Agent-based Approach to Retrospective Analysis of Aviation Accidents

Msc. Thesis
Dino Hajdarbegovic
1329820
Supervisor: A. Sharpanskykh
Graduate professor: Henk Blom

Delft University of Technology

Agent-based Approach to Retrospective Analysis of Aviation Accidents

By

D.Hajdarbegovic

in partial fulfilment of the requirements for the degree of

Master of Science

in Aeospace Engineering

at the Delft University of Technology,

to be defended publicly on

xxx

Delft University of Technology

Preface

Što manje, dijete, znaš - to si manje nesretan.

Delft University of Technology

List of Abbreviations

ABM	Agent Based Modeling
ATC	Air Traffic Control
Atco	Air traffic controller
ATM	Air Traffic Management
FAA	Federal Aviation Administration
FRAM	Functional Resonance Accident Modeling
HFACS	Human Factors Analysis and Classification System
KLM	Koninklijke Luchtvaar Maatschappij
MAREA	Mathematical Approach towards Resilience Engineering in *ATM*
MASA	Multi Agents Situation Awareness
NTSB	National Transportation Safety Board
PanAm	Pan American Airways
RA	Resolution Advisory
SA	Situation Awareness
SAD	Situation Awareness Difference
STAMP	System-Theoretic Accident Model and Processes
TA	Traffic Advisory
TCAS	Traffic Collision Avoidance System
ThoCo	Thomas Cook

List of Figures

Contents

1. Introduction

According to a ICAO Safety Report (ICAO, 2013), 2012 was the safest year in aviation since 2004. Overall there is a trend of fewer and fewer mishaps in aviation each year making it one of the safest modes of transportation. However, the report also states that the sector is ever growing. This can potentially have adverse effects to the safety as crowded airways and air terminals, become a more and more complex system. Accident analysis can serve two purposes: assigning blame and preventing future occurrences. Obviously, the second one is more important from an engineering standpoint. Regardless of the purpose, accident analysis is an important part of aviation. This chapter will show that there is a lack of review techniques in retrospective accident analysis and will explain the goal of this thesis. This will be achieved by giving a brief theoretical background on retrospective accident analysis, agent based modelling (ABM), human error in aviation and situation awareness (SA) in 1.1, 1.2, 1.3 and 1.4 respectively. Then in 1.5 these concepts will be summarised and the need for a new methodology will be explained.

1.1. Traditional accident analysis techniques and some limitations

Leveson (2004) argues that the linear nature of traditional approaches to accident analysis, like fault and event trees, fall short in describing the root causes of accidents because of the limited notion of causality that linear models pose. She also observes that this kind of analysis is mostly fit for component failure, so a whole slew of mishaps that can be attributed to systemic failure remains uncharted by these traditional methods. Furthermore, she argues that the chosen link between consecutive events is highly subjective and she calls for new approaches. Qureshi wrote a dissertation on accident modelling techniques for complex socio-technical systems (Qureshi, 2007). He argues that the complexity of modern socio-technical systems poses a great challenge for accident investigation as a multidisciplinary approach is required. In order to capture the complexity of said systems, researchers from engineering, social science and cognitive psychology need to collaborate in developing models that can analyse systemic failures. He concludes that approaches that model systemic failures and take emergence into account are more fit for mishap analysis. Similarly Johnson and Holloway (2003) note that formal techniques might be necessary to reduce bias that affects the interpretation of mishaps. They describe the most notable formal logics and how they can be used to prove whether certain events created the necessary causes for the mishap. The examples they mention, *classical logic*, *Bayesian logic*, *why-because analysis* etc. seem very useful for analyzing causal links between events. They are however not useful for formalizing the events themselves. A relatively new paradigm in modelling of complex systems, and one that Johnson and Holloway missed, is Agent-based modelling (ABM). ABM seems to posses all the properties that Leveson, Qureshi and Johnson and Holloway deem necessary for analysing accidents: ABM can formally model emergent systemic failures of complex systems.

1.2. Agent-based modelling

Agent-based modelling is a way to model complex systems from a microscopic view. Instead of modelling the entire system, ABM simulates individual entities that constitute the system. The approach has its roots in Distributed Artificial Intelligence which sought to solve problems by dividing knowledge between separate constituents who all possess different expertise (Gilbert & Terna, 2000). These constituents are called agents and are traditionally defined to be autonomous, reactive, social and goal driven (Janssen, 2005). The principal thought behind the approach is that the agents are easier to simulate than the entire system. By assigning simple behavioural rules to the agents, complex behaviour of the entire system could emerge. One of the earliest examples using this principle is perhaps Reynolds (1987) flock simulation. Reynolds managed to write an algorithm that successfully mimicked flocking behaviour of birds by programming a single particle with simple local rules like ‘stick together’ and ‘avoid collision’. A few years later, while ABM was fast gaining traction within social studies, Nagel et al. (1999) in association with the Los Alamos National Laboratory developed TRANSSIMS which effectively showed that agent-based simulation approach made it possible to simulate transportation issues like congestion. A survey of 279 published articles between 1998 and 2008 by Heath, Hill and Ciarallo (2009) shows how broadly the approach has been adopted: researches fields ranging from biology and social studies to engineering and economics have found a way to utilise the paradigm to gain an understanding of complex systems where emergent behaviour is key. Burmeister, Haddadi and Matylis (1997) note that traffic and transportation are well suited for ABM approach because of the geographic distribution of the subsystems and their highly dynamic nature.

Looking at aviation specifically, Xie, Shortle and Donohue (2004) used ABM to analyze the relationship between airport arrival capacity, delay and safety. With the model they could test the effects of increasing capacity on safety. They conclude that synchronizing arrival times is a potential way to increase capacity without sacrificing safety. Stroeve, Blom and Bakker (2009) also performed an agent based analysis of runway safety. They implement a Monte-Carlo scheme to introduce variability in the agents interactions. Pritchett et al. (2002) use agent-based modelling to capture the emergent behaviour of installing a clear air turbulence detector on aircraft. The model is quite encompassing simulating flight crew behaviour as well as air traffic controller behaviour and how these agents react to in incoming warning from the system at varying range. Similarly, Harper et al. (2002) have used agent based modelling to examine the effects of sensor accuracy on performance in a free flight scenario. They lay the groundwork for examining pilot errors more closely, more specifically, errors in maintaining situation awareness (more on this later on). Finally the MAREA-project warrants a mention. The objective of MAREA, or Mathematical Approach towards Resilience Engineering in ATM, is to develop a mathematical modelling and analysis approach for prospective analysis of resilience in Air Traffic Management (Stroeve, Everdij & Blom, 2011). For this an agent-based modelling approach was chosen and several model constructs have been identified that can together model a broad range of hazards in air traffic management (Bosse, Blom, Stroeve & Sharpanskykh, 2013).

The above mentioned papers shows that agent-based simulation is fit for a broad range of research topics and that it can be and has been implemented in modelling of complex socio-technical systems like transportation, and more specifically aviation and aviation safety. However little works seems to have been done in retrospective accident analysis. The need for an unbiased methodology is there, and agent-based modelling seems to fit the needs. Further research in the viability of agent-based modelling as a tool for retrospective accident analysis is an interesting topic well worth exploring.

1.3. Human error and event investigation

So why is it that linear event chain models are not sufficient anymore for retrospective accident analysis? Why is there a need for a systemic approach like ABM? The answer lies in the increased reliability of aircraft components. Whereas chain event models are good for analysing component failure they fall short when analysing system failures, and as technical reliability of aircraft increases the percentage of accidents that cannot be attributed to component failure is becoming less and less prominent. It is estimated that between 70% and 80% of all accidents in aviation can be attributed to human error (O'Hare et al. 1994). In their research O'Hare et al. (1994) more or less confirm this estimate. Their study is a twofold examination in human cognitive failure. The authors analyse a database of aviation mishaps occurring in New Zealand. In the first part they attribute to each entry in the database one of the error stages proposed by Nagel (1988). Nagel suggests that most human errors in the cockpit can be covered by three error stages, namely: information (acquiring, processing and communicating of information), decision (determining the best course of action) and action (execution of the chosen action). The results propose that decisional factors play the biggest role in fatal accidents. In the second part of the paper, the authors make use of an more expanded error taxonomy proposed by Rasmussen (1982). Rasmussen's taxonomy classifies human errors in six stages: information, diagnostic, goal, strategy and procedure. Using this taxonomy on the mishaps in the New Zealand database, the authors could correlate the error mode to the severity of the mishaps and to pilot characteristics. The results are quite complex for introduction and the interested reader is referred to the paper itself. What is important to mention here however is that a model for human error could be adopted to be used for aviation mishap investigation with the help of an accident database. O'Hare et al. tested this with two taxonomies of which they preferred the one proposed by Rasmussen. There have been several other frameworks proposed through the years. Kirwan (1998) has written a very thorough review on some of them. One human error framework that Kirwan does not mention is the Human Factors Analysis and Classification System (or HFACS), arguably because it was just rounding off development at the time Kirwan's dissertation was written (Shappell and Wiegmann, 1997). What makes HFACS important enough to merit some extra attention here is the fact that it was developed specifically with aviation in mind by the US navy (HFACS, 2009)

The HFACS taxonomy is based on Reason's concept of latent and active failures (Reason, 1990, referenced in Shappell and Wiegmann, 1997) and was developed throughout the mid nineties.

It positions human error into four possible categories:

1: Unsafe acts of operators: This includes *errors* and *violations* where the former can be considered unintentional and the latter is a wilful disregard for rules and regulations. Errors fall within three categories, namely decision errors skill based errors and perceptual errors. *Decision errors* can arise from misinterpretation or misuse of relevant information. *Skill based errors* includes lapses in observation and memory and forgotten intentions (Shapell & Wiegmann, 2001). Finally *perception errors* arise when sensory input is degraded by environmental conditions (bad weather, night time etc)

2: Preconditions for unsafe acts: These fall into three subdivisions. *Substandard conditions of the operator* includes adverse mental states, adverse physiological states and physical and mental limitations. The second subdivision, *personnel factors*, is comprised of crew resource mismanagement, for example failure to communicate, and personal readiness which accounts for factors that reduce the operating capabilities of the crew. Lastly there are the e*nvironmental factors* which include both the technological environment that deals with design of the equipment and the effectiveness of prescribed procedures, as well as the physical environment which includes the effect of the operational setting and the ambient environment.

3: Unsafe supervision: the third failure category accounts for the latent failures in Reason's model. Four subcategories are included here. *Inadequate supervision* refers to failure in training en leadership. *Planned inappropriate operations* deals with aspects of improper crew scheduling which affects the crews performance adversely. The third subcategory *failure to correct known problems* accounts for the cases where unsafe practices, which are not necessarily against regulations, are not corrected. The last subcategory, *supervisory violations*, on the other hand is failure to correct behaviour that is against regulations.

4: Organisational influences: Just like the previous category, the fourth one deals with latent failures. It accounts for fallible decision of upper level management that have a direct influence on supervisory practices (Shapell & Wiegmann, 2001).

The greatest difference between HFACS and the taxonomy adapted by O'Hare et al. (1994) from Rasmussen's model (1982) is the last two categories. Latent errors are not accounted for in the model used by O'Hare et al. Using the HFACS taxonomy Shappell et al. (2007) did a similar study to the one performed by O'Hare et al. While O'Hare et al. employ some tools that test the robustness of their method it falls short because of the need for experts. The classification of the errors is performed by the authors themselves, who do admittedly posses pilot licences. The study by Shappell et al. (2007) attempts to increase the robustness by employing the help of six pilots and two experts in human factors to code the accidents in the database. The database used is provided by the FAA and contains 1020 accidents between 1990 and 2002. The results show that a majority of the human factors can be attributed to *unsafe acts of operators* and *preconditions to unsafe acts*. In the *unsafe acts* category 56.5% of the accidents could be attributed to skill-based errors. In the *preconditions to unsafe acts* category 58% can be attributed to the physical environment. It is quite noticeable that the latent failures do not appear to have a significant contribution. This begs the question if the coding of the accidents by pilots is an adequate method because it is hard for a pilot to have a complete overview of the organization in question. This also raised the question whether the database is comprehensive enough to contain clues about potential organisational failures.

1.4. Situation awareness

One of the constructs that was identified in the MAREA project to cover all the hazards in air traffic management, is multi agent situation awareness (Bosse, Blom, Stroeve & Sharpanskykh, 2013). This means that there is already a formal description available for situation awareness in a multi agent environment. This will be explained later, but first situation awareness itself will be described.

Situation awareness (SA) as a construct has been around since World War I (Press, 1986, referenced in Endsley, 1995). When looking at the two models for human error discussed in 1.3, several elements of the SA construct can be recognized throughout the narrative even though it is not explicitly mentioned. In the HFACS taxonomy most errors result in loss or degradation of SA. In fact, it can be argued that all but *violations* can have a direct impact on the situation awareness of the operators. This gives credence to the assumption that a large part of accidents are in fact caused by a difference in SA among operators.

Several models exist to describe SA, however there is no consensus within the human factors community on which describes reality best. Stanton et al. (2010) mention that all models can be divided into three different schools of thought. The difference being in the main discipline used to describe it: psychology, engineering and human factors / system engineering. Those who view the subject from a psychology perspective believe that SA resides completely within the mind of the operator and that the complete SA can be considered a summation of the SA of the individual operators (Endsley, 1995). Other believe that that SA falls within the purview of engineering and that it resided within the world. For example, all the information that is present in the cockpit displays and dials is part of a pilots SA. The third school of thought, according to Stanton et al. is viewing SA from a system engineering side where it is defined as an emergent property that arises from interactions between the operators and the environment. As such it is something that resides in between agents and objects, a distributed cognition as it were.

Salmon et al. (2008) gives a comprehensive write-up of the most important frameworks. The authors note that the definition by Endsley (1995) is the most used. This model falls within the psychological school of thought.

Endsley's model describes situation awareness as a three levelled cognitive construct that enables a participant in a system to "know what is going on". Her exact definition is: *"The perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future" Endsley (1995)*. The three levelled nature of the construct can be recognized in this definition. The model shall be explained here briefly.

Level 1 SA: Perception of the Elements in the Environment. This is the first and most basic level of SA and it deals solely with the perception of relevant information in the environment. This includes direct audiovisual information, e.g. the presence of another aircraft in front of you on the runway within the pilot's sight or the communication from the runway operator notifying you of the other aircraft. It also includes information gathered from instruments, like the pilot's own velocity, or warning signals in the cockpit. Note that at this level of SA no processing of information occurs. To continue with the pilot example: with only level 1 SA the pilot does not comprehend the consequences of seeing another aircraft in front of his own aircraft. The pilot simply perceives another aircraft to be there without attributing any meaning to it

Level 2 SA: Comprehension of the current situation. At this level the information perceived during the acquisition of Level 1 SA is processed and its meaning is understood. To be more exact the effect of the perceived information on the pertinent goals of the operator is understood. The operator gains a holistic view of the environment. However, the view is a frozen frame of the current situation. In understanding what the information means for the near future, Level 3 SA comes into play.

Level 3 SA: Projection of Future Status. Both Level 1 and Level 2 SA are a prerequisite to have an understanding of the future state of the system. When Level 3 SA is acquired the operator knows of the future states and goals of other elements within the system as well as of himself.

Endsley (1995) goes on to describe some internal mechanisms that are used in the formation of SA and human factors that limit the formation. For example, she notes that limited attention capacity can impact Level 1 SA and working memory capacity impacts Level 2 SA. This can be offset by experience and training. Experience and training provide the operator with frameworks necessary to frame the information from level 1 SA in the correct context, to understand its meaning. She compares these frameworks with mental models as defined by Rouse and Morris (1985, referenced in Endsley 1995) which are essentially internal models of reality. Besides the human factors, Endsley also provides examples of system design factors that can influence the formation of SA, e.g. cockpit layout. She finalises the paper with a dissertation on how errors in SA can occur across the three levels. It is quite clear from this description how Stanton et al. (2010) view this model as psychological. According to this model SA is an internally held product of an individual. It is dependent on several cognitive processes like mental models, scripts, automation, working memory etc. (Endsley, 1995).

It is easy to understand why Endsley's model is widely used. It is very intuitive and easy to understand. Furthermore, the structure in three levels allows for easy implementation when trying to apply the model. Salmon et al. (2008) note however, that there are several points of critique to be found throughout the literature, like the fact that the construct is based on several unproven psychological constructs like mental models. Another point of critique is the fact that in this model, SA appears to be static and finite. Bedney and Meister (1999) have a similar three levelled approach, but their model seems to be of a more dynamic nature. It is based on the theory of activity which states that the goals of individuals in a system are the "ideal situation". Differences between this ideal situation and the perception of the current state of the system is what drives an individual to action(Bedney and Meister 1999). Activity consists of three stages. Firstly, there is the *orientational stage* where the individual gains perception and understanding of the current state. Then there is the *executive stage* where the individual decides on what actions need to be taken to get closer to the ideal and where the actions are executed. Finally the evaluative stage assesses the effects of the performed actions. Bedney and Meister use these stages to form a construct of various subsystems that influence each other to eventually form SA. Figure 1, taken from Salmon (2008), represents the model.

In the model information is processed in block 1. The framing of the information in the right context is influenced by the goal (block 2), the conceptual model of the current state of the system (block 8) and past experience (block 7). This seems similar to the concept of mental models in Endsley's framework which influence Level 2 SA. Endsley notes that the mental models themselves are influenced by training and experience (Endsley, 1995).

Moving on, the figure shows that the interpretation of the input also modifies the goal of the operator which is also influenced by the actions (block 5). So the main difference between this model and Endsley's approach is that this model deals more with the process of acquiring SA. All the subsystems in the Bedney and Meister model seem to continuously monitor reality, mix it with their past experiences (i.e. mental models) and adjust the SA accordingly to better cope with what they are perceiving. This is a far more dynamic approach than the single snapshot approach of Endsley.

According to Stanton et al. (2001) block 3 and block 8 are the most important functions where a faulty SA can arise from. Either the operator errs in deciding what the relevant tasks and conditions are or he does not possess the correct conceptual model to deal with the situation, for example because of lack of training or experience.

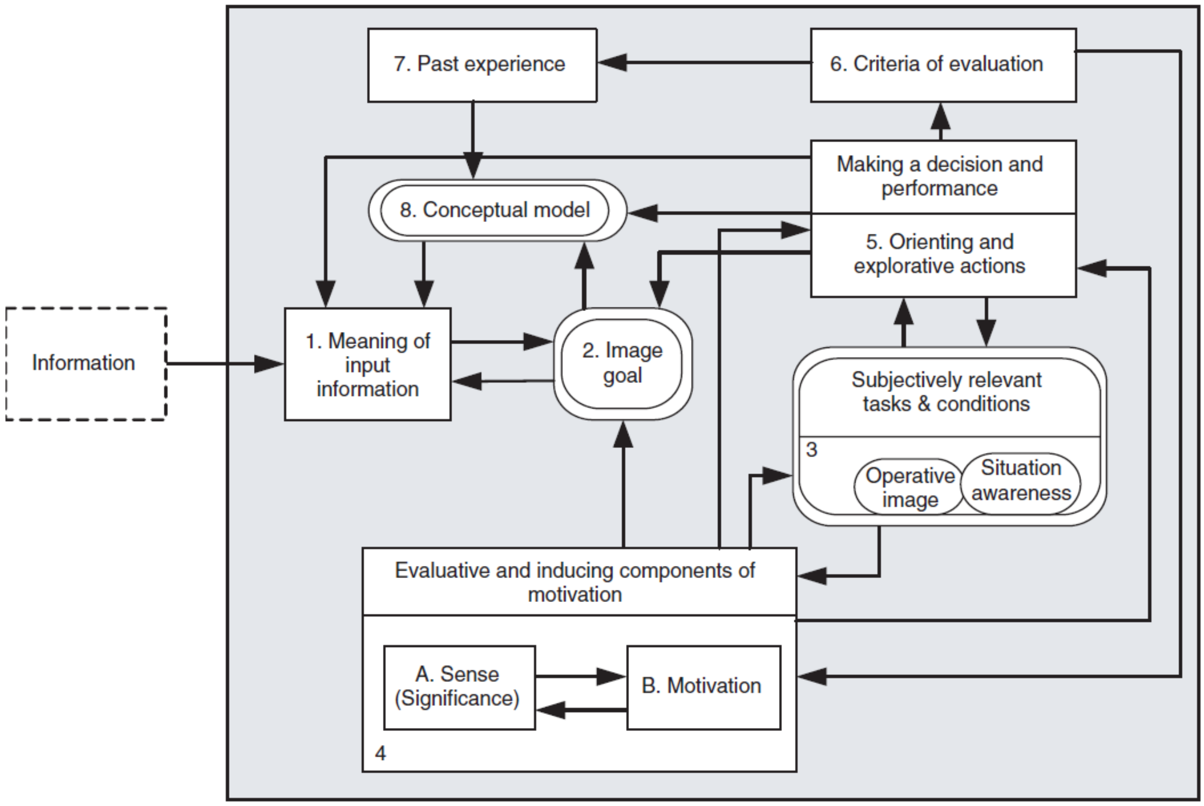

Figure 1: Sub-system approach to SA (Salmon et al. 2008)

Another different approach is proposed by Smith and Hancock (1995). This framework falls within the system engineering school of thought where SA is an emergent property arising from interactions between the agents that comprise the system. According to the model, SA is an "adaptive externally-directed consciousness". What this means is that Smith and Hancock define SA as both a product and a process. SA stems from both internally held models and external cues that direct the agents attention to a particular element. The perception of these elements in turn modify the internal models in an endless cycle similar to the dynamic approach of Bedney and Meister. The authors elevate SA from mere passive observation to active seeking of knowledge. They argue that the "picture" an agent holds of the system is not SA, rather it is formed by his SA.

The notion of SA existing outside of the operators as proposed by Smith and Hancock can also be found in the concept of Distributed Situation Awareness (DSA). Stanton et al. (2006) propose that situation awareness is a “dynamic and collaborative process” between agents on a moment to moment basis. The implications are that agents within the system don’t hold the same SA although they are compatible with each other. The knowledge required for SA is distributed among the agents of the socio-technical system. Here “knowledge” represents what the agents need to know in order to complete their goal. In other words, the agents, both human and non human, hold a piece of relevant information. Different agents hold different views on the situation, so their SA can differ as well. However the knowledge of the agents can overlap and complement each other depending on their goal. Stanton and colleagues (2006) suggest that the degraded SA of one agent can be updated by the other (e.g: verbal communication between two human agents, or the reading of an instrument). The authors note that it is important for the agents in the system to be aware of which agent holds what information in order to know which transactions of SA need to take place (meta SA).

This approach does not disqualify individual SA as proposed by Endsley. In fact it can be argued that the two models complement each other (Salmon et al., 2008). For example, in a takeoff scenario one agent is the radar that acquires Level 1 SA and relays the information to the Air Traffic Controller which proceeds to interpret the information (Level 2 SA) and project on the future state of the system (Level 3 SA). Rather, the distributed SA can be seen as an addition to individual SA, a system wide cognition that cannot be accounted for with individual cognitive elements (Stanton et al., 2006). In this authors opinion the combination of DSA with a model that defines individual SA, like Endsely’s model, is the most comprehensive way to look at SA in a cooperative environment. Specifically the concept of meta SA, the knowledge where to look for information to update your SA, is quite intriguing and warrants further research.

The above shows that there is no clear consensus on how SA should be perceived. Endsley’s approach seems to be the most elegant one and consequently it has drawn the most attention. All approaches have in common is that they lack any kind of formalism. At the beginning of this paragraph it was briefly mentioned that there is already a formal model for situation awareness that was used in the MAREA-project (Bosse, Blom, Stroeve & Sharpanskykh, 2013). This method will now be explained.

