

EUROCONTROL



Study of Latency of RA Downlink

Edition Number	:	1.5
Edition Date	:	15-05-2007
Status	:	Released Issue
Intended for	:	EATMP Stakeholders



DOCUMENT CHARACTERISTICS

TITLE		
Study of Latency of RA Downlink		
ALDA Reference:		07/05/31-27
Document Identifier	Edition Number:	1.5
QinetiQ/KI/C&IS/ADC/ATC/FARADS/02/11	Edition Date:	15-05-2007
Abstract		
<p>A model of the latencies associated with the downlink of ACAS RA information to ATC centres has been specified and implemented as a software application. The model has been used to investigate the performance of two candidate downlink methods (Mode S RA Report, and 1090 MHz extended squitter) in a set of scenarios representative of those that can be expected in European airspace.</p> <p>The preliminary results presented here show that RA downlink can greatly increase the timeliness of information for improved situational awareness in all contexts, and in safety critical situations will be in time to allow prevention of unwanted controller involvement in RA encounters.</p>		
Keywords		
ACAS Mode S RA Report	TCAS extended squitter	RA downlink Kairos
		latency FARADS
Contact Person(s)		Unit
Stanislaw Drozdowski		+32 (0) 2 729 3760
Harry Hutchinson		+44 (0) 1684 89 5839
		EUROCONTROL DAP/ATS QinetiQ

STATUS, AUDIENCE AND ACCESSIBILITY					
Status		Intended for		Accessible via	
Working Draft	<input type="checkbox"/>	General Public	<input type="checkbox"/>	Intranet	<input type="checkbox"/>
Draft	<input type="checkbox"/>	EATMP Stakeholders	<input checked="" type="checkbox"/>	Extranet	<input type="checkbox"/>
Proposed Issue	<input type="checkbox"/>	Restricted Audience	<input type="checkbox"/>	Internet (www.eurocontrol.int)	<input checked="" type="checkbox"/>
Released Issue	<input checked="" type="checkbox"/>				

ELECTRONIC SOURCE		
Host System	Software	Size
Windows_NT	Microsoft Word 10.0	

EUROCONTROL Agency, Library Documentation and Archives (ALDA)
EUROCONTROL Headquarters (50.703)
96 Rue de la Fusée
B-1130 BRUSSELS

Tel: +32 (0)2 729 11 52
E-mail: publications@eurocontrol.int

DOCUMENT APPROVAL

The following table identifies all management authorities who have successively approved the present issue of this document.

AUTHORITY	NAME AND SIGNATURE	DATE
QinetiQ	H. Hutchinson	2006-04-15
EUROCONTROL	S. Drozdowski	2006-05-03

DOCUMENT CHANGE RECORD

The following table records the complete history of the successive editions of the present document.

EDITION NUMBER	EDITION DATE	REASON FOR CHANGE	PAGES AFFECTED
0.1	2005-11-21	first draft of document	all
1.0	2005-11-28	internal QinetiQ review	various
1.1	2005-12-02	change of name of software to Kairos	various
1.2	2006-03-06	EUROCONTROL reviewers' comments incorporated	various
1.3	2006-03-27	Further EUROCONTROL reviewers' comments incorporated	various
1.4	2006-04-15	Editorial change to figures and glossary	41, 43, 44, & 70.
1.5	2007-05-15	Formatting changes	All

