and operations. The working hours of ATCOs are strictly regulated. However, similar constraints are not usually enforced for engineering teams. In consequence, technical staff can work for extended periods on critical infrastructures; fighting against levels of fatigue that undermine concentration and problem solving capabilities [7].

Figure 1 identifies two further landmarks in the roadmap that are strongly related to the competency of engineering staff, which support safety-critical systems. The first of these waypoints supports team resource management (TRM) training for systems engineering. This is particularly important given the communications barriers that previous paragraphs have identified between operations, management and technical staff. Significant advances have been made in helping promote cooperation between operational users – for instance through the use of simulation tools and training in the management of critical events. However, these techniques are not widely used to prepare engineers for the demands that are placed on them during infrastructure failures. In contrast, many technical staff learn 'on the job' even in safety-critical environments. It is normal practice to learn how to configure and maintain complex applications by setting up secondary or back-up systems that are only one-step away from live operations. This makes it difficult for engineers to deliberately simulate the faults that increasingly test the competence of infrastructure teams. It is, therefore, unsurprising that systems engineers often find it difficult to coordinate their response to adverse events with operational needs and the resources provided by their sub-contractors and systems integrators. In contrast, our approach argues for increased investment in training environments that help to develop team resource management skills in response to a range of simulated failures.

The final landmark in our roadmap focuses on the exchange of lessons learned from the systems engineering of safety-critical systems. Previous sections have used reports into the Linate and Überlingen accidents to identify the role that systems engineering has played in previous mishaps. These are rare documents; most infrastructure failures are seldom discussed beyond the organisations that suffer them. Political, economic and regulatory concerns limit the extent to which lessons are exchanged both within and between industries. This forms a strong contrast with the multiplicity of incident reporting sites that support the exchange of information about operational problems. The consequences of this can be seen in our third case study. There are similarities between the FDPS LAN failure and problems suffered by the FAA's Atlanta center during November 2009. Neither of the service providers knew of the problems suffered by their colleagues nor have there been any joint meetings to discuss the wider significance of these events for the safety of either SESAR or NextGen.

4. Conclusions

In the past, systems engineering had a limited impact upon the safety of Air Traffic Management. Operational staff had the ultimate responsibility for decision-making based on the information that was presented to them. Traffic flows were adjusted so that ATCOs could respond to potential failures; for instance by avoiding holding patterns and ultimately by 'closing the skies' during potential system failures. In the future this may not be possible. The complexity and integration of infrastructures within the NextGen and SESAR programmes will increase the operational reliance on computational tools. These innovations will significantly increase capacity. They will also increase the importance of technical staff in installing and maintaining the supporting systems that help preserve the safety of future operations.

This paper argues that by focussing on the end-users of complex applications, previous research has neglected the demands that are being placed on systems engineers. We have illustrated this argument by identifying the technical contribution to three case studies; the Linate and Überlingen accidents and a less publicised failure in an FDPS LAN. These incidents demonstrate a range of communications problems between systems engineers, operational staff and

more senior levels of management. This analysis has been used to develop a roadmap for safer-systems engineering based on increase attention by management and regulators, on improvements in competency and training for technical staff, on sustained support for team resource management including engineers and on the development of incident reporting systems for infrastructure failures.

Although our focus has been on Air Traffic Management, it is clear that many of our arguments can be applied to other safety-critical industries. In related work we have identified similarities with the role of software engineering teams in healthcare informatics. The increasing use of complex, programmable devices and the development of distributed electronic patient records has revealed the need for improved communication with clinical staff and with higher levels of management in order to improve patient safety.

References

[1] Agenzia Nazionale per la Sicurezza del Volo (ANSV), Milano Linate, ground collision between Boeing MD-87, registration SE-DMA and Cessna 525-A, registration D-IEVX, Reference A/1/04, (2004).

[2] L. Bainbridge, Ironies of Automation, *Automatica*, (19) 6:775-779 (1983).

[3] Bundesstelle für Flugunfalluntersuchung (BFU: German Federal Bureau of Aircraft Accidents Investigation), Accident on 1 July 2002, Near Überlingen/Lake Constance, Germany Involving Boeing B757-200 and Tupolev TU154M, Investigation Report AX001-1-2/02 (2004).

[4] EUROCONTROL, Guidelines for the Competence Assessment of Air Traffic Safety Electronics Personnel. http://www.eurocontrol.int/safety/gallery/content/public/library/L7.pdf (2006).

[5] C.J. Harris, C. G. Moore and M. Brown, *Intelligent Control: Aspects of Fuzzy Logic and Neural Networks*, World Scientific Press (1993).

[6] C.W. Johnson, *A Handbook of Accident and Incident Reporting*, Glasgow University Press (2003).

[7] C.W. Johnson, The Systemic Effects of Fatigue on Military Operations. In *2nd IET Systems Safety Conference*, The IET, Savoy Place, London, UK, (2007).

[8] C.W. Johnson, Identifying Common Problems in the Acquisition and Deployment of Large-Scale Software Projects in the US and UK Healthcare Systems, *Safety Science* (49)5:735-745, (2011).

[9] C.W. Johnson, B. Kirwan, T. Licu and P. Statsny, Recognition Primed Decision Making and the Organisational Response to Accidents: Überlingen and the Challenges of Safety Improvement in European Air Traffic Management, *Safety Science*, 47:853-872, (2009).

[10] E. Oprins, E. Burggraaff and H. van Weerdenburg, Reliability of Assessors' Competence Ratings in Air Traffic Control Training, *Human Factors and Aerospace Safety*, (6)4:305-321, (2006).

[11] F. Qi and B. Huang, Dynamic Bayesian Approach for Control Loop Diagnosis with Underlying Mode Dependency, *Ind. Eng. Chem. Res*. (49)18:8613–8623, (2010).

[12] N.B. Sarter, D. D. Woods, and C.E. Billings, Automation Surprises, in *Handbook of Human Factors & Ergonomics*, second edition, G. Salvendy (Ed.), Wiley (1997).

[13] C.D. Wickens, J.S. McCarley, A.L. Alexander, L.C. Thomas, M. Ambinder and S. Zheng, Attention-situation awareness (A-SA) model of pilot error. In D. C. Foyle & B. L. Hooey (Eds.), *Human performance modeling in aviation* (pp. 213-239). Taylor & Francis, Boca Raton, FL, (2008).