Stroeve, Blom and Van der Park (2003) propose a mathematical representation of multi agent situation awareness (MASA). This view is a lot like the view on Distributed Situation Awareness where all agents have their own SA about other agents. The representation of SA for any individual agent *k* about agent *j* at time *t* becomes a column of four sub processes:

$$\sigma_{t,k}^{j} = \begin{pmatrix} \hat{\iota}_{t,k}^{j} \\ \hat{x}_{t,k}^{j} \\ \hat{\theta}_{t,k}^{j} \\ \hat{\upsilon}_{t,k}^{j} \end{pmatrix}$$

with $j = 1, \dots, n$ where n denotes the number of agents. Furthermore:

- $\hat{\iota}_{t,k}^{j}$ denotes the awareness by agent *k* at time *t* of the identity of agent *j*.
- $\hat{x}_{t,k}^{j}$ denotes the awareness by agent *k* at time *t* of continuous-valued state components of agent *j*.
- $\hat{\theta}_{t,k}^{j}$ denotes the awareness by agent *k* at time *t* of discrete-valued state components (modes) of agent *j*.
- $\hat{\upsilon}_{t,k}^{j}$ denotes the awareness by agent *k* at time *t* of the intent of agents *j*.

The intent SA of agent k about agent j at time t can be represented as:

$$\upsilon_{t,k}^{j} = \begin{pmatrix} \bar{\theta}_{t,k}^{j} \\ \bar{x}_{t,k}^{j} \\ \bar{t}_{t,k}^{j} \end{pmatrix}$$

Here

- $\bar{\theta}_{t,k}^{j}$ denotes the mode of agent *j* that is anticipated by agent *k.*
- $\bar{x}_{t,k}^{j}$ denotes a continuous state of agent *j* as anticipated by agent *k.*
- $\bar{t}_{t,k}^{j}$ denotes the expectation of agent *k* about the time for which the continuous state or discrete state of agent *j* will be attained.

Updating each individual agents SA is assumed to occur through 3 different processes:

1: Observation

2: Communication

3: Reasoning

The *observation* process is a mapping function such that

$$\sigma_{t,k} = f_k^{obs}(\sigma_{t^-,k}, i_t, x_t, \theta_t, \varepsilon_{t,k})$$

Here i_t, x_t, and θ_t represent the actual identity, actual continues state component and actual discrete state component respectively. Furthermore t^- is the time just before the update occurs. Lastly there is the variable $\varepsilon_{t,k}$ which could for example model perception error.

The *communication* process is a mapping function such that

$$\sigma_{t,k} = f_{k1,k2}^{com}(\sigma_{t^-,k1}, \sigma_{t^-,k2}, \varepsilon_{t,k1,k2})$$

It depends on the previous SA of the agents that are exchanging information and again a variable that can represent noise or error.

Finally the *reasoning* process is a mapping function such that

$$\sigma_{t,k} = f_k^{rea}(\sigma_{t^-,k}, \varepsilon_{t,k})$$

It only depends on the previous SA of the agent in question and again a variable.

1.5. Concluding remarks and goal of this thesis

As the previous paragraphs have shown there seems to be a lack of systemic techniques in retrospective accident analysis. The need for new analysis tools stems mostly from the fact that improved reliability of components has rendered traditional analysis methods less applicable. It is estimated that the majority of accidents can be attributed to “human error” and as was shown, most of the categories in human error can be rephrased as errors in situation awareness. Therefore it can be assumed that the majority of aircraft incidents and accidents where there was no component failure can be attributed to a difference in SA among the players partaking in the scenario. Several models for Situation Awareness already exist of which Endsley’s model (1995) is most widely used. The concept of Situation Awareness can be incorporated into an agent based model as shown by Stroeve, Blom and Van der Park (2003) and it was shown that agent-based modelling and simulation is fit for a broad range of research topics and that it can be and has been implemented in modelling of complex socio-technical systems like transportation, and more specifically aviation and aviation safety. Right now the MAREA project seems to be leading the way in the use of agent based modelling for aviation. However little works seems to have been done in retrospective agent-based analysis in aviation safety. This thesis has the goal to find out how multi agent situation awareness can be used for retrospective accident analysis, so:

To create a structured, formal approach for retrospective modelling and analysis of safety occurrences by combining the concepts of Situation Awareness and agent-based modelling.

The research questions that need to be answered to attain this goal are as follows:

Q1: *How can a formal agent-based approach be used for systematic retrospective modeling and analysis of incidents and accidents?* And how can these findings be generalized?

Q2: *What does the agent-based approach add to the existing investigation reports and how does it compare to other existing approaches for retrospective modeling and analysis of safety occurrences?*

First the methodology will be outlined in chapter 2. Then these research questions will be answered along the way as three distinct accidents are analysed in Chapter 3, 4 and 5. The report starts off with the Los Rodeos accident of 1977 in 3 followed by another incident in Tenerife, namely the 2011 near miss. This analysis is found in chapter 4. The last case is the highly scrutinized attempted departure from a wrong runway in Lexington Kentucky. This case will be explored in chapter 5. Chapter 6 will show a comparison between MASA and two other retrospective analysis methods.

Finally chapter 7 will give the conclusions and recommendations.

2. Methodology for Retrospective MASA analysis

The agent based modelling approach from Stroeve, Blom and Van der Park (2003) will be used as a starting point for modelling three distinct cases in chapters 3, 4 and 5. This chapter will explain the methodology for the analysis and the steps that need to be taken. These steps do not necessarily have to follow up on one another. In fact, like many other engineering solutions, the retrospective MASA analysis should be viewed as an iterative process where previous steps are revisited as new insights are gained. The steps are:

- The orientation phase
  1: Identify likely SA differences
  2: Identify likely root causes of the SA differences
  3: Identify likely consequences of the SA differences
  4: Identify relevant agents in the scenario
  5: Identify likely non-nominal SA update types in the scenario
  6: Identify likely sub-models in the scenario
- Formalizing the SA differences and their progression
  1: Identification of relevant states and the ontology
  2: Identification of SA update cues
  3: Mapping of the intent vector
  4: Analysis of the SA updates at each cue
- Mapping of the propagation of the SA differences

These steps will now be explained.

2.1. Orientation phase

To get a feel for agent-based modelling, professor Blom (2012) introduces his students at the TU Delft University of Aerospace Engineering to a two step method whereby first possible SA differences between the agents are identified as well as their root causes and consequences. Secondly the method requires the identification of the relevant agents, non-nominal SA update types, and the underlying sub models, i.e. the reason why the SA update was non-nominal.

It serves as a good starting point for the entire analysis as it helps to roughly paint the overall picture of the accident or incident without getting lost in the details. The steps will be expanded upon below.

Identify likely SA differences

In this step possible SA differences are sought out. The question that needs answering is: which perceptions of the system by which agents seem to differ from the perception of another agent? A pilot might for example believe that there are no strong crosswinds and get caught by surprise by a sudden gust as he approaches the landing strip. Or an air traffic controller that believes an aircraft is on a certain taxiway whereas it is somewhere else entirely.

Identify likely root causes of the SA differences

Where, when and how did the previously found SA differences occur? What information was not available to the agents such that the difference occurred? In the crosswind example above, the root cause might be the absence of the information in the NOTAM, or the failure of the tower to communicate the current weather to the pilot. In the example of the aircraft that was not on the anticipated taxiway, the root cause could be that ATC did not observe the aircraft to take a wrong turn, and that he thus assumed nominal performance.

Identify likely consequences of the SA differences

At first sight this step might seem quite obvious: the most direct consequence of any difference in SA is that the agent with the difference has a view of the system that does not adhere to reality. So the pilot is not prepared and gets caught off guard by the sudden gust of wind, and ATC believes the aircraft to be somewhere else than it actually is. However it is also important to muse how the SA difference might have propagated through the system, i.e. what additional differences in SA were caused by the first difference. The air traffic controller in the example might believe that the runway is free from traffic while that is not the case.

Identify relevant agents in the scenario.

Janssen (2005) defined agents as constituents of a system that are autonomous, reactive, social, and goal driven. In the context of retrospective analysis of SA however, it is decided to create a broader definition of agents. Agents are:

All elements of a system that in some way can influence the performance of other elements by virtue of their current state.

This means that in a taxi or takeoff scenario the pilot of the aircraft is considered and agent as well as the air traffic controller. But also the runway: the runway can influence other elements, e.g. the pilot, by being in a state that is detrimental or positive to his performance. The runway can have the state *wet* or *slippery* with which it influences the takeoff procedure. The aircraft itself is also an agent by this definition, just as is other elements of the environment like the runway lights, the glide slope system, TCAS system, and even the weather. It should also be considered whether there is group of agents that can be considered homogenous by other agents within the system. For example, the pilot and the co-pilot could be considered a single agent by ATC because the controller has no way of discerning the individual agents.

It is likely that some agents selected in this early stage, turn out to be irrelevant, that relevant agents have been left out, or that a group is not homogenous after all. It underscores the importance of iterating and applying changes as new insights are gained.

Identify likely non-nominal SA update types in the scenario

As explained, Stroeve, Blom and Van der Park (2003) propose three ways the SA vector can get updated: observation, communication and reasoning. In this step the moments are observed where SA updates could have occurred but did not occur correctly or not at all. The tower failed to communicate the presence of gusts, and ATC failed to observe the aircraft taking a wrong turn.

Identify likely sub-models in the scenario

In this step the answer is sought to *why* the flow of information got impeded. These could be human errors like fatigue that have influence on the perception of elements in the environment, or technical issues like the radio having interference.

2.2. Formalizing SA differences and their progression

After the orientation phase it is time to formalize the SA differences. Formalizing helps to structure the case and can give important insights that were missed earlier. So once again it is important to note that the process is iterative.

Identify the relevant states and create the ontology

As seen in chapter 1, the SA vector can consist of any number of discrete or continuous states. From the orientation phase some clues can be gathered about what states are relevant for the case. Those elements that were perceived differently, i.e. the identified SA differences, need to be taken into account.

Returning again to our first examples the relevant states for the agent *weather* can be *windy* and *calm*. In the second example the agent *aircraft* could have the states *location A* and *location B.* The second example shows that the same state can be attributed to different agents: both the aircraft as well as the pilot can have the state *location A*. In these cases it is recommended to use the agent that represents the highest level of granularity.

Next for all agents and states a global ontology needs to be created. The different agents need an i.d and the states need an unambiguous representation.

Identify the SA update cues

In this step the moments are identified where critical SA updates *might* have occurred. Emphasis is on might because the availability of a cue does not necessarily mean the correct perception and processing as represented by ε.

In a perfect world all agents that comprise a socio-technical system have unlimited attention and can perceive and understand everything in the environment. Sadly attention capacity is limited in human agents as well as to some extent in non human agents. Managing ones consciousness to attend the relevant parts of an environment is thus important part of maintaining SA. Traditionally there is a dichotomy of how control of attention is handled in human beings. A distinction is made between endogenous and exogenous attention. Endogenous attention, as the name implies, originates internally and is determined by the goals of the agent. Exogenous control can be viewed as a stimulus from the environment that draws the attention of the agent. If this theory is tied to the SA construct, it implies that formation of SA is partly automatic and partly a conscious effort by the agent. The following 6 cues that might trigger an SA update are hypothesised:

Exogenous cues

- Exchange of information initiated by an agent other than the one whose SA is (about to be) updated. An example of this cue is an air traffic controller contacting the pilot flying and notifying him about the presence of another aircraft on the runway. The SA of the pilot is then likely updated via the communication function. However, this cue can also instigate the observation update function. For example a request for information by another agent might prompt the pilot to look around him so that he can gather information which he can relay back.

- Visual cues trigger the observation update function. For example: a sign with the runway number that suddenly comes into view will trigger an update, or confirmation, of the position of the aircraft. In retrospective modeling it can easily be determined if the cue was in view of the agent. The question remains however whether it was processed or not. This can be assumed from future actions. For example, if an agent makes a sudden stop after passing the cue or if he comments on it to another agent it can be safely assumed that it was processed.

- It is also proposed that there is a protocol enforced update cue. What this means is that by following a protocol the agent is forced to take note of his situation. For example, checking of flaps is part of the before takeoff procedure. Because it is mandatory it forces the pilots to actively update their SA.

Endogenous cues:

- Lack of SA, or more importantly, the realization of lack of SA by an agent is the third trigger. If a pilot suddenly realizes he is not sure whether he has take-off clearance he might contact the air traffic controller to have an update on his status. This one is by far the most difficult to infer in hindsight. Sudden change in the agent's behavior, without having been in visual contact with a cue and without outside communication would lead to believe that an update cue of this type was invoked.

- The attaining of the expected time of an intent vector is also considered a cue. If an agent assumes that he will reach a certain state in approximately five minutes, the realization of the passage of this time will prompt the agent to update his current state. The intent vector collapses so to speak unto the state SA. On the other hand, attaining an expected time but not observing the collapse of the intent vector to the state SA can be a cue as well. For example: a pilot expects to reach a hold short line for a particular runway in 5 minutes. If this time elapses but the hold short line is still not reached, the pilot might realize he took a wrong turn somewhere.

- Degradation is a form of passive and unwanted SA update. It is simply caused by human limits in retaining information. It is not really a cue as it is a process, and the exact moment it occurs is impossible to infer in hindsight. It can however trigger the update process, and some other cue can hint that somewhere in the process the agent's SA degraded.

All but the last of these cues triggers an update process through one of the functions detailed in chapter 1. Of course the cues can also have an additive effect increasing the certainty of an update. A visual cue that by itself would not be enough to ascertain a location could be a cue in conjunction with attaining the expected time of the intent vector.

Figure 3 shows how the two different types of cues interact with the system. The figure shows two agents, A1 and A2. Both have their own set of endogenous cues. They reside inside a world made up of exogenous cues and they themselves are exogenous cues as well. In this way an update prompted by an endogenous cue in the SA of A1 can propagate to update the SA of A2.

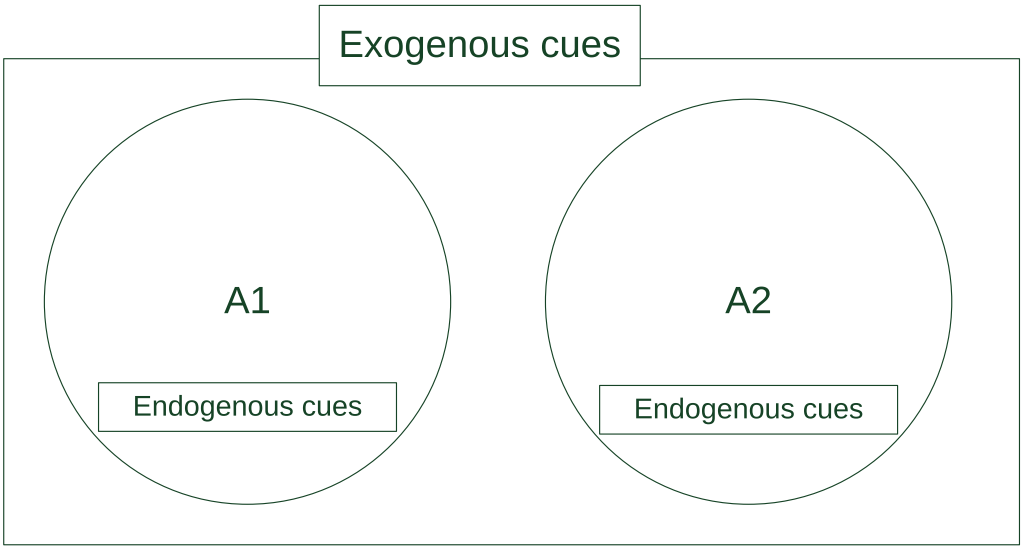

Figure 2: Endogenous and exogenous cues in a multi agent environment

These hypothesized cues will be used as a guidance to find the moments in the narrative where update cues might have occurred. The accuracy of these findings depends on the accuracy of the investigation reports.

Mapping of the intent vector

As seen above it is proposed that attaining the time of an intent vector can be an update cue. It is therefore important to map out the formation and development of the intent SA before the state SA is mapped. As in the definition by Stroeve, Blom and Van der Park (2003) the intent vector consists of a continuous state, discrete state as well as any expansions. Here $\bar{\theta}^{j}_{t,k}$ is the discrete state of agent j as perceived by agent k at times t to be achieved by agent j at time $\bar{t}^j_{t,k}$. And $\bar{x}^j_{t,k}$ is the continuous state of agent j that agent k expects him to achieve at time $\bar{t}^j_{t,k}$. In retrospective modelling the time t can be precisely determined from flight data recorders, cockpit voice recorders etc. The estimate by the agent on the other hand is hard to gauge: unless specifically spoken and recorded by the CVR it is nigh impossible to ascertain the time the agent believes a state will be attained. It is therefore far easier to either treat time qualitatively and comparative or as a precondition.

In the first case the time estimates made by the agents are assumed to be relative to previous notions of the time it would take. So, after an update the agent might be made aware that a state will be reached sooner than previously anticipated. The intent vector will then update from t to t_b where t_b is an earlier time point. How much earlier is hard to predict, so no quantization is made.

Secondly, time can be considered a precondition. For example a pilot might expect the aircraft to be in the discrete state *take_off*, at the time that permission is granted by ATC. In this scenario *permission granted* is the condition and the time can be designated as $t_{clearance}$ or $t_{clearance} + \delta$ indicating a non specified passage of time after takeoff clearance is granted.

If the analysis requires, the intent vector can be expanded with another state that can serve as a precondition. In the takeoff example the pilot might expect the aircraft to be in the discrete state *take_off* at the time permission is granted IF the cabin crew has been seated. There can be of course a copious amount of preconditions for every action during taxi operations which are not relevant to the analysis. The SA differences identified in the orientation phase can once again be used to determine which states could have relevance.

Analyse the SA update for every cue

The last step in formalising the SA difference is mapping how the perceived states of all agents changed at each identified cue. These changes are documented as well as any assumptions that are made.

Mapping the SA difference propagation

The last step in the methodology is mapping the propagation of the SA differences. It is hypothesized that SA differences can cause other SA differences through the various interactions that the agents perform. It is therefore important to muse in what way the differences might have influenced each other, and how the resolution of one difference could have affected the other.

The following three chapters will see the application of the steps outlined in this chapter. After every analysis the findings will be contrasted with the official findings to see if MASA analysis can give new insights and have a meaningful contribution to the base of knowledge about that particular accident.

3. Case 1: B742 / B741, Los Rodeos Tenerife, 1977

In this chapter Multi Agent Situation Awareness analysis is used on the Tenerife runway incursion of 1977. To this day this accident between a KLM B747 and a Pan American B747 remains the deadliest disaster in aviation. There were no component failures in either of the two aircraft, so this case is the perfect case to analyze from the point of SA differences. The online safety repository Skybrary, summarises the accident as follows (Skybrary):

On 27 March 1977, a KLM B747-200 commenced its daylight take off at Los Rodeos airport, Tenerife in very poor visibility, recorded as 300 metres three minutes earlier, after receiving only a departure clearance and continuing the take off roll even after ATC advised "standby for takeoff". This resulted in a collision with a Pan American Airways Boeing 747-100 which was taxiing on the runway in accordance with its ATC clearance issued on the same radio frequency. All 248 people on board the KLM aircraft died and only 61 of the 396 people on board the Pan American aircraft survived.

More on the event can be read at this reference and the official accident report by the Spanish government (Subsecretaria de Avion Civil, 1977). The official accident report cites the following causes of the accident:

The fundamental cause of this accident was the fact that the KLM captain:

- *Took off without clearance.*
- *Did not obey the "stand by for take-off" from the tower.*
- *Did not interrupt take-off when Pan Am reported that they were still on the runway.*
- *In reply to the flight engineer's query as to whether the Pan Am aeroplane had already left the runway, replied emphatically in the affirmative.*

Next to this the following contributing causes are cited*:*

- *It is considered that the captain might have been pressed by the strict work time regulations, any delay could mean an interruption of the flight.*
- *The runway visibility was around 1 kilometre with intermittent light rain and fog patches. This prevented the crews from seeing the other aircraft until the accident was imminent. The control tower also was not able to observe the runway.*

- *The "stand by for take-off … I will call you" from the tower coincided with Pan Am's "we are still taxiing down the runway", which caused a whistling sound which interfered with the communication.*
- *Inadequate language skills .*
- *Pan American flight missing 'the third' taxiway.*
- *Unusual traffic congestion at Los Rodeos airport.*

After the Multi Agent Situation Awareness (MASA) analysis is completed these causes will be contrasted with the findings of the analysis. As will be shown in the last section of this chapter, thinking in terms of SA differences leads additional information.

3.1. Multi Agent Situation Awareness Analysis of Case 1

First off, the steps of the orientation phase are performed.

Identify likely SA differences

SA difference 1: The captain of the KLM aircraft was not aware that take-off clearance had not been issued yet. ATC and the PanAm crew were aware of the fact that it was not granted yet. The KLM onboard engineer expressed doubt about this fact.

SA difference 2: The captain of the KLM aircraft was not aware that the PanAm aircraft was still taxiing on the runway. The PanAm crew and ATC were both aware of this fact.

SA difference 3: The PanAm aircraft had already passed the third taxiway. The PanAm crew was not aware of this fact.

SA difference 4: The PanAm aircraft had already passed the third taxiway. ATC did not realise this.

SA difference 5: The crew of the KLM aircraft had initiated takeoff. ATC did not realise this.

Identify likely root causes of the SA differences

SA difference 1: During ATC clearance the term "take-off" was used which possibly led the captain to believe that they had in fact received take-off clearance.

SA difference 2: The KLM captain misinterpreted ATC's flight path instructions as takeoff clearance thus believing the runway being clear. Furthermore, a shrill noise in the KLM cockpit further delayed the crew from being aware of the presence of the PanAm aircraft. Thirdly, there was the poor visibility on the runway.

SA difference 3: The runway visibility was rather poor due to fog and intermittent rain. Additionally, the runway centrelights were out of service.

SA difference 4: : A better visibility could have potentially allowed ATC to observe that the PanAM aircraft had missed the third taxiway.

SA difference 5: A better visibility could have potentially allowed ATC to observe that the KLM aircraft had initiated the takeoff. ATC did also not interpret the sentence "we are at take-off" as "we are taking off".

Identify likely consequences of SA differences

SA difference 1: The crew initiated the takeoff even though they were not given clearance.

SA difference 2: The crew initiated the takeoff even when the runway was not clear.

SA difference 3: The PanAm aircraft was longer present on the runway than intended.

SA difference 4: ATC did not intervene when the PanAm flight missed its exit.

SA difference 5: ATC did not intervene to stop the KLM aircraft from taking off.

Identify relevant agents in the scenario

From the identified differences in SA the following agents are found to be relevant to the scenario:

- KLM aircraft
- PanAm aircraft
- KLM captain

- KLM onboard engineer
- KLM co-pilot
- PanAm crew
- ATC
- Weather
- Runway

The KLM crew is identified as three separate agents instead of one agent as in the case of the PanAm crew. This is because the narrative of the incursion specifically mentions critical interactions between the KLM engineer, co-pilot and the captain.

Identify likely non-nominal SA update types in the scenario

- Communication between ATC and the KLM captain failed. ATC could not relay the message that the runway was occupied and that take off was to be postponed until a later moment.
- The PanAm crew failed to observe that they had passed the third taxiway.
- The KLM captain failed to correctly process the information regarding the occupied runway when the PanAm crew relayed this.
- ATC did not observe that the PanAm aircraft had missed their intended exit.
- The onboard engineer of the KLM aircraft failed to challenge the captain of the KLM aircraft when he noticed that he was taking off without clearance.

Identify likely sub-models in the scenario

KLM crew members:

- Human error in maintaining SA about the runway. The captain failed to process the information about the still present PanAm flight. The captain made a decision error possibly because of substandard conditions of the operator because the captain was being pressed by duty-time limit.
- Cognitive Control Mode switched to Opportunistic as everyone was in a hurry to leave the congested airport.