CONTENTS

Executive summary	1
1. Introduction	4
1.1 General.....	4
1.2 Background.....	4
1.2.1 ACAS.....	4
1.2.2 European mandate.....	4
1.2.3 Principle of operation of ACAS.....	4
1.2.4 ACAS studies	5
1.2.5 RA downlink.....	6
1.3 FARADS.....	6
1.3.1 Description.....	6
1.3.2 Technical study of RA downlink methods	7
1.3.3 Current study of RA downlink latency	8
2. Scope of the study	9
2.1 Assumptions.....	9
2.1.1 ACAS equippage.....	9
2.1.2 False alerts	9
2.1.3 Multiple aircraft encounters	9
2.1.4 Controllers	9
2.1.5 Multiple alerts	10
2.1.6 Radar coverage.....	10
2.1.7 Traffic advisories	10
2.1.8 Short term conflict alert	10
2.2 Defining moment of an encounter.....	10
2.2.1 Safety scenario.....	10
2.2.2 Operational scenario	11
2.3 Controlling regimes	11
2.3.1 Altitude dependence.....	11
2.3.2 Terminal regime.....	12
2.3.3 En-route regime.....	12
2.4 Radar and RA downlink coverage	12
2.4.1 Radar coverage.....	12
2.4.2 Extended Squitter coverage.....	13
2.5 Presentation of downlinked RA information.....	14

3. Latency model	15
3.1 Overall approach.....	15
3.1.1 Latencies	15
3.1.2 Timelines	15
3.1.3 Individual delays	15
3.1.4 Sampling the latencies	15
3.2 Timelines in the model	16
3.2.1 RA timeline	16
3.2.2 Downlink timeline.....	16
3.2.3 Current means timeline	17
3.2.4 Controller spontaneous communication timeline	18
3.2.5 Controller downlink response timeline	18
3.3 Distributions	19
3.3.1 Encounter geometry	19
3.3.2 Pilot response.....	20
3.3.3 Form of distributions.....	20
3.3.4 Human factors delays.....	21
3.4 Implementation of the model.....	22
4. Results	23
4.1 General.....	23
4.1.1 Use of the model	23
4.1.2 Downlink scenarios	23
4.1.3 Areas investigated.....	24
4.2 Situational awareness	24
4.2.1 Controller aware without RA downlink	24
4.2.2 Latency of RA downlink.....	25
4.2.3 Latency of controller awareness.....	28
4.3 Uninformed controller involvement	29
4.3.1 Background	29
4.3.2 Terminal regime.....	30
4.3.3 En-route regime.....	31
4.4 Informed controller intervention	33
4.4.1 Introduction.....	33
4.4.2 Terminal regime.....	34
4.4.3 En-route regime.....	34
5. Conclusions.....	35
5.1 RA downlink latency model.....	35
5.2 Summary of results	35
5.2.1 Situational awareness	35

5.2.2	Uninformed controller involvement.....	36
5.2.3	Informed controller intervention.....	36
5.2.4	Extended Squitter vs. Mode S RA Report.....	36
5.2.5	Immediate vs. synchronised display of RA downlink	36
5.3	Further work.....	37
6.	References.....	38
7.	Glossary.....	40
A.	Parameter values used in this study.....	45
A.1	Introduction	45
A.1.1	Background	45
A.1.2	Defining moment of encounter	45
A.2	Delays	45
A.2.1	ACAS tracker noise	45
A.2.2	Idealised warning time.....	46
A.2.3	Actual warning time	47
A.2.4	Clear of conflict.....	48
A.2.5	ACAS display processing time	48
A.2.6	Pilot reaction time.....	48
A.2.7	Pilot delay in contacting controller.....	49
A.2.8	Time until frequency becomes available	49
A.2.9	Duration of message on the RT	49
A.2.10	Controller reaction time to a pilot communication	49
A.2.11	Response to RA causes a significant apparent altitude deviation	50
A.2.12	ACAS clock cycle	51
A.2.13	Radar processing delay.....	51
A.2.14	Update of CWP display	51
A.2.15	Controller detection time of unexpected deviation (clearance not followed).....	52
A.2.16	Controller delay in attempting to contact pilot (unexpected deviation).....	52
A.2.17	RA information uploaded to Mode S transponder	53
A.2.18	RA data processing at ATC centre.....	53
A.2.19	Controller notices downlinked RA	53
A.2.20	Controller comprehends downlinked RA.....	54
A.2.21	Controller spontaneously issues an instruction (unaware of the RA)	54
A.2.22	Controller notices unexpected deviation (RA not followed)	54
A.2.23	Pilot reaction time to controller message	55
A.2.24	Change in altitude	56
A.2.25	Altitude data via Mode S SSR.....	56
A.2.26	RA information via Mode S SSR	57
A.2.27	RA information via Extended Squitter	57