ATCo:

- The intent of ATC differed from the KLM crew members.
- Human error in the phraseology. Mentioning the word "takeoff" during ATC clearance might have given the impression that takeoff clearance was granted. This is a form of latent error.

PanAm crew:

- Error in maintaining SA about the position of the PanAm aircraft.

As a next step, an ontology needs to be created that defines all the possible objects and events. The following tables show this. Table 1 shows all the identities for the relevant agents. Table 2 shows the relevant continuous state variables, as well as the relevant discrete state variables.

Table 1: The relevant agents for the scenario

Agent#	Agent	Identity
1	KLM aircraft	ac_klm
2	PanAm aircraft	ac_panam
3	KLM captain	captain_klm
4	KLM onboard engineer	engineer_klm
5	KLM copilot	copilot_klm
6	PanAm crew	crew_panam
7	Air traffic controller	atco
8	Weather	weather
9	Runway	runway

Table 2: The states of the agents relevant to the scenario

Agent	Continuous states	Discrete states
KLM aircraft	*not relevant*	hold, take_off, ground_run
PanAm aircraft	aft _exit, exit, pre_ exit	taxi, hold
KLM captain	*not relevant*	clearance, no_clearance
KLM onboard engineer	*not relevant*	clearance, no_clearance
KLM copilot	*not relevant*	clearance, no_clearance
PanAm crew	*not relevant*	*not relevant*
Air traffic controller	*not relevant*	*not relevant*
Weather	*not relevant*	clear, low_visibility
Runway	*not relevant*	lights_off, lights_on

nb: *hold* is the discrete state of the KLM aircraft where it is waiting for clearance. *Take_off* denotes the mode when clearance is granted but the aircraft has not started its groundrun. G*round_run* is the mode where the aircraft starts moving. *Clearance* is used as a discrete mode for the human agents within the KLM aircraft and denotes their awareness whether they are allowed to take off or not. Note that having the discrete mode *clearance* does not necessarily mean that clearance is actually granted. Some states are "not relevant" this means that the analysis has not found them to be part of a SA difference or contributing to an accident. For example a continuous state of the KLM aircraft could be the position on the airfield. However it is deemed that this state does not add any extra information on top of the already present discrete states, *hold*, *take_off*, and *ground_run*. These discrete states convey the same information. If the discrete state of the KLM aircraft is *hold* it conveys that the continuous state is *at the hold position*. This double information is withheld from the analysis.

Since the broad definition of agent is used it is important to consider which agents can have SA. Out of the agents identified for this example the following can acquire and maintain SA:

1: KLM aircraft
2: PanAm aircraft
3: KLM captain
4: KLM onboard engineer
5: KLM copilot
6: PanAm crew
7: Air traffic controller

Note that the two aircraft can only have SA about themselves. Now that the agents that can acquire and maintain SA are found, the points where their SA is updated must be identified. The cues hypothesized in chapter 2 are used as a guideline. The following moments have been selected:

t0: 1705:27.08, the point where the KLM aircraft has finished taxiing down the runway and has made a complete turn to face the take-off direction.
t1: 1705:50, the point where the KLM copilot has finished requesting ATC clearance.
t2: 1706:16.11, the point where the KLM copilot finished repeating ATC instructions.
t3: 1706:21.79, the point where ATC told the KLM crew to stand by for take-off clearance.
t4: 1706:25.47-1706:29.59, the point where ATC requested the PanAm crew to report when the PanAm aircraft had cleared the runway and they subsequently acknowledged this request.
t5: 1706:34.70, the point where the KLM onboard engineer questioned whether the PanAm aircraft had cleared the runway.

Finally the SA vector for these agents will be analyzed after every timestamp.

Timestamp t0

The state SA for the relevant agents at t0 can be found in table 3. The intent for all agents is in table 4 The following assumptions are made:

A1: At t0 there are no differences in SA among the crew of the KLM aircraft.

A2: The KLM crew believe the runway is not clear as they have only just turned around themselves. Furthermore, the fact that the runway lights were out of order was also relayed to the PanAm crew, so it is assumed that via the reasoning function, the SA about the presence of another aircraft is either updated or confirmed prior to this point.

A3: The weather does not change during the scenario nor does the perception of it by the other agents.

A4: The state of the runway does not change nor does the perception of it by the other agents

A5: The intent of the KLM crew is homogenous. This assumption is extended so that the anticipated intent of the entire crew as perceived by the other agents is the same as well.

A6: All the agents have the same anticipation for the KLM aircraft, namely to *hold, take_off, ground_run* in that order.

A7: All agents have the same anticipation for the PanAm aircraft, namely to *taxi* on the *runway* and to *hold* after the exit is passed.

A8: The KLM crew expects the aircraft to go into *take_off* mode as soon as the KLM crew goes into *clearance* mode. This time is denoted as $t_{clearance}$. Note that this does not mean the clearance has actually been granted.

A9: After a very short time the KLM crew anticipates the aircraft to move into *ground_run*. This time is denoted as $t_{clearance} + \delta$. Everything prior to that is indicated as $t_{clearance}^{-}$.

The time t_{exit} specifies the time at which the PanAm aircraft is at the exit and has cleared the runway. For all the agents it is assumed that $t_{clearance}$ is at a later or at least equal time as t_{exit}.

	KLM a/c	PanAm a/c	KLM captain	KLM engineer	KLM copilot	PanAm crew	ATC	Weather	Runway
KLM a/c	$\sigma^1_{t0,1} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ hold \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\{\cdot\}$	$\{\cdot\}$	$\{\cdot\}$	$\{\cdot\}$	$\{\cdot\}$	$\{\cdot\}$	$\{\cdot\}$	$\{\cdot\}$
PanAm a/c	$\{\cdot\}$	$\sigma^2_{t0,2} = \begin{pmatrix} ac\_panam \\ pre\_exit \\ taxi \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\{\cdot\}$	$\{\cdot\}$	$\{\cdot\}$	$\{\cdot\}$	$\{\cdot\}$	$\{\cdot\}$	$\{\cdot\}$
KLM captain	$\sigma^1_{t0,3} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ hold \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^2_{t0,3} = \begin{pmatrix} ac\_panam \\ \{\cdot\} \\ taxi \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^3_{t0,3} = \begin{pmatrix} captain\_klm \\ \{\cdot\} \\ no\_clearance \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^4_{t0,3} = \begin{pmatrix} engineer\_klm \\ \{\cdot\} \\ no\_clearance \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^5_{t0,3} = \begin{pmatrix} copilot\_klm \\ \{\cdot\} \\ no\_clearance \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^6_{t0,3} = \begin{pmatrix} crew\_panam \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^7_{t0,3} = \begin{pmatrix} atco \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^8_{t0,3} = \begin{pmatrix} weather \\ \{\cdot\} \\ low\_visibility \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^9_{t0,3} = \begin{pmatrix} runway \\ \{\cdot\} \\ lights\_off \\ \hat{v}^j_{t,k} \end{pmatrix}$
KLM engineer	$\sigma^1_{t0,4} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ hold \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^2_{t0,4} = \begin{pmatrix} ac\_panam \\ \{\cdot\} \\ taxi \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^3_{t0,4} = \begin{pmatrix} captain\_klm \\ \{\cdot\} \\ no\_clearance \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^4_{t0,4} = \begin{pmatrix} engineer\_klm \\ \{\cdot\} \\ no\_clearance \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^5_{t0,4} = \begin{pmatrix} copilot\_klm \\ \{\cdot\} \\ no\_clearance \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^6_{t0,4} = \begin{pmatrix} crew\_panam \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^7_{t0,4} = \begin{pmatrix} atco \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^8_{t0,4} = \begin{pmatrix} weather \\ \{\cdot\} \\ low\_visibility \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^9_{t0,4} = \begin{pmatrix} runway \\ \{\cdot\} \\ lights\_off \\ \hat{v}^j_{t,k} \end{pmatrix}$
KLM copilot	$\sigma^1_{t0,5} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ hold \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^2_{t0,5} = \begin{pmatrix} ac\_panam \\ \{\cdot\} \\ taxi \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^3_{t0,5} = \begin{pmatrix} captain\_klm \\ \{\cdot\} \\ no\_clearance \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^4_{t0,5} = \begin{pmatrix} engineer\_klm \\ \{\cdot\} \\ no\_clearance \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^5_{t0,5} = \begin{pmatrix} copilot\_klm \\ \{\cdot\} \\ no\_clearance \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^6_{t0,5} = \begin{pmatrix} crew\_panam \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^7_{t0,5} = \begin{pmatrix} atco \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^8_{t0,5} = \begin{pmatrix} weather \\ \{\cdot\} \\ low\_visibility \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^9_{t0,5} = \begin{pmatrix} runway \\ \{\cdot\} \\ lights\_off \\ \hat{v}^j_{t,k} \end{pmatrix}$
PanAM crew	$\sigma^1_{t0,6} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ taxi \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^2_{t0,6} = \begin{pmatrix} ac\_panam \\ pre\_exit \\ taxi \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^{3,4,5}_{t0,6} = \begin{pmatrix} crew\_klm \\ \{\cdot\} \\ no\_clearance \\ \hat{v}^j_{t,k} \end{pmatrix}$			$\sigma^6_{t0,6} = \begin{pmatrix} crew\_panam \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^7_{t0,6} = \begin{pmatrix} atco \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^8_{t0,6} = \begin{pmatrix} weather \\ \{\cdot\} \\ low\_visibility \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^9_{t0,6} = \begin{pmatrix} runway \\ \{\cdot\} \\ lights\_off \\ \hat{v}^j_{t,k} \end{pmatrix}$
ATC	$\sigma^1_{t0,7} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ taxi \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^2_{t0,7} = \begin{pmatrix} ac\_panam \\ \{\cdot\} \\ taxi \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^{3,4,5}_{t0,7} = \begin{pmatrix} crew\_klm \\ \{\cdot\} \\ no\_clearance \\ \hat{v}^j_{t,k} \end{pmatrix}$			$\sigma^6_{t0,7} = \begin{pmatrix} crew\_panam \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^7_{t0,7} = \begin{pmatrix} atco \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^8_{t0,7} = \begin{pmatrix} weather \\ \{\cdot\} \\ low\_visibility \\ \hat{v}^j_{t,k} \end{pmatrix}$	$\sigma^9_{t0,7} = \begin{pmatrix} runway \\ \{\cdot\} \\ lights\_off \\ \hat{v}^j_{t,k} \end{pmatrix}$

Table 3: state SA of the relevant agents at t0

Table 4: Intent SA of the relevant agents at t0

<table>
<tr><th></th><th>KLM a/c</th><th>PanAm a/c</th><th>K LM captain</th><th>KLM engineer</th><th>KLM copilot</th><th>PanAm crew</th><th>ATC</th></tr>
<tr><td>KLM captain</td><td>$(v^1_{t0,3})_1 = \begin{pmatrix} take\_off \\ \{\cdot\} \\ t_{clearance} \end{pmatrix}$
$(v^1_{t0,3})_2 = \begin{pmatrix} ground\_run \\ \{\cdot\} \\ t_{clearance} + \delta \end{pmatrix}$</td><td>$(v^2_{t0,3})_1 = \begin{pmatrix} hold \\ exit \\ t_{clearance} \end{pmatrix}$</td><td>$(v^3_{t0,3})_1 = \begin{pmatrix} clearance \\ \{\cdot\} \\ t_{clearance} \end{pmatrix}$</td><td>$(v^4_{t0,3})_1 = \begin{pmatrix} clearance \\ \{\cdot\} \\ t_{clearance} \end{pmatrix}$</td><td>$(v^5_{t0,3})_1 = \begin{pmatrix} clearance \\ \{\cdot\} \\ t_{clearance} \end{pmatrix}$</td><td>$v^6_{t0,3} = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix}$</td><td>$v^7_{t0,3} = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix}$</td></tr>
<tr><td>PanAM crew</td><td>$(v^1_{t0,6})_1 = \begin{pmatrix} hold \\ \{\cdot\} \\ t^-_{clearance} \end{pmatrix}$
$(v^1_{t0,6})_2 = \begin{pmatrix} take\_off \\ \{\cdot\} \\ t_{clearance} \end{pmatrix}$
$(v^1_{t0,6})_3 = \begin{pmatrix} ground\_run \\ \{\cdot\} \\ t_{clearance} + \delta \end{pmatrix}$</td><td>$(v^2_{t0,6})_1 = \begin{pmatrix} hold \\ exit \\ t_{exit} \end{pmatrix}$</td><td colspan="3">$(v^{3,4,5}_{t,6})_1 = \begin{pmatrix} clearance \\ \{\cdot\} \\ t_{clearance} \end{pmatrix}$</td><td>$(v^2_{t0,2})_1 = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix}$</td><td>$v^7_{t0,6} = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix}$</td></tr>
<tr><td>ATC</td><td>$(v^1_{t0,7})_1 = \begin{pmatrix} hold \\ \{\cdot\} \\ t^-_{clearance} \end{pmatrix}$
$(v^1_{t0,7})_2 = \begin{pmatrix} take\_off \\ \{\cdot\} \\ t_{clearance} \end{pmatrix}$
$(v^1_{t0,7})_3 = \begin{pmatrix} ground\ run \\ \{\cdot\} \\ t_{clearance} + \delta \end{pmatrix}$</td><td>$(v^2_{t0,7})_1 = \begin{pmatrix} hold \\ exit \\ t_{exit} \end{pmatrix}$</td><td colspan="3">$(v^{3,4,5}_{t0,7})_1 = \begin{pmatrix} clearance \\ \{\cdot\} \\ t_{clearance} \end{pmatrix}$</td><td>$(v^6_{t0,7})_1 = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix}$</td><td>$v^7_{t,7} = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix}$</td></tr>
</table>

Timestamp t1

The copilot of the KLM aircraft has requested ATC clearance and the message was heard in the PanAm cockpit as well. This communication updates the SA of ATC and the PanAm crew in the following way:

- Via the communication function the discrete state of the intent SA of the PanAm crew about the KLM aircraft is confirmed. The time for which they expect the aircraft to go in *take_off* is the moment that clearance is granted.

$$(v^1_{t0,6})_2 = \begin{pmatrix} take\_off \\ \{\cdot\} \\ t_{clearance} \end{pmatrix} \xrightarrow{communication} (v^1_{t1,6})_2 = \begin{pmatrix} take\_off \\ \{\cdot\} \\ t_{clearance} \end{pmatrix} \qquad \text{(confirmation)}$$

- Via the communication function the discrete state of the intent SA of ATC about the KLM aircraft is confirmed.

$$(v^1_{t0,7})_2 = \begin{pmatrix} take\_off \\ \{\cdot\} \\ t_{clearance} \end{pmatrix} \xrightarrow{communication} (v^1_{t1,7})_2 = \begin{pmatrix} take\_off \\ \{\cdot\} \\ t_{clearance} \end{pmatrix} \qquad \text{(confirmation)}$$

- Via the reasoning function the state SA of the PanAm crew about the KLM aircraft discrete state is updated from *taxi* to *hold.*

$$\sigma^1_{t0,6} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ taxi \\ \hat{v}^j_{t,k} \end{pmatrix} \xrightarrow{reasoning} \sigma^1_{t1,6} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ hold \\ \hat{v}^j_{t,k} \end{pmatrix} \qquad \text{(update)}$$

- Via the reasoning function the state SA of ATC about the KLM aircraft discrete state is updated from *taxi* to *hold.*

$$\sigma^1_{t0,7} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ taxi \\ \hat{v}^j_{t,k} \end{pmatrix} \xrightarrow{reasoning} \sigma^1_{t1,7} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ hold \\ \hat{v}^j_{t,k} \end{pmatrix} \text{(update)}$$

Timestamp t2

The important thing to note at this time point is the last sentence of the copilot: “We are now at take-off”. This communication had the following effects:

- The state SA of the PanAm crew about the KLM aircraft was updated from *hold* to *take_off* via the communication function.

$$\sigma_{t1,6}^{1} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ hold \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{communication} \sigma_{t2,6}^{1} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ take\_off \\ \hat{v}_{t,k}^{j} \end{pmatrix} \text{ (update)}$$

- The state SA of ATC about the KLM aircraft was updated from *hold* to *take_off* via the communication function.

$$\sigma_{t1,7}^{1} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ hold \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{communication} \sigma_{t2,7}^{1} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ take\_off \\ \hat{v}_{t,k}^{j} \end{pmatrix} \text{ (update)}$$

- The time in the intent SA of the PanAm crew about the KLM aircrafts intent for a ground run was updated to be a lot sooner than previously anticipated. This occurred via reasoning.

$$(v_{t1,6}^{1})_2 = \begin{pmatrix} ground\_run \\ \{\cdot\} \\ t_{clearance} + \delta \end{pmatrix} \xrightarrow{reasoning} (v_{t2,6}^{1})_2 = \begin{pmatrix} ground\_run \\ \{\cdot\} \\ t_{now} + \delta \end{pmatrix} \text{ (update)}$$

- The time in the intent SA of ATC about the KLM aircrafts intent for a ground run was updated to be a lot sooner than previously anticipated. This occurred via reasoning.

$$(v_{t1,7}^{1})_3 = \begin{pmatrix} ground\_run \\ \{\cdot\} \\ t_{clearance} + \delta \end{pmatrix} \xrightarrow{reasoning} (v_{t2,7}^{1})_3 = \begin{pmatrix} ground\_run \\ \{\cdot\} \\ t_{now} + \delta \end{pmatrix} \text{ (update)}$$

In the above updates, $t_{now} + \delta$ denotes a small time lapse after current moment that can be described as "imminent".

- The discrete state of the KLM aircraft was changed from *hold* to *take_off*.

$$\sigma_{t1,1}^{1} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ hold \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{communication} \sigma_{t2,1}^{1} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ take\_off \\ \hat{v}_{t,k}^{j} \end{pmatrix} \text{ (update)}$$

Timestamp t3

The crucial part in the communication at this timestamp occurs when ATC utters the words "stand by for take-off…I will call you". Almost simultaneously the PanAm crew notified that they were still taxiing down the runway. This communication was heard and processed by ATC, but not by the KLM cockpit. The result is that the SA of ATC about the location of the PanAm aircraft was confirmed or updated, and that ATC assumed that the SA of the KLM crew was as well. The changes in SA are then:

- The state SA of ATC about the location of the PanAm aircraft was updated.

$$\sigma^2_{t2,7} = \begin{pmatrix} ac\_panam \\ \{\cdot\} \\ taxi \\ \hat{v}^j_{t,k} \end{pmatrix} \xrightarrow{communication} \sigma^2_{t3,7} = \begin{pmatrix} ac\_panam \\ pre\_exit \\ taxi \\ \hat{v}^j_{t,k} \end{pmatrix} \text{ (update)}$$

- Through reasoning the KLM crew updated their state SA of the PanAm aircraft.

$$\sigma^2_{t2,i} = \begin{pmatrix} ac\_panam \\ pre\_exit \\ taxi \\ \hat{v}^j_{t,k} \end{pmatrix} \xrightarrow{reasoning} \sigma^2_{t3,i} = \begin{pmatrix} ac\_panam \\ exit \\ hold \\ \hat{v}^j_{t,k} \end{pmatrix} \text{ (update)}$$

For *i = 3, 4, 5*

- The second intent SA vector of the PanAm crew about the KLM crew (all three) and their intent for clearance was updated to a new later time because they assumed the message was relayed correctly.

$$(v^i_{t2,6})_2 = \begin{pmatrix} clearance \\ \{\cdot\} \\ t_{now} \end{pmatrix} \xrightarrow{reasoning} (v^i_{t3,6})_2 = \begin{pmatrix} clearance \\ \{\cdot\} \\ t_{exit} \end{pmatrix} \text{ (update)}$$

for *i = 3, 4, 5*

- The second intent SA vector of ATC about the KLM crew (all three) and their intent for clearance was updated to a new later time because he assumed the message was relayed correctly.

$$(v^i_{t2,7})_2 = \begin{pmatrix} clearance \\ \{\cdot\} \\ t_{now} \end{pmatrix} \xrightarrow{reasoning} (v^i_{t3,7})_2 = \begin{pmatrix} clearance \\ \{\cdot\} \\ t_{exit} \end{pmatrix} \text{ (update)}$$

for *i = 3, 4, 5*

In the above two updates t_{exit} denotes the time at which the PanAm aircraft has cleared the runway.

- The third intent SA of the PanAm crew about the KLM aircraft was updated: the time for *ground_run* was updated to a new later time.

$$(v_{t2,6}^{1})_{3} = \begin{pmatrix} ground\_run \\ \{\cdot\} \\ t_{now} + \delta \end{pmatrix} \xrightarrow{reasoning} (v_{t3,6}^{1})_{2} = \begin{pmatrix} ground\_run \\ \{\cdot\} \\ t_{exit} + \delta \end{pmatrix} \text{ (update)}$$

- The third intent SA vector of ATC about the KLM aircraft was updated: the time for *ground_run* was updated to a new later time.

$$(v_{t2,7}^{1})_{3} = \begin{pmatrix} ground\_run \\ \{\cdot\} \\ t_{now} + \delta \end{pmatrix} \xrightarrow{reasoning} (v_{t3,7}^{1})_{3} = \begin{pmatrix} ground\_run \\ \{\cdot\} \\ t_{exit} + \delta \end{pmatrix} \text{ (update)}$$

In the above two updates $t_{exit} + \delta$ denotes a short time after the PanAm aircraft has cleared the runway.

- The state SA of the PanAm crew about the KLM aircraft was updated: the discrete state of the KLM aircraft changed from *take_off* to *hold*.

$$\sigma_{t2,6}^{1} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ take\_off \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{reasoning} \sigma_{t3,6}^{1} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ hold \\ \hat{v}_{t,k}^{j} \end{pmatrix} \text{ (update)}$$

- The state SA of ATC about the KLM aircraft was updated: the discrete state of the KLM aircraft changed from *take_off* to *hold*.

$$\sigma_{t2,7}^{1} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ take\_off \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{reasoning} \sigma_{t3,7}^{1} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ hold \\ \hat{v}_{t,k}^{j} \end{pmatrix} \text{ (update)}$$

- In the meantime the discrete state of the KLM aircraft had changed from *take_off* to *ground_run*.

$$\sigma_{t2,1}^{1} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ take\_off \\ \hat{v}_{t,k}^{j} \end{pmatrix} \longrightarrow \sigma_{t3,1}^{1} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ ground\_run \\ \hat{v}_{t,k}^{j} \end{pmatrix} \text{ (update)}$$

Timestamp t4

This timestamp is the point at which ATC requests the PanAm crew to report once they had cleared the runway. Subsequently the PanAm crew acknowledged the request. It is another SA update opportunity which reassured the PanAm crew about the intent of the KLM aircraft, as they assumed that the crew was now fully aware of the position of the PanAm aircraft. It also marks the first time where there is difference in SA among the KLM crew. The update of the state SA of the KLM onboard engineer about the PanAm aircraft did go through, while the SA of the KLM captain was not updated.

- The SA of the KLM engineer about the location and mode of the PanAm aircraft was updated.

$$\sigma_{t3,4}^{2} = \begin{pmatrix} ac\_panam \\ exit \\ hold \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{communication} \sigma_{t4,4}^{2} = \begin{pmatrix} ac\_panam \\ pre\_exit \\ taxi \\ \hat{v}_{t,k}^{j} \end{pmatrix} \text{ (update)}$$

- The SA of the KLM engineer about the KLM crew was updated. He realized that they had no clearance.

$$\sigma_{t3,4}^{i} = \begin{pmatrix} crew\_klm \\ \{\cdot\} \\ clearance \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{communication} \sigma_{t4,4}^{i} = \begin{pmatrix} crew\_klm \\ \{\cdot\} \\ no\_clearance \\ \hat{v}_{t,k}^{j} \end{pmatrix} \text{ (update)}$$

for *i = 3, 4, 5*

- The intent SA of the PanAm crew about the KLM crew was confirmed.

$$(v_{t3,6}^{i})_2 = \begin{pmatrix} clearance \\ \{\cdot\} \\ t_{exit} \end{pmatrix} \xrightarrow{reasoning} (v_{t4,6}^{i})_2 = \begin{pmatrix} clearance \\ \{\cdot\} \\ t_{exit} \end{pmatrix} \text{ (confirmation)}$$

for *i = 3, 4, 5*

- The state SA of the PanAm crew about the KLM aircraft was confirmed

$$\sigma_{t3,7}^{1} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ hold \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{reasoning} \sigma_{t4,7}^{1} = \begin{pmatrix} ac\_klm \\ \{\cdot\} \\ hold \\ \hat{v}_{t,k}^{j} \end{pmatrix} \text{ (confirmation)}$$

Timestamp t5

The sixth timestamp, moments before the accident, shows that there was a difference in SA between the KLM engineer and captain that was not resolved. The cockpit voice recorder shows that the engineer was expressing doubt about the location and mode of the PanAm aircraft. This was the last opportunity for the captain to update his SA. The captain responded to the engineer's doubts by ensuring him that the PanAm aircraft had in fact already cleared the runway, so quite possibly the opposite happened: instead of updating the captain with the correct SA, the engineer's SA was updated incorrectly.

- The SA of the KLM engineer about the PanAm aircraft was updated.

$$\sigma_{t4,4}^{2} = \begin{pmatrix} ac\_panam \\ runway \\ taxi \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{communication} \sigma_{t5,4}^{2} = \begin{pmatrix} ac\_panam \\ exit \\ hold \\ \hat{v}_{t,k}^{j} \end{pmatrix} \text{ (update)}$$

- The SA of the KLM engineer about the KLM crew was updated.

$$\sigma_{t4,4}^{i} = \begin{pmatrix} crew\_klm \\ \{\cdot\} \\ no\_clearance \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{communication} \sigma_{t5,4}^{i} = \begin{pmatrix} crew\_klm \\ \{\cdot\} \\ clearance \\ \hat{v}_{t,k}^{j} \end{pmatrix} \text{ (update)}$$

for *i = 3, 4, 5*

Propagation of SA differences

Now that all the updates have been mapped it is time to look at the propagations of the SA differences. First off it can be assumed that SA difference 1 was a direct cause for the accident. If SA difference 1 would have been resolved in a timely fashion it is reasonable to believe that the KLM aircraft would not attempt a take-off and that the accident would therefore have not occurred. It was shown that this SA difference occurred at timestamp t2 where instruction for the flight path were probably misinterpreted as take-off clearance. It is reasonable to assume that this SA difference caused SA difference 2 to arise or at least to cement it: if clearance is granted, this means the runway is clear of other aircraft. Arguing the other way around is also possible, the assumption that the runway is clear can cause the agent to believe that it had clearance. At t3 SA difference 5 occurs.

Most likely because ATC assumed that the communication was properly received. In other words, ATC was not aware that SA difference 1 existed, and in this way it can be argued that SA difference 1 also caused SA difference 5. Finally SA differences 3 and 4 seem to have been self contained. They were neither caused nor exacerbated by the other SA differences. However, as will be explained below, the resolution of these differences could have resolved difference 1, ultimately preventing the accident from happening. Figure 3 visualizes these relations for case 1.

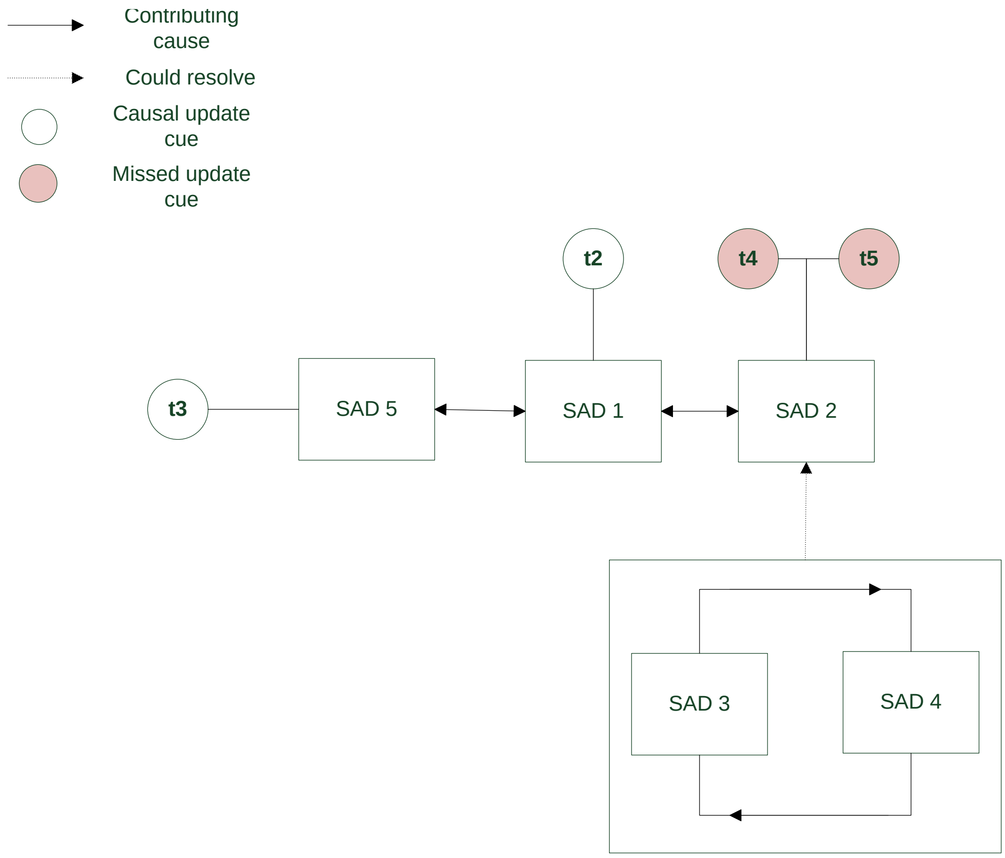

Figure 3: Relations between the SA differences in case 1

The full arrows represent a potential causal relation between the identified SA differences. A → B should be read as: A was a contributing cause of B. The dashed arrows represent that the resolution of one difference could have resolved the other. A → B should be read as: the resolution of A could lead to a resolution of B.

3.2. Results from the analysis of case 1

The above analysis shows how several differences in SA occurred and propagated. This section will contrast these findings with the conclusion of the official investigation report of the accident. The differences in SA that are found are reiterated and correlated with the found causes.

SA difference 1: The captain of the KLM aircraft was not aware that take-off clearance had not been issued yet. ATC and the PanAm crew were aware of the fact that it was not granted yet. The KLM onboard engineer expressed doubt about this fact.

This coincides with the first fundamental cause identified by the Spanish report: the KLM captain took off without clearance. It is also related to the second fundamental cause, namely that the KLM captain ignored the “stand-by for takeoff” issued by ATC. As contributing causes the report cites inadequate language because ATC clearance might have been interpreted as take-off clearance. Consequently, one of the recommendations in the report is the avoidance of the word “take-off” when issuing ATC clearances.

The difference in SA between the KLM crew and the PanAm crew was further exacerbated when ATC requested the PanAm crew to report when the runway was clear confirming their belief that the KLM crew did in fact not have clearance. The official report fails to show the importance of this fact. Simply stating that one of the contributing causes was that the KLM captain took off without clearance narrows down the search for preventive measures. If SA difference 1 is considered a fundamental cause, the recommendations would focus on the PanAm crew and the KLM engineer specifically as well as the KLM captain and ATC. All these agents might have helped resolve the MA-SA differences, and thus help prevent the accident from happening, so why only focus on ATC and the KLM captain?

SA difference 2: The captain of the KLM aircraft was not aware that the PanAm aircraft was still taxiing on the runway. The PanAm crew and ATC were both aware of this fact.

This SA differences ties into the third and fourth fundamental cause as identified by the Spanish government. By all accounts it is clear that the crew of the KLM aircraft was not aware of the location of the PanAm aircraft when they initiated their ground run. As a contributing cause the Spanish government cites the two interfering voice communications from ATC and the PanAm crew which caused a whistling sound in the KLM cockpit. The last opportunity for the KLM captain to become aware of the PanAm aircraft’s location happens just a few seconds later when ATC requests the PanAm aircraft to report when they had cleared the runway. This exogenous cue was not processed by the KLM crew with maybe the exception of the KLM onboard engineer whose doubts were quickly "corrected" by the captain.

The multi agent SA analysis shows more clearly how SA difference 2 occurred, and more importantly how it is a follow up from SA difference 1. When looking at the cues that resulted in SA difference 2 one can see that it was a result of an erroneous update through the reasoning function. The cue in question can be considered a *collapse of intent* as explained in chapter 2. This collapse was reached because the KLM crew *thought* that they had received take-off clearance. This is an example of an endogenous cue that was set in motion by an exogenous cue, namely ATC issuing ATC clearance and this being misinterpreted for take-off clearance. To state it differently: because lack of exogenous cues the KLM crew's SA was falsely updated by an endogenous cue set in motion by the first difference in SA. This is where the multi agent SA analysis sets itself apart from the accident report where the correlation between the two contributing factors is not specifically mentioned or explored.

SA difference 3: The PanAm aircraft had already passed the third taxiway. The PanAm crew was not aware of this fact.

SA difference 4: The PanAm aircraft had already passed the third taxiway. ATC did not realise this.

These SA differences correspond to one of the contributing causes from the report, namely that the PanAm crew had missed the third taxiway. It is implicitly stated that due to this the PanAm aircraft was longer on the runway than intended. Had they taken the correct exit, this event would have been merely an incident where a crew took off without clearance.

Thinking in terms of multi agent SA differences shows an interesting interplay between the relevant agents that shows how the accident could have been prevented even though the PanAm crew stayed on the runway for a longer than intended timeframe. It shows how resolving this SA difference would have been sufficient to prevent the accident. There are two possible scenarios:
1: SA difference 3 is the first one to be resolved. The PanAm crew could have had some exogenous cue like a sign, or an endogenous cue like the collapse of an intent vector. Once they resolved their SA difference it is reasonable to assume that they would have contacted ATC to relay them this fact and in turn allowing ATC to resolve SA difference 4.
2: SA difference 4 is the first one to be resolved. This would depend entirely on the visibility of the PanAm aircraft and the likelihood of ATC monitoring the taxiing aircraft. Had this happened, ATC

would have contacted the PanAm crew notifying them of this fact and hence SA difference 3 would be resolved.

Regardless of the order, the SA analysis shows how resolving difference 3 and 4 could have resolved difference 1, and 2. Firstly the extra chatter over the radio between ATC and the PanAm crew could have been a sufficient cue for the KLM crew to halt their take-off. Secondly, it is not unreasonable to assume that ATC would have contacted the KLM crew to tell them that they would receive take-off clearance at a later timeframe because of the lost PanAm aircraft. So whereas the official report only mentions the errors that are made, thinking in terms of SA differences gives a clue about possible contingencies that could have prevented the error in becoming something catastrophic.

SA difference 5: The crew of the KLM aircraft had initiated takeoff. ATC did not realise this.

This difference in SA does not correlate with any of the causes in the official accident report, while it can be clearly seen how the accident would not have occurred had the difference not existed. Assuming ATC had realized the difference in SA and assuming no malicious intent, it is reasonable to expect that he would have intervened in a more stringent manner. As it is now, ATC interpreted the sentence "we are now at take-off" from the KLM crew as "we are ready to take-off at your leisure". Following this he replied with a "O.K...I will call you". Had ATC interpreted the KLM crew's notification as "we are now taking off" he most likely would have replied with a more urgent message. Emphasizing that the crew does not have clearance and that the runway is not clear. Consequently this would have resolved SA difference 1 and SA difference 2.

To summarize the findings: the case shows very clearly the added benefit of approaching accidents with multi agent SA differences in mind. It allows for easy identification of the relevant agents and it shows interesting interplays between these agents and how SA differences propagated and formed other SA differences. A multi agent SA analysis allows for better musing about not only prevention of these differences, but also about contingencies that could have prevented the accident even though the differences occurred.

4. Case 2: B752/B752, en route, north of Tenerife 2011

For the second example a recent incident that occurred at Tenerife involving the Traffic Collision Avoidance System (TCAS) will be analyzed. This case was selected to see how MASA analysis fairs when there is a non-human agent involved that can maintain SA about agents other than itself. The TCAS is a fully automated system that minimizes the risks of mid air collisions by issuing either Traffic Advisories (TA) or Resolution Advisories (RA) to the pilots of the aircraft that are impinging on each other's airspace. So this near miss becomes a showcase for how non-human proactive and autonomous agents can also be analyzed in terms of SA differences.

The incident is summarized as follows by the skybrary database:

On 20 November 2011, a problem in reading the altitude labels on ATC radar control display led to a Finnair Boeing 757 being cleared to make a descent which brought it into proximity with a Thomas Cook Boeing 757 in day VMC. Co-ordinated TCAS RAs were generated onboard both aircraft but when the Finnair aircraft failed to respond to its Climb RA and continued descent, the other aircraft, which had responded correctly to its initial RA, received a further RA to reverse their descent to a climb. The Finnair aircraft reported retained visual contact with the other aircraft throughout.

The entire case can be read at the reference (Skybrary, B) and the official investigation report (CIAIAC, 2012). The report cites the following contributing causes:

- *A possible garbling problem that made it possible for the label on the radar display for (the Thomas Cook) aircraft to show that it was flying at flight level 405.*
- *The controller's failure to detect the fault that existed with the labels.*
- *The improper response by the crew of (the Finnair) aircraft to the climb advisory issued by its TCAS.*

Following these conclusions two recommendations were issued:

- *It is recommended that AENA ensure that its controllers are aware of the faults that can occur affecting the auxiliary systems used over the course of their jobs and of the way to detect and address said faults.*

- *It is recommended that the operator, FINNAIR, as part of the refresher training given to its crews, enhance the training on the procedures to be followed in the event of aircraft encounters resulting in a TCAS activation.*

These causes will be contrasted with the findings of the Multi Agent SA difference analysis.

4.1. Multi Agent Situation Awareness Analysis of Case 2

As previously, the steps of the orientation phase are performed first.

Identify likely SA differences

SA difference 1**:** ATC believed that the Thomas Cook (ThoCo) was at flight level 405. The aircraft was actually at flight level 370 , and this is also what the crew believed.

SA difference 2: The Finnair crew was not aware of the severity of the conflicting location of the Thomas Cook aircraft. The Thomas Cook crew and the TCAS system were aware of this fact.

SA difference 3: ATC was not aware of the loss of separation. The Thomas Cook Crew and the TCAS were aware of this fact.

Identify likely root causes of the SA difference

SA difference 1: ATC failed to deduce the correct flight levels from the radar screen.

SA difference 2: The Finnair crew failed to process the advisory and resolution of the TCAS system.

SA difference 3: Due to SA difference 1 ATC was not aware that the Finnair aircraft would pass through the flightlevel of the Thomas Cook aircraft.

Identify likely consequences of the SA difference

SA difference 1: ATC believes that the Finnair aircraft is below the Thomas Cook aircraft.

SA difference 2: The Finnair crew continued their descent despite the conflict.

SA difference 3: ATC issued descent clearance even though this would cause a potentially hazardous situation.

Identify relevant agents in the scenario

The TCAS is treated as a single agent that can obtain and maintain SA about agents other than itself. There appears to be no difference in crew SA, so both the crews are treated as a homogenous entity. Then the agents are:

- Finnair aircraft
- Thomas Cook aircraft
- Finnair crew
- Thomas Cook crew
- TCAS
- ATC

Identify likely non-nominal SA update types

1: The air traffic controller got relieved just prior to the event. The new controller might have not processed the information on the flight progress strip.

2: Due to a radar error the flight level of the Thomas Cook aircraft was displayed incorrectly.

3: The TA that the Finnair aircraft received was followed immediately by an RA potentially drowning out the TA warning. Pilots are used to first receive a TA followed by an RA. This could have been one of the reasons the RA was questioned by the crew.

Identify likely submodels in the scenario

Finnair crew:

- Human error in following protocols. The Finnair crew did not follow the TCAS instruction.
- Human error in maintaining SA. Specifically the second and third level of SA (Endsley, 1995) as the crew failed to understand the potential hazard of the location of the Thomas Cook aircraft.

ATC:

- Human error in maintaining SA. ATC failed to process the information on the flight progress strip.

As a next step, an ontology needs to be created that defines all the possible objects and events. The following tables show this. Table 5 shows all the identities for the relevant agents. Table 6 shows the relevant continuous state variables, as well as the relevant discrete state variables.

Table 5: The relevant agents in the second scenario

Agent#	Agent	Identity
1	Finnair aircraft	ac_finnair
2	Thomas Cook aircraft	ac_thoco
3	Finnair crew	crew_finnair
4	Thomas Cook crew	crew_thoco
5	*TCAS*	*tcas*
6	ATC	atco

Table 6: The relevant states of the agents in the second scenario

Agent	Continuous states	Discrete states
Finnair aircraft	increasing_conflict, decreasing_conflict, no_conflict	top_fl, bot_fl, equal_fl
Thomas Cook aircraft	increasing_conflict, decreasing_conflict, no_conflict	top_fl, bot_fl, equal_fl
Finnair crew	*n.a*	ascend, descend, level
Thomas Cook crew	*n.a*	ascend, descend, level
TCAS	*n.a*	ra_descend, ra_ascend, ra_level, ta
ATC	*n.a*	*n.a*

The discrete states of both aircraft refer to their relative height. If the Finnair aircraft is factually in top_fl, this means that it is at a higher flight level. This automatically means the Thomas Cook aircraft is in bot_fl. The continuous states of both aircraft pertains to their changing position and indicates whether it increases the conflict (i.e getting closer) or decreases it (getting further away). The discrete states of both the crews indicate whether they are in want to make the aircraft descend, ascend or maintain level flight. The discrete states of the TCAS are the traffic advisory and the resolution advisories ascend, descend and maintain level. It also needs to be noted that the TCAS is modeled as a single agent. The result of this is that different agents can have differing SA about the TCAS yet both still be correct, because the TCAS can have different modes for different agents.

The following cues have been selected as moments where updates in SA occurred.

t0: 12:48.10, the Thomas Cook crew requests descent clearance and is told to stand by .

t1: 12:48.20, the point at which descent clearance is granted to the Finnair crew. At the same time the radar error had occurred.

t2: Here the TCAS "notices" that the Finnair aircraft has started descent.

t3: 12:49.32, ATC requests the Thomas Cook aircraft to descend to FL390 after which the Thomas Cook crew informs ATC that they are already at FL370.

t4: Here the TCAS on the Thomas Cook aircraft issued a TA and the TCAS on the Finnair aircraft issued an RA to maintain level flight.

t5: One second later from t4 the TCAS issued a descend RA to the Thomas Cook aircraft and a climb RA to the Finnair.

t6: The Finnair crew has their descent clearance confirmed by ATC.

t7: Here the TCAS issued a climb RA to the Thomas Cook aircraft and a maintain level RA to the Finnair.

Now the SA at every cue will be formalized.

Timestamp t0

The state SA of the relevant agents at this time point can be found in table 5. The following assumptions are made:

A1: The Finnair crew reported retained visual contact with the Thomas Cook aircraft throughout. It is unknown whether they were aware of the identity.

A2: It is assumed the Thomas Cook crew did not have visual contact with the Finnair aircraft.

A3: It is assumed that both aircraft are maintaining level flight at this point.

A4: It is assumed that ATC believes the Thomas Cook crew will honor his request to standby and that he thus perceives them to maintain level flight.

A5: It is assumed that, upon being contacted by the Thomas Cook crew, ATC glanced at his radar and perceived the flight level.

A6: It is assumed that there is no downlink for the TCAS, so ATC is not aware of the states of the system.

A7: It is assumed that ATC intends to grant descent clearance to the Thomas Cook aircraft a time δ after descent clearance has been granted to the Finnair a/c (t1).

Furthermore, since no states were defined for ATC, this agent will be held out of the analysis concerning the SA of other agents about ATC.

The intent SA vector can be found in table 8. It is decided to expand the intent vector with a condition. This condition pertains to the state no_conflict attributed to the aircraft. So, for example the intent vector of the Finnair crew about themselves would become

$$\left(v_{t0,3}^{3}\right)_1 = \begin{pmatrix} descend \\ \{\cdot\} \\ t_{clearance} \\ no\_conflict \end{pmatrix}$$

This means that the crew will go into descend mode as soon as clearance is granted at time $t_{clearance}$ when the condition of no_conflict is true. The condition pertains to a belief and is not a representation of reality: as will be made clear later, if no_conflict is true it does not necessarily mean that there is in fact no conflict.

Table 7: state SA of the relevant agents at t0

	Finnair a/c	Thomas Cook a/c	Finnair crew	Thomas Cook crew	TCAS
Finnair a/c	$\sigma_{t0,1}^{1}=\begin{pmatrix} ac\_finnair \\ no\_conflict \\ top\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$n.a$	$n.a$	$n.a$	$n.a$
Thomas Cook a/c	$n.a$	$\sigma_{t0,2}^{2}=\begin{pmatrix} ac\_thoco \\ no\_conflict \\ bot\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$n.a$	$n.a$	$n.a$
Finnair crew	$\sigma_{t0,3}^{1}=\begin{pmatrix} ac\_finnair \\ no\_conflict \\ top\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,3}^{2}=\begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ bot\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,3}^{3}=\begin{pmatrix} crew\_finnair \\ \{\cdot\} \\ level \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,3}^{4}=\begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,3}^{5}=\begin{pmatrix} tcas \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$
Thomas Cook crew	$\sigma_{t0,4}^{1}=\begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,4}^{2}=\begin{pmatrix} ac\_thoco \\ no\_conflict \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,4}^{3}=\begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,4}^{4}=\begin{pmatrix} crew\_thoco \\ \{\cdot\} \\ level \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,4}^{5}=\begin{pmatrix} tcas \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$
TCAS	$\sigma_{t0,5}^{1}=\begin{pmatrix} ac\_finnair \\ \{\cdot\} \\ top\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,5}^{2}=\begin{pmatrix} ac\_thoco \\ \{\cdot\} \\ bot\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$n.a$	$n.a$	$n.a$
ATC	$\sigma_{t0,6}^{1}=\begin{pmatrix} ac\_finnair \\ \{\cdot\} \\ top\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,6}^{2}=\begin{pmatrix} ac\_thoco \\ \{\cdot\} \\ bot\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,6}^{3}=\begin{pmatrix} crew\_finnair \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,6}^{4}=\begin{pmatrix} crew\_thoco \\ \{\cdot\} \\ level \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,6}^{4}=\begin{pmatrix} tcas \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$

Table 8: Intent SA of the relevant agents at t0

	Finnair a/c	Thomas Cook a/c	Finnair crew	Thomas Cook crew	TCAS
Finnair crew	$\left(v^1_{t0,3}\right)_1 = \begin{pmatrix}\{\cdot\}\\ \{\cdot\}\\ \{\cdot\}\end{pmatrix}$	$\left(v^2_{t0,3}\right)_1 = \begin{pmatrix}\{\cdot\}\\ \{\cdot\}\\ \{\cdot\}\end{pmatrix}$	$\left(v^3_{t0,3}\right)_1 = \begin{pmatrix}descend\\ \{\cdot\}\\ t_{clearance}\\ no\_conflict\end{pmatrix}$	$\left(v^4_{t0,3}\right)_1 = \begin{pmatrix}\{\cdot\}\\ \{\cdot\}\\ \{\cdot\}\end{pmatrix}$	$\left(v^5_{t0,3}\right)_1 = \begin{pmatrix}\{\cdot\}\\ \{\cdot\}\\ \{\cdot\}\end{pmatrix}$
Thomas Cook crew	$\left(v^1_{t0,4}\right)_1 = \begin{pmatrix}\{\cdot\}\\ \{\cdot\}\\ \{\cdot\}\end{pmatrix}$	$\left(v^2_{t0,4}\right)_1 = \begin{pmatrix}\{\cdot\}\\ \{\cdot\}\\ \{\cdot\}\end{pmatrix}$	$\left(v^3_{t0,4}\right)_1 = \begin{pmatrix}\{\cdot\}\\ \{\cdot\}\\ \{\cdot\}\end{pmatrix}$	$\left(v^4_{t0,4}\right)_1 = \begin{pmatrix}descend\\ \{\cdot\}\\ t_{clearance}\\ no\_conflict\end{pmatrix}$	$\left(v^5_{t0,4}\right)_1 = \begin{pmatrix}\{\cdot\}\\ \{\cdot\}\\ \{\cdot\}\end{pmatrix}$
TCAS	$\left(v^1_{t0,5}\right)_1 = \begin{pmatrix}\{\cdot\}\\ \{\cdot\}\\ \{\cdot\}\end{pmatrix}$	$\left(v^2_{t0,5}\right)_1 = \begin{pmatrix}\{\cdot\}\\ \{\cdot\}\\ \{\cdot\}\end{pmatrix}$	*n.a.*	*n.a.*	*n.a.*
ATC	$\left(v^1_{t0,6}\right)_1 = \begin{pmatrix}\{\cdot\}\\ \{\cdot\}\\ \{\cdot\}\end{pmatrix}$	$\left(v^1_{t0,6}\right)_1 = \begin{pmatrix}\{\cdot\}\\ \{\cdot\}\\ \{\cdot\}\end{pmatrix}$	$\left(v^3_{t0,6}\right)_1 = \begin{pmatrix}descend\\ \{\cdot\}\\ t1\\ no\_conflict\end{pmatrix}$	$\left(v^4_{t0,6}\right)_1 = \begin{pmatrix}descend\\ \{\cdot\}\\ t1 + \delta\\ no\_conflict\end{pmatrix}$	*n.a*

Timestamp t1

Here descent clearance is granted to the Finnair aircraft. Arround the same time the radar error occurs. This prompts the following changes in SA:

- ATC has observed the wrong flight level for the Thomas Cook aircraft. This prompts an update in ATC's SA for the discrete states of both aircraft. ATC perceives the state of the Thomas Cook aircraft to be *top_fl.*

$$\sigma_{t0,6}^{2} = \begin{pmatrix} ac\_thoco \\ \{\cdot\} \\ bot\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{observation} \sigma_{t1,6}^{2} = \begin{pmatrix} ac\_thoco \\ \{\cdot\} \\ top\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix} \quad \text{(update)}$$

- ATC perceives the discrete state of the Finnair aircraft to be *bot_fl.*

$$\sigma_{t0,6}^{1} = \begin{pmatrix} ac\_finnair \\ \{\cdot\} \\ top\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{observation} \sigma_{t1,6}^{1} = \begin{pmatrix} ac\_finnair \\ \{\cdot\} \\ bot\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix} \quad \text{(update)}$$

- Via the communication function the state SA of ATC about the Finnair crew is updated. The discrete state changes to *descend*.

$$\sigma_{t0,3}^{3} = \begin{pmatrix} crew\_finnair \\ \{\cdot\} \\ level \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{communication} \sigma_{t1,3}^{3} = \begin{pmatrix} crew\_finnair \\ \{\cdot\} \\ descend \\ \hat{v}_{t,k}^{j} \end{pmatrix} \quad \text{(update)}$$

- Via the communication function the state SA of the Finnair crew about the Finnair crew is updated. The discrete state changes to *descend*.

$$\sigma_{t0,6}^{3} = \begin{pmatrix} crew\_finnair \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{communication} \sigma_{t1,6}^{3} = \begin{pmatrix} crew\_finnair \\ \{\cdot\} \\ descend \\ \hat{v}_{t,k}^{j} \end{pmatrix} \quad \text{(update)}$$

Timestamp t2

The Finnair a/c starts its descent. The following changes occur:

- The state SA of the TCAS system regarding the Finnair a/c was updated to *increasing_conflict.*

$$\sigma_{t1,5}^{1} = \begin{pmatrix} ac\_finnair \\ \{\cdot\} \\ top\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{observation} \sigma_{t2,5}^{1} = \begin{pmatrix} ac\_finnair \\ increasing\_conflict \\ top\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix} \quad \text{(update)}$$

- The discrete state in the intent SA of the TCAS system regarding the Finnair aircraft was updated to *equal_fl.*

$$\left(v^{1}_{t1,5}\right)_1 = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix} \xrightarrow{reasoning} \left(v^{1}_{t2,5}\right)_1 = \begin{pmatrix} equal\_fl \\ \{\cdot\} \\ t_{efl} \end{pmatrix} \text{ (update)}$$

- The discrete state in the intent SA of the TCAS system regarding the Thomas Cook aircraft was updated to *equal_fl.*

$$\left(v^{2}_{t1,5}\right)_1 = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix} \xrightarrow{reasoning} \left(v^{2}_{t2,5}\right)_1 = \begin{pmatrix} equal\_fl \\ \{\cdot\} \\ t_{efl} \end{pmatrix} \text{ (update)}$$

Here t_{efl} is the time for which the TCAS system projects the aircraft to be on an equal flight level barring no changes.

Timestamp t3

At time t3 the error caused by the radar is resolved. Upon requesting the Thomas Cook crew to descend to FL390 ATC is informed by the crew that the aircraft is already at FL370. This updates the state SA of ATC regarding both aircraft:

- The relative flight level for both aircraft is flipped around in the state SA of ATC. The perceived discrete state of the Finnair aircraft is now *top_fl.*

$$\sigma^{1}_{t2,6} = \begin{pmatrix} ac\_finnair \\ \{\cdot\} \\ bot\_fl \\ \hat{v}^{j}_{t,k} \end{pmatrix} \xrightarrow{reasoning} \sigma^{1}_{t3,6} = \begin{pmatrix} ac\_finnair \\ \{\cdot\} \\ top\_fl \\ \hat{v}^{j}_{t,k} \end{pmatrix} \text{ (update)}$$

- The perceived discrete state of the Thomas Cook aircraft is now *bot_fl.*

$$\sigma^{2}_{t2,6} = \begin{pmatrix} ac\_finnair \\ \{\cdot\} \\ top\_fl \\ \hat{v}^{j}_{t,k} \end{pmatrix} \xrightarrow{reasoning} \sigma^{2}_{t3,6} = \begin{pmatrix} ac\_thoco \\ \{\cdot\} \\ bot\_fl \\ \hat{v}^{j}_{t,k} \end{pmatrix} \text{ (update)}$$

Timestamp t4

The TCAS system issues a TA to the Thomas Cook a/c and a remain level flight RA to the Finnair a/c. The result of the TA is that it can be assumed beyond reasonable doubt that the ThoCo crew was now aware of the Finnair. However, there was a coordinated RA just one second later at t5, so it is assumed that most SA updates occurred after this point.

- The state SA of the Finnair crew regarding the TCAS changes from unknown to ra_level.

$$\sigma^5_{t3,3} = \begin{pmatrix} tcas \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix} \xrightarrow{observation} \sigma^5_{t4,3} = \begin{pmatrix} tcas \\ \{\cdot\} \\ ra\_level \\ \hat{v}^j_{t,k} \end{pmatrix} \quad \text{(update)}$$

- The state SA of the Thomas Cook crew regarding the TCAS changes from unknown to *ta.*

$$. \ \sigma^5_{t3,4} = \begin{pmatrix} tcas \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix} \xrightarrow{observation} \sigma^5_{t4,4} = \begin{pmatrix} tcas \\ \{\cdot\} \\ ta \\ \hat{v}^j_{t,k} \end{pmatrix} \quad \text{(update)}$$

- The state SA of the Thomas Cook crew is updated so that they are now aware of an aircraft above them. It is unknown whether they know the identity of the aircraft.

$$\sigma^1_{t3,4} = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix} \xrightarrow{reasoning} \sigma^1_{t4,4} = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ top\_fl \\ \hat{v}^j_{t,k} \end{pmatrix} \quad \text{(update)}$$

Timestamp t5

Only one second later a coordinated RA is issued to both aircrafts. The Finnair receives an RA to climb and the Thomas Cook to descend. The RA issued to the Finnair crew means that the condition to keep descending was not met anymore. This means that the intent SA of the crew about the crew itself was unknown at this point. The reason for this assumption is because at timepoint t6 the crew asks ATC to confirm their clearance. The changes are:

- The state SA of both crew regarding the TCAS is updated. The Finnair crew perceives that the discrete state of the TCAS changes from *ra_level* to *ra_ascend.*

$$\sigma^5_{t4,3} = \begin{pmatrix} tcas \\ \{\cdot\} \\ ra\_level \\ \hat{v}^j_{t,k} \end{pmatrix} \xrightarrow{observation} \sigma^5_{t5,3} = \begin{pmatrix} tcas \\ \{\cdot\} \\ ra\_ascend \\ \hat{v}^j_{t,k} \end{pmatrix} \quad \text{(update)}$$

- The Thomas Cook crew perceives that the discrete state of the TCAS changes from *ta* to *ra_descend.*

$$\sigma^5_{t4,4} = \begin{pmatrix} tcas \\ \{\cdot\} \\ ta \\ \hat{v}^j_{t,k} \end{pmatrix} \xrightarrow{observation} \sigma^5_{t5,4} = \begin{pmatrix} tcas \\ \{\cdot\} \\ ra\_descend \\ \hat{v}^j_{t,k} \end{pmatrix} \quad \text{(update)}$$

- The intent SA of the Finnair crew regarding their condition becomes unknown.

$$\left(v_{t4,3}^{3}\right)_1 = \begin{pmatrix} descend \\ \{\cdot\} \\ t_{clearance} \\ no\_conflict \end{pmatrix} \xrightarrow{reasoning} \left(v_{t5,3}^{3}\right)_1 = \begin{pmatrix} descend \\ \{\cdot\} \\ t_{clearance} \\ \{\cdot\} \end{pmatrix} \text{ (update)}$$

Timestamp t6

Upon receiving the RA the Finnair crew has called into question their descent clearance. After their inquiry, ATC confirms that they in fact have clearance. This updates their intent SA vector once again. Around the same time the Thomas Cook aircraft adheres to the RA and starts to descend.

- The Finnair crew now again believes there is no conflict and the continues their descent.

$$\left(v_{t5,3}^{3}\right)_1 = \begin{pmatrix} descend \\ \{\cdot\} \\ t_{clearance} \\ \{\cdot\} \end{pmatrix} \xrightarrow{reasoning} \left(v_{t6,3}^{3}\right)_1 = \begin{pmatrix} descend \\ \{\cdot\} \\ t_{clearance} \\ no\_conflict \end{pmatrix} \text{ (update)}$$

- The TCAS observes that the ThoCo a/c has started to descend.

$$\sigma_{t5,5}^{2} = \begin{pmatrix} ac\_thoco \\ \{\cdot\} \\ bot\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{observation} \sigma_{t6,5}^{2} = \begin{pmatrix} ac\_thoco \\ decreasing\_conflict \\ bot\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix} \text{ (update)}$$

Timestamp t7

Because the Finnair crew remained in *descend* mode instead of following the RA, and because they were descending at a faster rate they past the flight level of the Thomas Cook aircraft. The TCAS issued two new RA's.

- The state SA of the TCAS about both aircrafts was updated. The continuous state of the Finnair aircraft changes from *increasing_conflict* to *decreasing_conflict.*

$$\sigma_{t6,5}^{1} = \begin{pmatrix} ac\_finnair \\ increasing\_conflict \\ top\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{observation} \sigma_{t7,5}^{1} = \begin{pmatrix} ac\_finnair \\ decreasing\_conflict \\ bot\_fl \\ \hat{v}_{t,k}^{j} \end{pmatrix} \text{ (update)}$$

- The continuous state of the Thomas Cook aircraft changes from *decreasing_conflict* to *increasing_conflict.*

$$\sigma^2_{t6,5} = \begin{pmatrix} ac\_thoco \\ decreasing\_conflict \\ bot\_fl \\ \hat{v}^j_{t,k} \end{pmatrix} \xrightarrow{observation} \sigma^2_{t7,5} = \begin{pmatrix} ac\_thoco \\ increasing\_conflict \\ top\_fl \\ \hat{v}^j_{t,k} \end{pmatrix} \quad \text{(update)}$$

- The state SA of both crews about the TCAS was updated. The discrete state perceived by the Finnair crew changed from *ra_ascend* to *ra_level.*

$$\sigma^5_{t6,3} = \begin{pmatrix} tcas \\ \{\cdot\} \\ ra\_ascend \\ \hat{v}^j_{t,k} \end{pmatrix} \xrightarrow{observation} \sigma^5_{t7,3} = \begin{pmatrix} tcas \\ \{\cdot\} \\ ra\_level \\ \hat{v}^j_{t,k} \end{pmatrix} \quad \text{(update)}$$

- The discrete state perceived by the Thomas Cook crew changed from *ra_descend* to *ra_ascend.*

$$\sigma^5_{t6,4} = \begin{pmatrix} tcas \\ \{\cdot\} \\ ra\_descend \\ \hat{v}^j_{t,k} \end{pmatrix} \xrightarrow{observation} \sigma^5_{t7,4} = \begin{pmatrix} tcas \\ \{\cdot\} \\ ra\_ascend \\ \hat{v}^j_{t,k} \end{pmatrix}$$

Propagation of SA differences

Now the propagation of the SA differences will be analyzed. It is reasonable to assume that ATC was not aware of the loss of separation, SA difference 3, because the controller misread the radar and believed that the Thomas Cook aircraft was at a higher flight level, SA difference 1. Difference 1 occurred at t1 and was resolved at t3. However, this did not resolve difference 3. Difference 3 could have been resolved however by the resolution of 2. If the Finnair crew had been aware of the severity they might have contacted ATC to inform them of this. At the same time the resolution of 3 could have resolved difference 2. Figure 3 shows the relations.

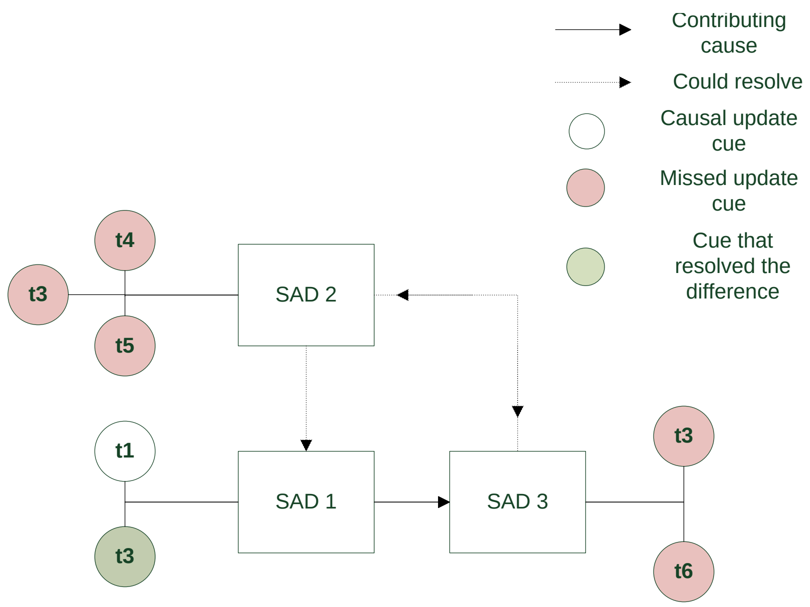

Figure 4: SA difference propagation for case 2

4.2. Results from the analysis of case 2

The above analysis shows how several differences in SA occurred. This section will contrast these findings with the conclusion of the official investigation report of the incident. The differences in SA that are found are reiterated and correlated with the found causes.

SA difference 1: ATC believed that the Thomas Cook was at flight level 405. The aircraft was actually at flight level 370, and this is also what the crew believed.

This finding corresponds to the first and second contributing cause from the official report, namely the possible garbling on the radar display and the controllers inability to detect the issue. Consequently, one of the recommendations is that AENA ensures that its controllers are aware of this problem. The MASA analysis shows that this difference was resolved rather quickly. At t3 the Thomas Cook crew informs that contrary to the controllers belief, their aircraft was at flightlevel 370. Seventeen seconds later ATC confirms descent clearance to the Finnair crew, showing that although the first SA difference was resolved, this did not carry over to resolve the third SA difference.

SA difference 2: The Finnair crew was not aware of the severity of the conflicting location of the Thomas Cook aircraft. The Thomas Cook crew and the TCAS system were aware of this fact.

This corresponds nicely to the third contributing cause from the official report: the improper response by the Finnair crew to the TCAS advisory. Focussing only on the missed TCAS advisory however does not explain the entirety of the SA difference. It was shown that there were two additional cues that could have been used to resolve SA difference 1. Firstly, the Finnair crew reported maintained visual contact with the Thomas Cook aircraft. This shows that they failed in step two and three of Endsley's model, the comprehension of the perceived elements and the projection in the future. Secondly the crew contacted ATC at t6 to have their descent clearance confirmed. Had it not been for SA difference 3, SA difference 2 could have been resolved here. So to summarise: there were 3 opportunities to resolve SA difference 2, of which only one was related to the TCAS system. Focussing only on the improper response to the TCAS does not cover two of the three contributing causes which are not mentioned in the official report.

SA difference 3: ATC was not aware of the loss of separation. The Thomas Cook Crew and the TCAS were aware of this fact.

As shown above, SA difference 3 contributed to SA difference 2. Had SA difference 3 been resolved, SA difference 2 could have been resolved. There was no exogenous cue to resolve this difference. The situation depended entirely on the controller updating his SA through reasoning: upon resolving SA difference 1, the controller might have reasoned how his previous decision based on SA difference 1 could have caused issues. In the short time span this was understandably too much to ask. Having an exogenous cue available to resolve difference 3 could have gone a long way in resolving the incident quicker. For example a downlink of the TCAS system so that the controller was alerted as well.

So, MASA shows how not all facets are captured in the official report. Assuming that the resolving of SA differences can prevent incidents, the analysis shows that the found causes and recommendations are not all encompassing.

5. Case 3: Attempted takeoff from wrong runway, Comair flight 5191, Lexington Kentucky, USA

For the final case to be analysed with multi agent situation awareness in mind, the attempted takeoff from a wrong runway in Lexington airport will be used. There is no component failure, and it can be argued that a difference in SA was the direct cause for the tragic accident.

Skybrary summarizes the accident as follows (Skybrary, C):

"On August 27, 2006, Comair flight 5191 crashed during takeoff from Blue Grass Airport in Lexington, Kentucky. The flight had been cleared for takeoff on runway 221; however, the flightcrew inadvertently taxied onto runway 262 and ran off the end of the runway during the takeoff roll. The airplane was destroyed and of the 47 passengers and 3 flightcrew members aboard the airplane, 49 were killed and 1 received serious injuries."

In the months following the accident the National Transportation Safety Board (NTSB) did a thorough investigation. Their findings are summarised below (National Transportation Safety Board, 2007):

"The National Transportation Safety Board determines that the probable cause of this accident was the flight crewmembers' failure to use available cues and aids to identify the airplane's location on the airport surface during taxi and their failure to cross-check and verify that the airplane was on the correct runway before takeoff. Contributing to the accident were the flight crew's nonpertinent conversation during taxi, which resulted in a loss of positional awareness, and the Federal Aviation Administration's failure to require that all runway crossings be authorized only by specific air traffic control clearances."

Based on these findings the following recommendations were made:

- That all crewmembers are required to positively confirm the aircraft's location before take-off.
- That all operators install moving map displays on the aircraft or a system that alerts pilots when takeoff from a wrong runway is attempted.

- That all airports implement enhanced taxiway centreline markings and surface painted holding position signs at all runway entrances.
- That issuance of takeoff clearance should be prohibited until after the airplane has crossed all intersecting runways.
- That Air Traffic Control should refrain from performing administrative tasks when moving aircraft are in their area of responsibility.

Now a Multi Agent Situation Awareness analysis will be performed to see if these conclusions and recommendations corroborate with the new findings.

5.1. Multi Agent Situation Awareness Analysis of Case 3

First off the steps of the orientation phase are performed.

Identify likely SA differences

SA difference 1: The Comair crew thought the aircraft was at the hold short line for runway 22 while it was actually at the hold short line for runway 26.

SA difference 2: ATC thought the Comair aircraft was at the hold short line for runway 22 while it was actually at the hold short line for runway 26.

SA difference 3: The crew thought the runway lights were not functioning while in fact they were.

SA difference 4: The Comair crew thought the aircraft was taking off from runway 22 while it was actually taking off from runway 26.

SA difference 5: ATC thought the aircraft was taking off from runway 22 while it was actually taking off from runway 26.

SA difference 6: The crew was not aware that the aircraft needed to cross a runway. ATC was aware of this fact.

Identify likely root causes of the SA difference

SA difference 1: The crew failed to observe crucial signage that would have clued them about their whereabouts. Furthermore, they anticipated a “short” taxi run.

SA difference 2: The air traffic controller did not monitor the aircraft during its taxi operation.

SA difference 3: The first officer recalled how the lights were out at an earlier visit. This created the belief that they were still out of order. The NOTAM made no mention about the lights.

SA difference 4: The crew failed to observe crucial signage that would have clued them about their whereabouts. Furthermore they expected the runway lights to be off and were thus not alarmed when there were no lights.

SA difference 5: The air traffic controller did not monitor the aircraft during the take-off run.

SA difference 6: The crew failed to properly read the maps provided.

Identify likely consequences of the SA difference

SA difference 1: The crew assumed they were lining up on runway 22 when in fact they were positioned for runway 26.

SA difference 2: ATC assumed that the aircraft was lining up on runway 22 when in fact they were positioned for runway 26.

SA difference 3: The crew did not get alarmed by the fact that the runway was not illuminated.

SA difference 4: The crew took off from a runway that was too short.

SA difference 5: ATC could not prevent the crew from taking off from a runway that was too short.

SA difference 6: Because they did not expect to cross a runway, the crew did not feel alarmed when they had not crossed, resulting in SA difference 1 and 4.

Identify relevant agents in the scenario

From the identified differences in SA the following agents are found to be relevant to the scenario:

- The Lexington airport air traffic controller (*ATC)*
- The Comair aircraft (*com_ac*)
- The Comair first officer (*com_fo)*
- The Comair flight captain (com_cap)
- Runway lights (*lights)*

Identify likely non-nominal SA update types in the scenario

- Night time conditions potentially impeded visual SA updates.
- Changes in the airport layout might have impeded visual SA updates.
- The air traffic controller decided not to monitor the aircraft during the taxi run and the take-off preventing any visual SA updates during that time.

Identify likely sub-models in the scenario

Comair crew:

- Human error in maintaining SA about the aircraft position.
- First Officers cognitive mode possibly switched to opportunistic. The CVR captured the first officer yawning which indicated he might have been tired.
- Non pertinent conversation between the captain and the first officer during the taxi run might have hindered them in maintaining SA. Control mode potentially switched to scrambled.

ATC:

- The controllers mode was possibly switched to opportunistic. The controller indicated he was tired, however he also stated that he was “fine” and “alert”. He might have felt pressured to finish his administrative duties before the end of his shift.

Now, an ontology needs to be created that defines all the possible objects and events. The following tables show this. Table 9 shows all the identities for the relevant agents. Table 10 shows the relevant continuous state variables, as well as the relevant discrete state variables.

Table 9: The relevant agents in the third scenario

Agent#	Agent	Identity
1	Air Traffic Controller	atco
2	Comair captain	com_cap
3	Comair first officer	com_fo
4	Comair aircraft	com_ac
5	Runway lights	rw_lights

Table 10: The states of the relevant agents in the third scenario

Agent	Continuous states	Discrete states
Air Traffic Controller	*n.a.*	*n.a.*
Comair aircraft	hs_22, hs_26, rwy_22, rwy_26, tw, park, crossing	*cl_22, cl_26,*
Comair captain	*n.a.*	*n.a.*
Comair first officer	*n.a.*	*n.a.*
Runway lights	*n.a.*	lights_on, lights_off

nb: hs_22 and hs_26 refer to the hold short lines for runways 22 and 26 respectively, whereas rwy_22 and rwy_26 refer to the runways themselves. *Tw* stands for taxiway and *park* denotes the spot the aircraft is parked at the beginning. *Crossing* refers to the point where the taxiway crosses runway 26. The discrete states of the aircraft pertain to the belief of other agents about the runway for which the aircraft has clearance to take off from. So if the discrete state of the aircraft as perceived by the Comair captain is *cl_22*, this means that he holds the belief that he is to take-off from runway 22.

The following timestamps have been selected as moments in the narrative where SA updates likely occurred:

t0: At 05:48 the Automatic Terminal Information System (ATIS) broadcast is received in the cockpit. The most relevant information here is that the active runway is runway 22.

t1: Less than a minute later at 05:49 the air traffic controller is contacted by the crew. Their intention to depart for Atlanta is let known as well as their acknowledgement of the aforementioned ATIS information.

t2 : At 05:56 the Cockpit Voice Recorder (CVR) shows that the first officer's SA was either not updated properly during the ATIS broadcast or that it had already degraded, as he seems unsure about the runway clearance and he assumes that it is runway 24. The captain points out that runway 22 is the correct runway.

t3 : A few seconds later the first officer makes a remark about the ATIS mentioning that the glideslope for the runway is out of service. This prompts him to recall that the lights were not functioning "the other night".

t4 : At 05:57 the first officer states: "two two's a short taxi" indicating his belief about the travel distance to the runway.

t5: At approximately 06:04 the FDR shows that the aircraft stops. This is most likely upon reaching the hold short line for runway 26. This turned out to be an erroneous update cue.

t6: At 06:05 the accident flight is granted permission for takeoff.

t7: The accident flight starts moving again, and as the follows the centreline leading to runway 26 the SA of the crew members is updated.

t8: The captain exclaims "V one, rotate" even though FDR indicates that these speeds had not been reached yet. This prompts the assumption that the captain had realized that something was wrong. It was likely the end of the runway functioning as a visual cue.

Timestamp t0

The ATIS broadcast is received in the cockpit. The following assumptions are made:

A1: Both the crewmembers have their intent SA regarding the takeoff clearance updated from unknown to *cl_22*.

A2: At this point the crew assumed nominal condition, i.e: runway lights working properly.

A3: The crew is homogenous from the outside, and their SA is perceived as such from the outside.

The SA for the relevant agents at this point can be seen in table 11. The intent SA is found in table 12.

Table 11: The state vectors of the relevant agents at t0

	ATC	Comair Captain	Comair First Officer	Comair a/c	Runway lights
ATC	$\sigma_{t0,1}^{1} = \begin{pmatrix} atco \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,1}^{2:3} = \begin{pmatrix} com\_crew \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$		$\sigma_{t0,1}^{4} = \begin{pmatrix} com\_ac \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,1}^{5} = \begin{pmatrix} rw\_lights \\ \{\cdot\} \\ lights\_on \\ \hat{v}_{t,k}^{j} \end{pmatrix}$
Comair captain	$\sigma_{t0,2}^{1} = \begin{pmatrix} atco \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,2}^{2} = \begin{pmatrix} com\_cap \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,2}^{3} = \begin{pmatrix} com\_fo \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,2}^{4} = \begin{pmatrix} com\_ac \\ park \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,2}^{5} = \begin{pmatrix} rw\_lights \\ \{\cdot\} \\ lights\_on \\ \hat{v}_{t,k}^{j} \end{pmatrix}$
Comair first officer	$\sigma_{t0,3}^{1} = \begin{pmatrix} atco \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,3}^{2} = \begin{pmatrix} com\_cap \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,3}^{3} = \begin{pmatrix} com\_fo \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,3}^{4} = \begin{pmatrix} com\_ac \\ park \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\sigma_{t0,3}^{5} = \begin{pmatrix} rw\_lights \\ \{\cdot\} \\ lights\_on \\ \hat{v}_{t,k}^{j} \end{pmatrix}$
Comair a/c	$\{\cdot\}$	$\{\cdot\}$		$\sigma_{t0,4}^{4} = \begin{pmatrix} com\_ac \\ park \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$	$\{\cdot\}$
Runway lights	$\{\cdot\}$	$\{\cdot\}$		$\{\cdot\}$	$\sigma_{t0,5}^{5} = \begin{pmatrix} rw\_lights \\ \{\cdot\} \\ lights\_on \\ \hat{v}_{t,k}^{j} \end{pmatrix}$

Table 12: The intent vectors for the relevant agents at t0

	ATC	Comair Captain	Comair First Officer	Comair a/c	Runway lights
ATC	$(v^1_{t0,1})_1 = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix}$	$(v^{2:3}_{t0,1})_1 = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix}$		$(v^4_{t0,1})_1 = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix}$	$\{\cdot\}$
Comair captain	$(v^1_{2t0,2})_1 = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix}$	$(v^2_{t0,2})_1 = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix}$	$(v^3_{t0,2})_1 = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix}$	$(v^4_{t0,2})_1 = \begin{pmatrix} \{\cdot\} \\ hs\_22 \\ t_{hs22} \end{pmatrix}$ $(v^4_{t0,2})_2 = \begin{pmatrix} cl\_22 \\ hs\_22 \\ t_{clearance} \end{pmatrix}$ $(v^4_{t0,2})_3 = \begin{pmatrix} cl\_22 \\ rwy\_22 \\ t_{clearance} + \delta \end{pmatrix}$	$\{\cdot\}$
Comair first officer	$(v^1_{t0,3})_1 = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix}$	$(v^2_{t0,3})_1 = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix}$	$(v^3_{t0,3})_1 = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix}$	$(v^4_{t0,3})_1 = \begin{pmatrix} \{\cdot\} \\ hs\_22 \\ t_{hs22} \end{pmatrix}$ $(v^4_{t0,3})_2 = \begin{pmatrix} cl\_22 \\ hs\_22 \\ t_{clearance} \end{pmatrix}$ $(v^4_{t0,3})_3 = \begin{pmatrix} cl\_22 \\ rwy\_22 \\ t_{clearance} + \delta \end{pmatrix}$	$\{\cdot\}$

Timestamp t1

Here the SA of ATC is updated. The information in the ATIS is acknowledged. It is also assumed that somewhere in between t1 and t2 the SA of the first officer degraded to the point that he no longer knew what runway was in use. The changes are:

- Via the communication function the attention of ATC is drawn towards the aircraft. The identity and continuous state is updated to *com_ac* and *park* respectively.

$$\sigma^4_{t0,1} = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix} \xrightarrow{communication} \sigma^4_{t1,1} \begin{pmatrix} com\_ac \\ park \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix} \qquad \text{(update)}$$

- Via the reasoning function the intent SA of ATC about the Comair aircraft is updated. At time $t_{crossing}$ the aircraft's continuous state will be *crossing*.

$$\left(v^4_{t0,1}\right)_1 = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix} \xrightarrow{reasoning} \left(v^4_{t1,1}\right)_1 = \begin{pmatrix} \{\cdot\} \\ crossing \\ t_{crossing} \end{pmatrix} \text{(update)}$$

- Via the reasoning function the intent SA of ATC about the Comair aircraft is updated. At time t_{hs22} the aircraft's continuous state will be hs_22.

$$\left(v^4_{t0,1}\right)_2 = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix} \xrightarrow{reasoning} \left(v^4_{t1,1}\right)_1 = \begin{pmatrix} \{\cdot\} \\ hs\_22 \\ t_{hs22} \end{pmatrix} \qquad \text{(update)}$$

- Via the reasoning function the intent SA of ATC about the Comair aircraft is updated. At time $t_{clearance}$ the aircraft's discrete state will be *clearance_22.*

$$\left(v^4_{t0,1}\right)_3 = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix} \xrightarrow{reasoning} \left(v^4_{t1,1}\right)_2 = \begin{pmatrix} cl\_22 \\ hs\_22 \\ t_{clearance} \end{pmatrix} \text{(update)}$$

- Via the reasoning function the intent SA of ATC about the Comair aircraft is updated. At time $t_{clearance} + \delta$ the aircraft's continuous state will be rwy*_22.*

$$\left(v_{t0,1}^{4}\right)_{4} = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix} \xrightarrow{reasoning} \left(v_{t1,1}^{4}\right)_{3} = \begin{pmatrix} cl\_22 \\ rwy\_22 \\ t_{clearance} + \delta \end{pmatrix}$$

- The First officer forgot that the active runway was 22. Thus his intent vector for the aircraft was updated from *clearance_22* to unknown.

$$\left(v_{t0,3}^{4}\right)_{1} = \begin{pmatrix} cl\_22 \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix} \xrightarrow{degredation} \left(v_{t1,3}^{4}\right)_{1} = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix} \quad \text{(update)}$$

While the last update is a difference in SA it is decided not to take it into consideration for the analysis because it was so quickly repaired by the captain.

Timestamp t2

The captain points out to the first officer that runway 22 is the runway currently in use upgrading his SA as follows:

- The intent SA of the first officer is once again updated. This time from *unknown* to *clearance_22*

$$\left(v_{t1,3}^{4}\right)_{1} = \begin{pmatrix} \{\cdot\} \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix} \xrightarrow{communication} \left(v_{t2,3}^{4}\right)_{1} = \begin{pmatrix} cl\_22 \\ \{\cdot\} \\ \{\cdot\} \end{pmatrix} \qquad \text{(update)}$$

Timestamp t3

The first officer makes a remark about the runway lights being out of order earlier, possibly casting doubt about their current state.

- The state SA of the captain about the runway lights is updated. The discrete state is set from *lights_on* to *unknown.*

$$\sigma_{t2,2}^{5} = \begin{pmatrix} rw\_lights \\ \{\cdot\} \\ lights\_on \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{communication} \sigma_{t3,2}^{5} = \begin{pmatrix} rw\_lights \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix} \qquad \text{(update)}$$

- The state SA of the first officer about the runway lights is updated. The discrete state is set from *lights_on* to *unknown.*

$$\sigma_{t2,3}^{5} = \begin{pmatrix} rw\_lights \\ \{\cdot\} \\ lights\_on \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{reasoning} \sigma_{t3,3}^{5} = \begin{pmatrix} rw\_lights \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix} \qquad \text{(update)}$$

Timestamp t4

The phrase "two two's a short taxi" uttered by the first officer potentially decreases the time expected by both the crew members for the aircraft to attain the continuous state *hs_22.* Previously, as seen in table 12, this time was denoted by t_{hs22} . The new expected time that the state will be attained is denoted by $t_{hs22}b$ where $t_{hs22}b$ is an earlier time point than t_{hs22}. The aircraft proceeded to do the taxi run. The following changes occur:

- Through reasoning the intent SA of the captain about the aircraft is updated. The time for which he expects the aircraft to reach the continuous state *hs_22* is updated from t_{hs22} to $t_{hs22}b$.

$$\left(v_{t3,2}^{4}\right)_1 = \begin{pmatrix} \{\cdot\} \\ hs\_22 \\ t_{hs22} \end{pmatrix} \xrightarrow{reasoning} \left(v_{t4,2}^{4}\right)_1 = \begin{pmatrix} \{\cdot\} \\ hs\_22 \\ t_{hs22}b \end{pmatrix} \qquad \text{(update)}$$

- Through reasoning the intent SA of the first officer about the aircraft is updated. The time for which he expects the aircraft to reach the continuous state *hs_22* is updated from t_{hs22} to $t_{hs22}b$.

$$\left(v_{t3,3}^{4}\right)_1 = \begin{pmatrix} \{\cdot\} \\ hs\_22 \\ t_{hs22} \end{pmatrix} \xrightarrow{reasoning} \left(v_{t4,3}^{4}\right)_1 = \begin{pmatrix} \{\cdot\} \\ hs\_22 \\ t_{hs22}b \end{pmatrix} \qquad \text{(update)}$$

- The state SA of the aircraft about itself was updated. The continuous state changed from *park* to *tw.*

$$\sigma_{t3,4}^{4} = \begin{pmatrix} com\_ac \\ park \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix} \longrightarrow \sigma_{t4,4}^{4} = \begin{pmatrix} com\_ac \\ tw \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$$

- The state SA of the captain about the aircraft was updated. The continuous state changed from *park* to *tw.*

$$\sigma_{t3,2}^{4} = \begin{pmatrix} com\_ac \\ park \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{observation} \sigma_{t4,2}^{4} = \begin{pmatrix} com\_ac \\ tw \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix} \qquad \text{(update)}$$

- The state SA of the first officer about the aircraft was updated. The continuous state changed from *park* to *tw.*

$$\sigma_{t3,3}^{4} = \begin{pmatrix} com\_ac \\ park \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{observation} \sigma_{t4,3}^{4} = \begin{pmatrix} com\_ac \\ tw \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix} \qquad \text{(update)}$$

Timestamp t5

The aircraft had reached the hold short line for runway 26. The markings were visible and illuminated as was made clear during the taxi demonstration that was performed the day after the accident (page 25 of the accident report). However, the cue was wrongly interpreted: it is assumed that the crew members thought that they had now reached the hold short line for runway 22 when in fact they were at the hold short for runway 26. In the crew's interpretation t5 had become $t_{hs22}b$ and thus the time at which the first intent vector for the crew agents is attained had been reached:

- The state SA of the aircraft about itself was updated. The continuous state changed from *tw* to *hs_26.*

$$\sigma_{t4,4}^{4} = \begin{pmatrix} com\_ac \\ tw \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix} \longrightarrow \sigma_{t5,4}^{4} = \begin{pmatrix} com\_ac \\ hs\_26 \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix}$$

- Through reasoning the first intent SA vector of the captain about the aircraft was updated. The expected time for the aircraft to reach the state *hs_22* changed from $t_{hs22}b$ to t5.

$$\left(v_{t4,2}^{4}\right)_1 = \begin{pmatrix} \{\cdot\} \\ hs\_22 \\ t_{hs22}b \end{pmatrix} \xrightarrow{reasoning} \left(v_{t5,2}^{4}\right)_1 = \begin{pmatrix} \{\cdot\} \\ hs\_22 \\ t5 \end{pmatrix} \qquad \text{(update)}$$

- Through reasoning and observation the state SA of the captain about the aircraft was updated. The continuous state changed from *tw* to *hs_22*.

$$\sigma^4_{t4,2} = \begin{pmatrix} com\_ac \\ tw \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix} \xrightarrow{reasoning\ /\ observation} \sigma^4_{t5,2} = \begin{pmatrix} com\_ac \\ hs\_22 \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix} \qquad \text{(update)}$$

- Through reasoning the first intent SA vector of the first officer about the aircraft was updated. The expected time for the aircraft to reach the state *hs_22* changed from $t_{hs22}b$ to t5.

$$\left(v^4_{t4,3}\right)_1 = \begin{pmatrix} \{\cdot\} \\ hs\_22 \\ t_{hs22}b \end{pmatrix} \xrightarrow{reasoning} \left(v^4_{t5,3}\right)_1 = \begin{pmatrix} \{\cdot\} \\ hs\_22 \\ t5 \end{pmatrix} \qquad \text{(update)}$$

- Through reasoning and observation the state SA of the first officer about the aircraft was updated. The continuous state changed from *tw* to *hs_22*.

$$\sigma^4_{t4,3} = \begin{pmatrix} com\_ac \\ tw \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix} \xrightarrow{observation\ /\ reasoning} \sigma^4_{t5,3} = \begin{pmatrix} com\_ac \\ hs\_22 \\ \{\cdot\} \\ \hat{v}^j_{t,k} \end{pmatrix} \qquad \text{(update)}$$

N.B: Since the captain was doing the taxiing, it is not known whether the first officer had also observed the markings for the hold short line. It is assumed that he had not as he was busy with the checklist, and that his SA was updated via the reasoning function: observing that the aircraft had halted, he must have assumed that the correct runway was reached.

Timestamp t6

The crew contacts ATC and notifies him that they are ready to go. Takeoff clearance is granted. The time of the second intent vector for the crew is attained. Furthermore, ATC now assumes that the aircraft was at the hold short line for runway 22:

- The state SA of the aircraft about the aircraft was updated to reflect the granted clearance.

$$\sigma_{t5,4}^{4} = \begin{pmatrix} com\_ac \\ hs\_26 \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix} \longrightarrow \sigma_{t6,4}^{4} = \begin{pmatrix} com\_ac \\ hs\_26 \\ cl\_22 \\ \hat{v}_{t,k}^{j} \end{pmatrix}$$

- Through reasoning the second intent vector of the captain about the aircraft was updated. The time for which the state *clearance_22* was expected to be attained changed from $t_{clearance}$ to *t6*.

$$\left(v_{t5,2}^{4}\right)_2 = \begin{pmatrix} cl\_22 \\ hs\_22 \\ t_{clearance} \end{pmatrix} \xrightarrow{reasoning} \left(v_{t5,2}^{4}\right)_2 = \begin{pmatrix} cl\_22 \\ hs\_22 \\ t6 \end{pmatrix} \qquad \text{(update)}$$

- Through reasoning the third intent vector of the captain about the aircraft was updated. The time for which the state *rwy_22* was expected to be attained changed from $t_{clearance} + \delta$ to *t6* $+ \delta$.

$$\left(v_{t5,2}^{4}\right)_3 = \begin{pmatrix} cl\_22 \\ rwy\_22 \\ t_{clearance} + \delta \end{pmatrix} \xrightarrow{reasoning} \left(v_{t5,2}^{4}\right)_3 = \begin{pmatrix} cl\_22 \\ rwy\_22 \\ t6 + \delta \end{pmatrix} \qquad \text{(update)}$$

- Through reasoning the second intent vector of the first officer about the aircraft was updated. The time for which the state *clearance_22* was expected to be attained changed from $t_{clearance}$ to *t6*.

$$\left(v_{t5,3}^{4}\right)_2 = \begin{pmatrix} cl\_22 \\ hs\_22 \\ t_{clearance} \end{pmatrix} \xrightarrow{reasoning} \left(v_{t5,3}^{4}\right)_2 = \begin{pmatrix} cl\_22 \\ hs\_22 \\ t6 \end{pmatrix} \qquad \text{(update)}$$

- Through reasoning the third intent vector of the first offcicer about the aircraft was updated. The time for which the state *rwy_22* was expected to be attained changed from $t_{clearance} +$ δ to *t6* $+ \delta$.

$$\left(v_{t5,3}^{4}\right)_3 = \begin{pmatrix} cl\_22 \\ rwy\_22 \\ t_{clearance} + \delta \end{pmatrix} \xrightarrow{reasoning} (v_{t532}^{4})_3 = \begin{pmatrix} cl\_22 \\ rwy\_22 \\ t6 + \delta \end{pmatrix} \qquad \text{(update)}$$

- The state SA of the captain about the aircraft was updated to reflect the granted clearance.

$$\sigma_{t5,2}^{4} = \begin{pmatrix} com\_ac \\ hs\_22 \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{communication} \sigma_{t6,2}^{4} = \begin{pmatrix} com\_ac \\ hs\_22 \\ cl\_22 \\ \hat{v}_{t,k}^{j} \end{pmatrix} \qquad \text{(update)}$$

- The state SA of the first officer about the aircraft was updated to reflect the granted clearance.

$$\sigma_{t5,3}^{4} = \begin{pmatrix} com\_ac \\ hs\_22 \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{communication} \sigma_{t6,3}^{4} = \begin{pmatrix} com\_ac \\ hs\_22 \\ cl\_22 \\ \hat{v}_{t,k}^{j} \end{pmatrix} \qquad \text{(update)}$$

- The state SA of ATC was updated through reasoning. The continuous state changed from *unknown* to *hs_22.* The discrete state was updated to *clearance_22*.

$$\sigma_{t5,1}^{4} = \begin{pmatrix} com\_ac \\ \{\cdot\} \\ \{\cdot\} \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{reasoning} \sigma_{t6,1}^{4} = \begin{pmatrix} com\_ac \\ hs\_22 \\ cl\_22 \\ \hat{v}_{t,k}^{j} \end{pmatrix} \quad \text{(update)}$$

- Through reasoning the third intent SA vector of ATC about the aircraft was updated to reflect the time at which the aircraft would attain the continuous state *crossing.*

$$\left(v_{t5,1}^{4}\right)_{3} = \begin{pmatrix} cl\_22 \\ rwy\_22 \\ t_{clearance} + \delta \end{pmatrix} \xrightarrow{reasoning} \left(v_{t6,1}^{4}\right)_{2} = \begin{pmatrix} cl\_22 \\ crossing \\ t6 + \delta \end{pmatrix} \qquad \text{(update)}$$

Timestamp t7

After the hold short line for runway 26 the taxi centreline split into three directions. The crew followed the left, which arched unto runway 26. This prompted them to update their SA.

- The state SA of the aircraft about itself was updated to reflect the new position.

$$\sigma_{t6,4}^{4} = \begin{pmatrix} com\_ac \\ hs\_26 \\ cl\_22 \\ \hat{v}_{t,k}^{j} \end{pmatrix} \longrightarrow \sigma_{t7,4}^{4} = \begin{pmatrix} com\_ac \\ crossing \\ cl\_22 \\ \hat{v}_{t,k}^{j} \end{pmatrix}$$

- Through observation the state SA of the captain about the aircraft was updated. The continuous state changed from *hs_22* to *rwy_22.*

$$\sigma_{t6,2}^{4} = \begin{pmatrix} com\_ac \\ hs\_22 \\ cl\_22 \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{observation} \sigma_{t7,2}^{4} = \begin{pmatrix} com\_ac \\ rwy\_22 \\ cl\_22 \\ \hat{v}_{t,k}^{j} \end{pmatrix} \text{ (update)}$$

- Through observation the state SA of the first officer about the aircraft was updated. The continuous state changed from *hs_22* to *rwy_22.*

$$\sigma_{t6,3}^{4} = \begin{pmatrix} com\_ac \\ hs\_22 \\ cl\_22 \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{observation} \sigma_{t7,3}^{4} = \begin{pmatrix} com\_ac \\ rwy\_22 \\ cl\_22 \\ \hat{v}_{t,k}^{j} \end{pmatrix} \text{ (update)}$$

- Through reasoning the state SA of ATC about the aircraft was updated. The continuous state changed from hs_22 to rwy_22.

$$\sigma_{t6,1}^{4} = \begin{pmatrix} com\_ac \\ hs\_22 \\ cl\_22 \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{reasoning} \sigma_{t7,1}^{4} = \begin{pmatrix} com\_ac \\ rwy\_22 \\ cl\_22 \\ \hat{v}_{t,k}^{j} \end{pmatrix} \text{ (update)}$$

Timestamp t8

The exclamation by the captain "V one, rotate" which was made too soon, suggests a last minute realization that something was not correct where the end of the runway was possibly functioning as a cue. It is assumed that the crew did not have enough time to fully realize the extent of the error. It is therefore decided that timestamp t8 led to a loss of positional SA rather than an update. It is assumed that this last minute realization holds for both crewmembers.

- Through observation the state SA of the captain about the aircraft was updated. The continuous state changed from *rwy_22* to *unknown.*

$$\sigma_{t7,2}^{4} = \begin{pmatrix} com\_ac \\ rwy\_22 \\ cl\_22 \\ \hat{v}_{t,k}^{j} \end{pmatrix} \xrightarrow{observation} \sigma_{t8,2}^{4} = \begin{pmatrix} com\_ac \\ \{\cdot\} \\ cl\_22 \\ \hat{v}_{t,k}^{j} \end{pmatrix} \text{ (update)}$$

- Through observation the state SA of the first officer about the aircraft was updated. The continuous state changed from *rwy_22* to *unknown.*

$$\sigma^4_{t7,3} = \begin{pmatrix} com\_ac \\ rwy\_22 \\ cl\_22 \\ \hat{v}^j_{t,k} \end{pmatrix} \xrightarrow{observation} \sigma^4_{t8,3} = \begin{pmatrix} com\_ac \\ \{\cdot\} \\ cl\_22 \\ \hat{v}^j_{t,k} \end{pmatrix} \quad \text{(update)}$$

Propagation of SA differences

The above gives clues about the propagation of the SA differences and their relations. It is quite clear that the fact that the crew did not know they had a runway crossing on their taxi route, SA difference 6, was a contributing cause for them to believe they were at the hold short line for runway 22 while they were actually at the hold short line for runway 26 (SA difference 1). This difference evidently propagated to form SAD 2 as the controller relied solely on the communication from the crew to ascertain their location, thus the controller believed the aircraft was at the hold short line for runway 22 as well. Difference 1 also contributed to difference 4 as the crew logically assumed they were now at runway 22. Difference 2 went on to cause difference 5 as the controller only relied on reasoning and in that case difference 5 follows logically from difference 2. Had the controller used another cue to resolve difference 2, it is likely that difference 1 would be resolved as well. Same goes for difference 5 and 4. Finally difference 3 was a contributing cause to difference 4. Figure 5 shows a schematic representation of the SA difference relations.

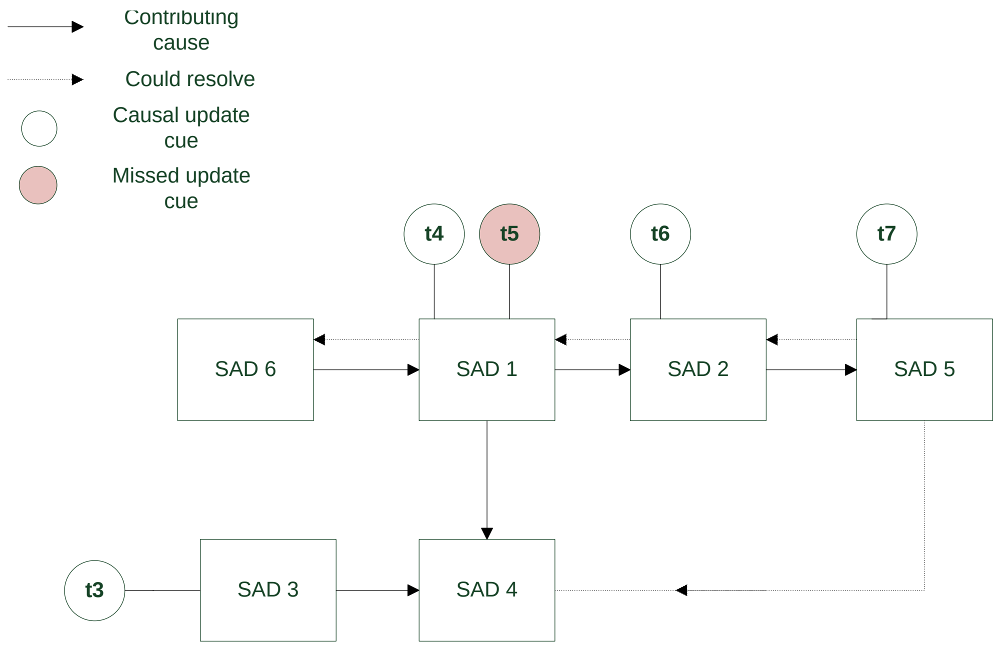

Figure 5: SA difference propagation for case 3

5.2. Results from the analysis of case 3

It is shown that a difference in SA can be considered a cause of an accident. In an accident like the Commair crash it is quite intuitive to gravitate towards *human error* as the reason it occurred. This is also the conclusion of the official accident report (National Transportation Safety Board, 2007) which implicitly blames the accident on failure of the crew to maintain SA. So how do the findings of the multi agent situation awareness, where SA is explicitly analysed, compare to the official findings? The found differences in SA are revisited.

SA difference 1: The Comair crew thought the aircraft was at the hold short line for runway 22 while it was actually at the hold short line for runway 26.

The official report states that the accident was caused by the crew's failure to use the available cues to verify their location during taxiing. From 06.03:16 to 06.03:56 the CVR recorded non pertinent conversation between the captain en the first officer. The report cites this as a contributing cause to the accident. The reason for this line of thinking is quite clear: the conversation potentially impeded both the agents' SA update processes leading to a difference in SA. This implies however that during the time of the conversation some cues were missed. So which cues are these? The report fails to specify but presumably the investigators referred to airport signage. However, the taxi demonstration performed by the investigators shows that arguably the most salient visual cues were still available when the aircraft had reached the hold short line for runway 26, i.e when the non-pertinent conversation had already ended. For one, the runway hold sign for runway 26 was visible to the left of the plane when it was waiting for take-off clearance. Secondly, the red and white runway holding position sign for runway 22 (which was located on taxiway A on the north side of runway 26) was visible. The sign was illuminated. The investigators assume that the crew missed some signs due to the conversation, but that does not explain why these two cues were missed, even though the conversation had ended. It casts some doubt on the investigators' conclusion that the non-pertinent conversation weighed so heavily on the loss of positional awareness.

Thinking in terms of multi agent SA one can imagine another type of cue that might have been missed due to the non-pertinent conversation. Being engrossed, the crew might have not noticed that they had not crossed runway 26 while taxiing. This however means that the crew should have expected a runway crossing. The fact that the taxi briefing was abbreviated suggests that the crew did not expect a crossing, resulting in SA difference 6. Had SA difference 6 not occurred or had it been resolved, SA difference 1 would have not occurred. This all shows that the sterile cockpit rule should be upheld without a doubt and that the crew broke this rule with their conversation. On the other hand, the above shows that the investigators might have put too much emphasis on this element since it shows that the crew did not pass any salient cues about their whereabouts during the time their conversation took place.

SA difference 2: ATC thought the Comair aircraft was at the hold short line for runway 22 while it was actually at the hold short line for runway 26.

The official report recommends that ATC should refrain from administrative tasks while moving aircraft are in their area of responsibility. The MASA analysis findings corroborate this. There was only one exogenous cue available for ATC, and that was the crew of the aircraft notified ATC that they were ready for take-off. This triggered an update cue through reasoning that resulted in the SA difference. Had ATC not been busy with administrative tasks, he might have used one of the other available cues, like simply looking outside of the cabin window. MASA shows that SA difference 1 caused SA difference 2, but also that if SA difference 2 was resolved, SA difference 1 would likely be resolved as well.

SA difference 3: The crew thought the runway lights were not functioning while in fact they were.

This SA difference does not correspond to any of the causes nor is it addressed by any of the recommendations in the NTSB report. By contrast, MASA shows that it can in fact be considered a contributing cause. It is not unlikely that resolving SA difference 3 could have resolved SA difference 4. During the aligning of the aircraft, and during take-off the lack of lights did not alarm the crew. Had they expected the lights to be in order like they were, the lack of lights could have casted doubt about their position. It could have prompted the crew to contact ATC to confirm their location and subsequently it could have prevented the accident.

The SA difference occurred when the first officer was reading back the NOTAM which stated that the glideslope for the runway was out. This prompted him to recall that on his previous visit to the airport the runway lights were out as well. So the difference did not occur due to a mistake in the NOTAM but rather due to an assumption by the crew. A recommendation for this finding could be that the reliability of the NOTAM should be increased and that crewmembers should have enough trust to base their awareness on the NOTAM instead of previous experiences.

SA difference 4: The Comair crew thought the aircraft was taking off from runway 22 while it was actually taking off from runway 26.

SA difference 5: ATC thought the aircraft was taking off from runway 22 while it was actually taking off from runway 26.

SA difference 4 and 5 can be considered the most proximate SA differences and a direct cause to the accident. This finding corroborates with the official findings. The differences are a result of SA difference 1 and 2 and it was already shown how resolving SA difference 3 could have resolved SA differences 4 and 5.

SA difference 6: The crew was not aware that the aircraft needed to cross a runway. ATC was aware of this fact.

The last SA difference corresponds to the last finding of the NTSB, namely that the Federal Aviation Administration's failed to require that all runway crossings are authorized only by specific air traffic control clearances. MASA shows how an extra exogenous cue, where the crew would specifically receive clearance for a runway crossing, could have resolved SA difference 6 which would in turn resolve SA difference 1, which would in turn resolve all other SA differences eventually preventing the accident.

6. The Comair accident and other retrospective analysis techniques

There exists two (partial) analyses of the Comair accident performed with other retrospective analysis techniques. The first one is FRAM, or functional resonance accident modelling. The second one is STAMP, System-Theoretic Accident Model and Processes. Both techniques will shortly be explained followed by a comparison of the results.

6.1. FRAM

FRAM, or functional resonance accident modelling, is a qualitative ,systemic accident investigation method based on resilience engineering (Holnagel, Pruchnicki, Woltjer, & Etcher, 2008). It treats failure as a consequence of *functional resonance*, meaning that performance of the system is variable and that this variation can resonate throughout the rest of the system. This variability depends on the nature of the functions and the context.

The following four steps are performed:

1. Identify essential system functions with nominal performance as a baseline. The functions are characterized by their *input, output, preconditions, resources, time* and *control.*
2. Characterize the observed variability of system functions. This step entails finding out in what way the identified system functions can vary in performance. These variations can be attributed to one of the six characteristics mentioned in step one.
3. Identify and describe the functional resonance. First the dependencies of the functions are observed, meaning that the output of one function may be the input of a second function. Once these dependencies are mapped the progression of variabilities throughout the system, i.e. resonance, can be described.
4. Identify barriers for variability. Barriers are those functions and constraints that serve as damping factors and prevent an unwanted event.

A graphical representation of the interdependencies can be made by the use of so called FRAM modules where a function is represented by a hexagon and every corner is a characteristic of that function. See figure 6.

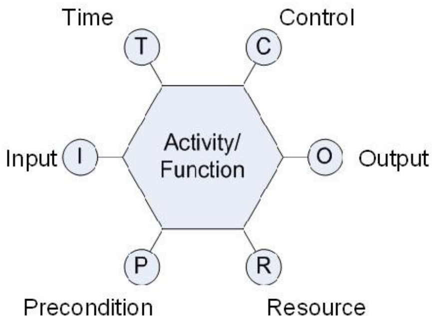

Figure 6: A generalized FRAM module

A (partial) FRAM analysis was performed on the Comair accident in Holnagel, Pruchnicki, Woltjer, & Etcher (2008). The graphical representation can be seen in appendix A.

The following points show how a more in dept understanding of SA differenced could have benefitted the FRAM analysis.

- If the function "Taxi to runway" is observed closer, it can be seen that the authors did incorporate the concept of Situation Awareness into the analysis despite Hollnagel's dismissal of the construct. The "control" part of the function module states that the non pertinant conversation of the crew and the incomplete NOTAM led to the variability of the function. While the inclusion of the changed taxiway in the NOTAM could have prompted an update of the SA through the "Lack of SA" cue, the non pertinent conversation is unlikely to have contributed to the variability. As is shown in chapter 5.2 during the time this conversation took place there were no update cues that could have aided in the pilot's situation awareness. A case can be made that the non pertinent conversation prevented an update cue of the form "lack of SA" that would occur if either crewmember realized that there was no runway crossing while there should have been one. The question remains however wheter the crew was explicitly aware of a required crossing as this was never vocalized by either because the taxi briefing was abbreviated.

- The FRAM analysis also mentions the expectation of low lightning. The expectation of low lightning is presented as an output of the "Taxi briefing" as well as the "Perform take-off checklist" functions. It is indeed true that the first officer remarked that "lights were out all over the place the other night" during the taxi briefing (National Transportation Safety Board, 2007). The FRAM analysis however does not demarcate the two crewmembers. In not doing so it does not show the complete picture of how the error in SA progressed. The MASA analysis shows the likely reasons the expectations occurred and that it occurred for different reasons in both the crew members. For the first officer a protocol enforced update cue was found: while doing the taxi briefing he mentions that the glideslope is out of function which prompts him to recall that on a prior flight at the same airport the lights were out. He vocalizes this recollection and through that the expectation by the captain of the lights being possibly out of functions is created as well. This demarcation become important when looking at preventive measures which could either focus on how the expecations arose for the first officer or how it arose for the captain. The result is the same in both the analyses: the expectation of reduced runway lightning is set as the controll part of the “Taxi to runway” and “Take-off” functions in the Fram analysis, implying that due to the expectation the crew could not reflect on their position. The MASA analysis treats this as a missed update cue that could have prompted the crew to abort the takeoff.

- The authors also make note of the expecations that the crew might have had for the lenght of the taxi run. The expectation of a short taxi run is mentioned, and MASA analysis concludes as well that this has contributed to the error in SA. But like in the previous point the FRAM analysis does not show how this expecation arose whereas the MASA analysis indicates what cues triggered the SA updates showing more exactly what preventive meassures should focus on.

Overall it can be stated that a MASA analysis shows at a more precise level how errors occurred and progressed through the system. It is easier to see at a glance which mechanisms allowed the system to fail and it shows ways to improve it by pointing out possible missed update cues. Therefore a MASA analysis would be a great precursor to a FRAM analysis. Three ways in which MASA methodology can be used in conjuncture with FRAM are proposed.

SA as a function

“Maintaining SA” can be considered a function within any system. The six characteristics of the function fit well into the mould of MASA:

Input: this can be the identified information flow in the scenario leading to updates in SA.
Preconditions: the preconditions of the functions is the perceptibility of the cues and the receptiveness of the agents like mental modes or technical barriers.
Resources: the identified update cues can be considered a resource.
Control: control is the situation awareness of the agent about himself.

Time: like in any systems function time constraints can play a role in maintaining SA. The correct awareness needs to be attained before the time it is needed.

As for the *output* of the function, this is naturally Situation Awareness and it can be tied in to other functions by either making it their resource, input, precondition or control depending on the function. For example, if the function is “taxi to runway” Situation Awareness about the position can be considered a resource or precondition of that function. If the function is “abort takeoff” the awareness of another aircraft on the runway is the input.

Instead of treating the SA function as a block box a MASA analysis can be used to better understand the ultimate causes of an event.

SA as a precondition

A second way in which the two methods can enhance each other is when correct SA is considered a precondition for other functions. The upside of this line of thinking compared to SA as a function is that it reduces the complexity of graphical representation of the analysis. They way “SA as a function” is described above, it requires that all SA elements are represented as a separate function. This would mean a myriad of functions like “Agent A maintaining awareness of the discrete state of agent B”. Whereas, if you consider it a precondition it suffices to mention a lack of SA or incorrect SA.

SA for characterizing the variability of functions

The second step of the FRAM analysis is the characterization of the performance variability. The FRAM methodology identifies several performance conditions that affect the variability. These can be assigned a rating: stable or variable but adequate, stable or variable but inadequate and unpredictable. However, there appear to be no standards for these ratings and imposing standards can prove to be quite hard. Because of this classifying variability can be highly subjective. Having a proper view of the SA at every point in the narrative can at least partially abate the subjectivity in performance variability: if variability is considered as lapses in SA a MASA analysis can be used as a standard.

6.2. STAMP

Several factors including an ever increasing complexity and intercoupling of human machine interfaces due to the fast pace of technological change have led Nancy Leveson to believe that traditional event chain models pose a serious limitation on mishap analysis (Leveson, 2004). She argues that a model based on system theory can provide a more objective way of looking at accidents. In this view accidents occur when “external disturbances, component failures, or dysfunctional interactions among system components are not adequately handled by the control system [...] they result from inadequate control or enforcement of safety-related constraints on the development, design, and operation of the system”. From this the STAMP methodology was created. System-Theoretic Accident Model and Processes.

In STAMP the basic concepts are *constraints*, *control loops and process models*, and *levels of control*. *Constraints*, not events, are seen as the causation of accidents. That is to say an accident occurs because the system fails to impose adequate constraints on the interactions between the components of a system. These constraints are imposed by *control loops and process models*. Systems are described by control based on feedback mechanisms. Meaning that controllers, be they human or machine, observe the state of the system and affect its variables in such a way that goals and constraints are met. To do this, the controller need to have information about the required states of the variables and their relationships. In other words, the controller needs a *model of the process.* The final element of STAMP pertains to the hierarchy of the system. Each level imposes constraints on the level beneath it. Leveson argues that constraints in the higher structures need to be analyzed as well to obtain a complete view of the causal factors. All this highly implies that accidents occur because of inadequate control. Leveson provides a general classification of system flaws that lead to accidents. It can be seen in Appendix B.

From this very short description of STAMP it is already obvious that the creator has sought to look much further than the proximate causes of an incident. The hierarchy element and the focus on constraints instead of events of a STAMP investigation assure that the top layers of the system, like management and regulatory bodies get scrutinized as well. In fact, in most applications of STAMP, the proximate factors seem to take a backseat (Song, 2002; Leveson Dulac & Marais, 2003 ; Nelson, 2008). This does not mean that proximate factors do not matter as much. It can be argued that proximate factors are needed to uncover the ultimate factors. This is where MASA can benefit a STAMP analysis.

A MASA analysis in itself deals only with the proximate causes that can lead to loss of SA and thus an incident. By the definition of SA that is used in this thesis, MASA cannot be used for the higher levels of hierarchy that the STAMP analysis investigates. It is assumed that an understanding of the constraints on the lower levels of hierarchy leads to a better understanding of the inadequate enforcements of constraints on the higher levels.

Paul Nelson of Lund University performed a STAMP analysis on the Comair 5191 accident as his Master thesis. The result is a quite comprehensive report detailing several levels of hierarchy (Nelson, 2008):

- The flight crew
- Comair Airlines
- Blue Grass Airport Authority
- Airport Safety & Standards District Office
- Air Line Pilots Association
- Airport Liaison Representative
- LEX Air Trafic Control Facility
- FAA Air Trafic Organization: Terminal Services
- National Flight Data Center
- Jeppesen-Charting Division
- Federal Aviation Administration

The MASA analysis in this writing has not touched upon the higher level of hierarchies, it only dealt with those proximally closest to the accident. As such it will only be compared to the lowest level of hierarchy in the STAMP analysis and the repercussions of the MASA analysis for the higher hierarchies will be examined.

When looking at the crew, Nelson identifies the following inadequate control actions:

- The crew taxied to runway 26 instead of runway 22.
- The crew did not use the airport signage to confirm their position.
- The crew did not confirm whether the runway heading and compass heading matched.

He attributes these inadequate actions to flaws in the mental model. The crew:

- thought the route to the runway was the same as previously experienced.
- believed the airport chart accurately depicted the taxi route.
- believed high-threat taxi procedures were unnecessary.
- believed they were on runway 22 when the takeoff was initiated.
- believed “lights were out all over the place”

These findings generally agree with the findings of the MASA analysis, and once again it can be seen that the SA construct is as defined in this thesis is present, yet not explicitly mentioned.

Moving on to the other proximate agent, ATC, Nelson's findings corroborate the MASA analysis once again when it comes to the SA. He mentions that ATC:

- did not monitor and confirm that 5191 had taxied to runway 22
- issued takeoff clearance while 5191 was holding short of the wrong runway

It can be seen that MASA analysis provides a deeper level of understanding of how the errors occurred and progressed. A STAMP analysis could greatly benefit from a formal description of SA differences. The SA construct is highly pervasive in the lower levels of hierarchy, albeit not explicitly mentioned, and a full STAMP analysis could greatly benefit from the insights that MASA provides. Two mutually exclusive ways are proposed.

SA as a process model

According to STAMP, accidents occur when constraints are not adequately enforced to counter disturbances. The controllers in the process need to monitor the relevant states, compare them with desirable states and act upon it. This seems very much similar to the first two levels of SA that Endsely proposes (Endsley, 1995). Namely perception of elements in the environment and comprehension of their meaning. To be a bit more precise, the reader again is referred to Appendix B which shows the general classification of system flaws as proposed by Leveson. The sub-sub flaw 1.2.2 states "inconsistent, incomplete or incorrect process model" as a reason for inadequate enforcements of constraints. Leveson suggests that this can be due to flaws in creation, or the maintenance of the model. The parallels between the creation process of the model and the initial SA on one hand and the update of the process and the SA updates on the other hand is quite clear now. Within this framework SA is simply the process model that the controllers require to enforce the constraints. Once the third level of SA is added, projection of the future status of the system, the model of the process is completed.

SA as part of the control loop

The constraints that are considered the main element of an accident in STAMP are imposed by control loops. These loops are nothing more but flows of information and control actions that keep the system in the desired state. The third general flaw points at "inadequate or missing feedback" as an accident causation. Feedback can be considered a flow of information relating the current state of the system to the controllers, i.e. agents. Leveson suggests that feedback loops can be inadequate because they are either not provided in the system design, because of communication flaws, time lag or because of inadequate sensor operations. These can all be likened to the update cues of the SA analysis. Some visual cue might have been omitted during airport planning that could have benefitted the controllers / agents. Or some static noise could have prevented ATC from relaying critical information. The bottom line is that both the feedback loops from the STAMP analysis and the update cues from MASA are updates on the current state of the system.

7. Conclusions and recommendations

The aerospace industry has had tremendous growth resulting in congested airways and crowded terminals. Despite all this, aviation remains the safest mode of transportation. For the greatest part this is achieved through the industries willingness to innovate and learn from past mistakes. No other industry has such a great willingness to document every event, no matter how small, in hopes of further improving the safety of air travel. Perhaps it is because aviation has played such a huge role in the last couple of decades. Socially and economically, aviation has become a fundamental basis of human life, at least in the developed west. Because of it, it has become quite mundane and the loss of human lives has become less and less socially acceptable. It is part of everyday life and should therefore be as safe as possible.

At the same time advancements in aviation technology have significantly improved the reliability of aircraft components over the years. It is estimated that the majority of accidents can be attributed to "human error", and as was shown, most of the categories in human error can be rephrased as differences in situation awareness. Therefore it can be assumed that the majority of aircraft incidents and accidents where there was no component failure can be attributed to a difference in SA among the players partaking in the scenario. The relative growth of these systemic failures has rendered traditional analysis methods less applicable. There is a need for a methodology that can retrospectively analyze accidents and incidents in an unbiased systematic fashion. Agent based modeling proves to be an ideal way to model and analyze emergent systemic failure of complex socio-technical systems. Situation awareness is just one aspect of a socio-technical system and this thesis has sought to find out just how a multi agent situation awareness analysis can better the view on past accidents and incidents in aviation.

It is time to revisit the research questions from chapter one:

Q1: *How can a formal agent-based approach be used for systematic retrospective modeling and analysis of incidents and accidents?* And how can these findings be generalized?

To answer this question it was decided to use the formal description of multi agent situation awareness as proposed by Stroeve, Blom and Van der Park (2003). Retrospectively modeling how the agents might have perceived the system they are part of, means some assumptions would have to be made. To minimize the assumptions it was sought to generalize the moments in a narrative where updates to situation awareness might have occurred. This would greatly ease the search for these moments in the actual cases. In total, six types of update cues were hypothesized:

- Exchange of information
- Visual cues
- Protocol enforced
- Realization of lack of SA
- The attaining of the time of an intent vector
- Degradation

The multi agent situation awareness method was used in combination with these cues to analyze the cases.

The second question was:

Q2: *What does the agent-based approach add to the existing investigation reports and how does it compare to other existing approaches for retrospective modeling and analysis of safety occurrences?*

This question was answered by comparing the findings of the MASA analysis with the findings of official accident reports of three distinct cases.

The first case was the Tenerife runway incursion of 1977. This case was complex due to the shear amount of agents involved. A total of 4 SA differences was found and their interrelations were mapped. The most salient conclusion is that the MASA analysis shows some oversights in the official report. By solely focusing on the fact that the KLM captain took off without permission, the investigators ignored other potential ways one of the fundamental causes could have been prevented. Furthermore, the relation between the identified fundamental causes was not explored, whereas MASA analysis shows the interplay of the SA differences. The final oversight in the official report was SA difference 5: the fact that ATC was not aware that the KLM crew had initiated take-off. This is without a doubt a contributing cause as the controller could have prevented the accident had this difference in SA been resolved.

The second case was an incident in Tenerife involving a coordinated TCAS maneuver. This case contained a non human agent that could nonetheless acquire and maintain SA about other agents.

The case was also complex due to the large amount of identified update cues and the smaller investigation report from which less facts could be drawn. Nevertheless, three SA differences were identified. Again MASA analysis shows some oversights in the investigation report. First off the report recommends that Finnair trains their crew for better responses to TCAS warnings. While this is a recommendation that would follow from the MASA analysis as well, MASA also shows that the TCAS warning was not the only opportunity for the crew to resolve the SA difference and intervene in a more timely fashion. The second finding is how a TCAS downlink could have resolved the incident sooner as well.

The last case is the attempted take-off from a wrong runway at Lexington airport. This case shows that an incident involving only one aircraft and few total agents can me analyzed as well. A total of six SA differences was found. The first thing that stands out when comparing the findings is weight that the investigators have put on the non pertinent conversation, while MASA shows this not to be as unequivocal. There were other better ways in which the crew could have updated their positional awareness, and these cues were missed as well, hence the doubt about the non pertinent conversation being a fundamental cause to the accident. Just as in case 1, this case identified a contributing cause that was not represented in the official report. This was identified as SA difference 3: the crew thought the runway lights were not functioning while in fact they were. Because this was not mentioned as a contributing cause in the official report no recommendations were made. MASA shows how taking away this contributing cause could have potentially prevented the accident.

The Comair case was also analyzed by two other retrospective analysis methods, FRAM and STAMP. Overall it can be stated that a MASA analysis shows at a more precise level how errors occurred and progressed through the system. It is easier to see at a glance which mechanisms allowed the system to fail and it shows ways to improve it by pointing out possible missed update cues. Therefore a MASA analysis would be a great precursor to a FRAM analysis. When comparing the MASA findings to the STAMP results it can be seen that MASA analysis provides a deeper level of understanding of how the errors occurred and progressed. A STAMP analysis could greatly benefit from a formal description of SA differences as the SA construct is highly pervasive in the lower levels of hierarchy, albeit not explicitly mentioned. A full STAMP analysis could greatly benefit from the insights that MASA provides.

The three cases show how complex systems can benefit from MASA analysis, how not only human agents but also technical agents can be captured and it shows how seemingly simple cases can have oversights as well. The MASA analysis also brings to light some facts that are missed by STAMP and

FRAM and it was shown how MASA could benefit these two methods. Overall it can be said that a MASA can add to the official findings in three ways:

1: By reframing the fundamental causes as differences in SA the searching field for preventive measures is widened. Instead of thinking about how to prevent “human error”, a MASA analysis poses the question how to resolve SA differences. All agents in a system can potentially resolve an SA difference. Preventive measures therefore target all agents instead of the single agent that had a “human error”. In other words, MASA can lead to more and different recommendations on preventing similar events in the future.

2: Mapping the interdependencies between the identified SA differences leads to new insights on contributing causes and how one cause played into the other.

3: By formally describing the SA during the event, MASA can help gain new insights and identifying completely new contributing causes.

Nevertheless a whole lot more can be done to improve the method. The following recommendations are made:

- The generalized update cues were sufficient to capture all update moments in these three cases. More types of cues are possible and it is recommended that more cases are studied to find these types.
- On site inspection of the surroundings can greatly reduce assumptions in visual update cues.
- A multidisciplinary team working on the cases would further reduce the amount of assumptions. Experts in psychology and human factors could help determine whether a cue was processed or not.
- A better technique in visualizing the SA difference progression would benefit the readability of a MASA analysis.
- The interdependencies between the SA differences warrants some further research such that causal relations can be attributed more rigorously.
- SA differences can be represented in several ways. These should be performed on the same incident to check if the findings hold up.
- A “what if?” analysis could be performed within a completed multi agent framework to test the assumptions of the missed update cues.

8. References

Blom, H. (2012). Agent-based safety risk analysus lecture series . *Lecture 7 through 9*. Faculty of Aerospace Engineering, TU Delft.

Blom, H. (2013). Agent Based Safety Risk Analysis Lecture Series. *Lecture 9*. Delft: TU Delft.

Bosse, T., Blom, H., Stroeve, S., & Sharpanskykh, A. (2013). An Integrated Multi-Agent Model for Modelling Hazards within Air Traffic Management. *Web Intelligence (WI) and Intelligent Agent Technologies* , (pp. 179-186). Atlanta.

CIAIAC. (2012). *REPORT IN-050/2011.*

Endsley, M. R. (1995). Toward a Theory of Situation Awareness in Dynamic Systems. *Human Factors 37(1)*, 31-64.

Gerstmeir, T., Ladkin, P., & Loer, K. (1997). Analysing the Cali Accident With a WB-Graph. *Human Error and Systems Development Workshop.* Glasgow.

Gilbert, N., & Terna, P. (2000). How to Build and Use Agent-Based Models in Social Science. *Mind & Society Vol. 1*, 57-72.

Harper, K. A., Mulgund, S. S., Zacharias, G. L., & Kuchar, J. K. (2002). *arc.aiaa.org.* Retrieved July 4, 2013, from http://arc.aiaa.org/doi/pdf/10.2514/6.1998-4101

Heath, B., Hill, R., & Ciarallo, F. (2009). A Survey of Agent-Based Modeling Practices (January 1998 to July 2008). *Journal of Artificial Societies and Social Simulation 12 (4) 9*.

HFACS. (n.d.). Retrieved June 20, 2013, from www.hfacs.com: http://hfacs.com/about-hfacs-framework

Holnagel, E., Pruchnicki, S., Woltjer, R., & Etcher, S. (2008). Analysis of Comair flight 5191 with the. *8th International Symposium of the Australian Aviation Psychology Association.* Sydney.

ICAO. (2013). *2013 Safety Report.* Montreal: International Civil Aviation Organization.

Janssen, M. A. (2005). Agent-Based Modelling. In J. Proops, & P. Safonov, *Modeling in Ecological Economics* (pp. 155-172). Cheltenham.

Johnson, C., & Holloway, C. M. (2003). A survey of logic formalisms to support mishap analysis. *Reliability Engineering and System Safety 80*, 271-291.

Kirwan, B. (1998). Human error identification techniques for risk assessment of high risk systems - Part 1: review and evaluation of techniques. *Applied Ergonomics Vol. 29, No. 3*, 157-177.

Ladkin, P. (2005). The PARDIA Classification. In P. Ladkin, *Causal System Analysis* (pp. 263-278). Springer London Ltd.

Ladkin, P. B. (2000). Causal Reasoning about Aircraft Accidents. *Lecture Notes in Computer Science Volume 1943*, 344-360.

Leveson, N. (2004). A new accident model for engineering safer systems. *Safety Science 42*, 237-270.

Leveson, N., Dulac, N., & Marais, K. (2003, Juni). *Applying STAMP in accident investigation.* Retrieved april 2014, from mit.edu: http://esd.mit.edu/wps/esd-wp-2003-02.pdf

Meyer, J. C., Wieringa, R. J., & Dignum, F. P. (1998). The Role of Deontic Logic in the Specification of Information Systems. In *Logics for Databases and Information Systems* (pp. 71-115). Springer.

Nagel, D., & Wiener, E. L. (1988). Human error In aviation operations. *Human factors in Aviation*, 263-303.

Nagel, K., Beckman, R. J., & Barrett, C. L. (1999). TRANSSIMS for urban planning. *6th International Conference in Urban Planning and Urban Management.* Venice.

National Transportation Safety Board. (2007). *Attempted Takeoff From Wrong Runway Bombardier CL-600-2B19, N431CA Lexington, Kentucky August 27, 2006.* Washington .

National Transportation Safety Board. (2007). *Attempted Takeoff From Wrong Runway Comair Flight 5191.* Washington, D.C: National Transportation Safety Board.

Nelson, P. (2008). *A STAMP Analysis of the LEX Comair 5191 Accident.* Lund University Sweden.

O'Hare, D., Wiggins, M., Batt, R., & Morrison, D. (1994). Cognitive failure analysis for aircraft accident investigation. *Ergonomics 34:11*, 1855-1869.

Primiero, G., Raimondi, F., & Rungta, N. (2014). Model Checking Degrees of Belief in a System of Agents . *AAMAS '14*, (pp. 133-140).
Pritchett, A. R., Lee, S. M., Abkin, M. H., Gilgur, R. C., Corker, K. M., Verma, S., et al. (2002). Examining air transportation safety issues through agent-based simulation incorporating human performance models. *Digital Avionics Systems Conference, 2002. Proceedings. The 21st (Volume:2) .*
Qureshi, Z. H. (2007). A Review of Accident Modelling Approaches for Complex Socio-Technical Systems. *Proc. 12th Australian Conference on Safety-Related Programmable Systems.* Adelaide.
Rasmussen, J. (1982). Human errors: a taxonomy for describing human malfunction in industrial installations. *Journal of Occupational Accidents 4*, 311-333.
Reason, J. (1990). *Human Error.* New York: Cambridge University Press.
Reynolds, C. W. (1987). Flocks, Herds and Shools: A Distributed Behavioral Model. *Computer Graphics, 21:4*.
Salmon, P., Stanton, N. A., Walker, G. H., Baber, C., Jenkins, D. P., McMaster, R., et al. (2008). What really is going on? Review of situation awareness models for individuals and teams. *Ergonomics Science, 9:4*, 297-323.
Shappell, S. A., & Wiegmann, D. A. (1997). A Human Error Approach to Accident Investigation: The Taxonomy of unsafe Operations. *The International Journal of Aviation Psychology 7:4*, 269-291.
Shappell, S. A., & Wiegmann, D. A. (2001). *A Human Error Analysis of Commercial Aviation Accidents Using the Human Factors Analysis and Classification System (HFACS).* Washington, DC: Federal Aviation Administration.
Shappell, S., Detweiler, C., Holcomb, K., Hackworth, C., Boquet, A., & Wiegmann, D. A. (2007). Human Error and Commercial Aviation Accidents: An Analysis Using the Human Factors Analysis and Classification System. *Human Factors, Vol.49, No. 2*, 227-242.
Skybrary. (n.d.). Retrieved July 4, 2013, from http://www.skybrary.aero/index.php/B742_/_B741,_Los_Rodeos_Tenerife,_1977_%28RI_AGC_WX%29
Skybrary. (n.d.). Retrieved from http://www.skybrary.aero/index.php/B752/B752,_en_route,_north_of_Tenerife_Spain_2011_%28LOS_HF%29?utm_source=SKYbrary&utm_campaign=adbdbffac3-SKYbrary_Highlight_04_07_2013&utm_medium=email&utm_term=0_e405169b04-adbdbffac3-264064085
Skybrary, C. (n.d.). Retrieved from http://www.skybrary.aero/index.php/CRJ1,_Lexington_KY_USA,_2006_%28HF_RE_FIRE%29
Smith, K., & Hancock, P. A. (1995). Situation Awareness is Adaptive, Externally-Directed Consciousness. *Human Factors 37(1)*, 137-148.
Song, Y. (2012). *Applying System-Theoretic Accident Model and Processes (STAMP) to Hazard Analysis.* Retrieved from Open Access Dissertations and Theses.
Stanton, N. A., Chambers, P. R., & Piggott, J. (2001). Situational awareness and safety. *Safety Science 39*, 189-204.
Stanton, N. A., Salmon, P. M., Walker, G. H., & Jenkins, D. P. (2010). Is situation awareness all in the mind? *Theoretical Issues in Ergonomics Science 11:1-2*, 29-40.
Stanton, N. A., Stewart, R., Harris, D., Houghton, R. J., Baber, C., McMaster, R., et al. (2006). Distributed situation awareness in dynamic systems: theoretical development and application of an ergonomics methodology. *Ergonomics 49:12-13*, 1288-1311.
Stroeve, S. H., Blom, H. A., & Bakker, G. J. (2009). Systemic accident risk assessment in air traffic by Monte Carlo simulation. *Safety Science, 47:2*, 238-249.
Stroeve, S., Blom, H., & Van der Park, M. (2003). Multi-Agent Situation Awareness Error Evolution in Accident Risk Modelling. *5th USA/Europe Air Traffic Management R&D Seminar.* Budapest.

Stroeve, S., Everdij, M., & Blom, H. (2011). Studying hazards for resilience modelling in ATM. *SESAR Innovation Days.* Toulouse.

Subsecretaria de Avion Civil. (1977). *Colision Aeronaves Boeing 747 PH-BUF de K.L.M y Boeing 747 N 737 de PANAM en Los Rodeos (Tenerife) en el 27 de marzo de 1977.*

Wiegmann, D. A., & Shappell, S. A. (2001). A Human Error Analysis of Commercial Aviation Accidents Using the Human Factors Analysis and Classification System (HFACS). *Aviation, Space, and Environmental Medicine, 72*, 1006-1016.

Xie, Y., John, S., & Donohue, G. (2004). Airport terminal-approach safety and capacity analysis using an agent-based model. *Proceedings of the 2004 Winter Simulation Conference .*

A. Partial FRAM analysis of the Comair accident

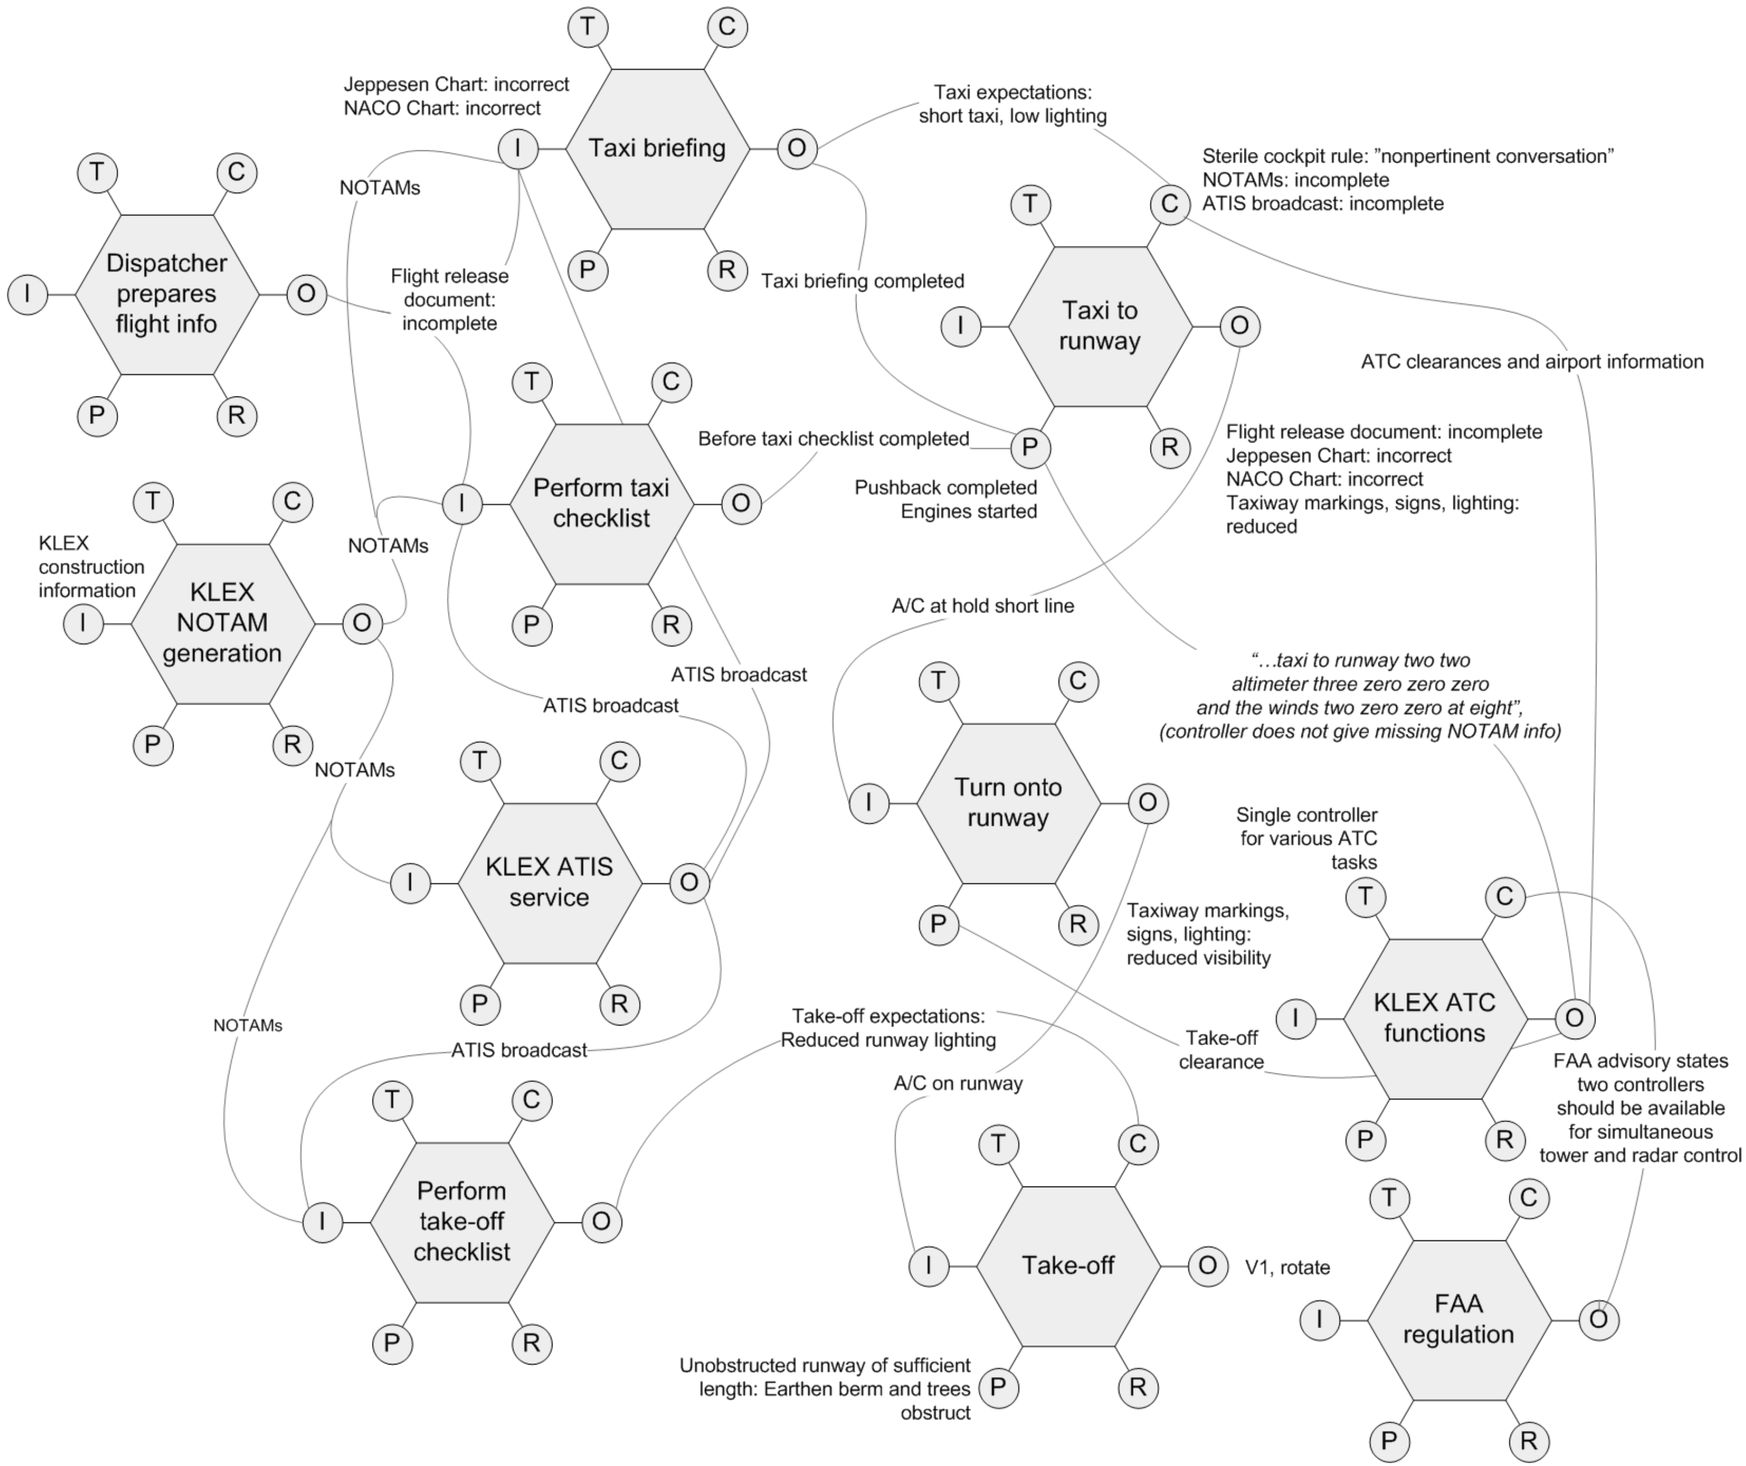

Taxi briefing
Jeppesen Chart: incorrect
NACO Chart: incorrect
NOTAMs
Flight release document: incomplete
Taxi expectations: short taxi, low lighting
Taxi briefing completed
Dispatcher prepares flight info
Taxi to runway
Sterile cockpit rule: "nonpertinent conversation"
NOTAMs: incomplete
ATIS broadcast: incomplete
ATC clearances and airport information
Perform taxi checklist
Before taxi checklist completed
Pushback completed
Engines started
Flight release document: incomplete
Jeppesen Chart: incorrect
NACO Chart: incorrect
Taxiway markings, signs, lighting: reduced
KLEX construction information
KLEX NOTAM generation
NOTAMs
A/C at hold short line
ATIS broadcast
ATIS broadcast
"…taxi to runway two two altimeter three zero zero zero and the winds two zero zero at eight", (controller does not give missing NOTAM info)
Turn onto runway
NOTAMs
KLEX ATIS service
Single controller for various ATC tasks
Taxiway markings, signs, lighting: reduced visibility
KLEX ATC functions
NOTAMs
Take-off expectations: Reduced runway lighting
ATIS broadcast
Take-off clearance
A/C on runway
FAA advisory states two controllers should be available for simultaneous tower and radar control
Perform take-off checklist
Take-off
V1, rotate
FAA regulation
Unobstructed runway of sufficient length: Earthen berm and trees obstruct
T
C
I
O
P
R

B. General classification of system flaws in STAMP

1. Inadequate Enforcement of Constraints (Control Actions)

1.1 Unidentified hazards

1.2 Inappropriate, ineffective, or missing control actions for identified hazards

1.2.1 Design of control algorithm (process) does not enforce constraints

- Flaw(s) in creation process
- Process changes without appropriate change in control algorithm (asynchronous evolution)
- Incorrect modification or adaptation

1.2.2 Process models inconsistent, incomplete, or incorrect (lack of linkup)

- Flaw(s) in creation process
- Flaws(s) in updating process (asynchronous evolution)
- Time lags and measurement inaccuracies not accounted for

1.2.3 Inadequate coordination among controllers and decision makers (boundary and overlap areas)

2. Inadequate Execution of Control Action

2.1 Communication flaw

2.2 Inadequate actuator operation

2.3 Time lag

3. Inadequate or missing feedback

3.1 Not provided in system design

3.2 Communication flaw

3.3 Time lag

3.4 Inadequate sensor operation (incorrect or no information provided)