A.2.28	Detection by Mode S SSR.....	58
A.2.29	Processing of altitude data by SSR.....	60
A.2.30	Detection by ES ground station.....	60
A.2.31	Data transmitted through ground network.....	61
A.2.32	Processing of RA information by SSR	61
A.3	Probabilities.....	62
A.3.1	Probability of ACAS equipage	62
A.3.2	Pilot decides to inform controller of RA	62
A.3.3	RA requires a deviation from clearance	62
A.3.4	RT frequency busy when pilot attempts communication.....	62
A.3.5	Pilot complies with RA.....	62
A.3.6	Aircraft altitude reports quantised to 25-ft precision.....	63
A.3.7	RT frequency busy when controller attempts communication	63
A.3.8	Controller (unaware of RA) issues an instruction.....	63
A.3.9	Probability of horizontal manoeuvre	64
B.	Sensitivity Analysis.....	65
B.1	Introduction	65
B.1.1	General.....	65
B.1.2	Key results.....	65
B.2	Approach.....	66
B.2.1	Parameters of interest	66
B.2.2	Latencies of interest	67
B.3	Sensitivity measures	68
B.3.1	Baseline scenario	68
B.3.2	Variation in delay parameters.....	68
B.3.3	Variation in probability	68
B.3.4	Interpretation of sensitivity measures.....	69
B.4	Results	69
B.4.1	Scatter plot	69
B.4.2	Sensitive combinations.....	70
B.4.3	Summary	71

EXECUTIVE SUMMARY

Background

The deployment of the Airborne Collision Avoidance System (ACAS) II has been demonstrated to reduce the risk of mid-air collision, and its carriage is mandatory in European airspace.

ACAS II is an airborne avionics system that works by interrogating and tracking nearby transponder equipped aircraft and issuing Resolution Advisories (RAs) to the flight crew when there is a diagnosed risk of impending collision. An RA provides the pilot with advice on how to regulate or adjust the vertical speed of his aircraft so as to avoid a collision.

In many instances complying with an ACAS RA will cause an aircraft to deviate from its ATC clearance. Currently, a controller can only become aware of this deviation if informed by the pilot or by noticing an unexpected variation in altitude. This awareness can occur many seconds after the RA is issued on board the aircraft, engendering a lack of situational awareness in the controller. Instances of controller instructions to manoeuvre in a sense contrary to an ACAS RA are frequent and can have catastrophic consequences.

Methods exist by which RA information could be downlinked (via the aircraft's Mode S transponder) and provided to controllers in a more timely fashion. All aspects of the feasibility of downlinking RAs are being addressed by the EUROCONTROL FARADS (Feasibility of ACAS RA Downlink Study) project. A crucial aspect in the assessment of any proposed method is the latency of the downlinked information, *i.e.* the time delay between an RA being presented to the pilot and the downlinked RA information being provided to the controller.

Latency of RA downlink

A technical study within the FARADS project identified two candidate technologies that satisfied the requirements for RA downlink (*viz.* Mode S RA Reports and 1090 MHz Extended Squitter) and recommended further work to confirm that they could meet minimum latency requirements.

The current document reports a theoretical study of the latency of RA downlink for the two candidate technologies. The study also goes beyond this to consider other latencies related to the ATC/RA interaction:

- Situational awareness – how quickly can downlinked RA information be presented to the controller and how quickly will the controller become aware of the RA?
- Uninformed controller involvement – can RA information be downlinked quickly enough to prevent the controller (without knowledge of the RA) becoming involved in an encounter?
- Informed controller intervention – can RA information be downlinked quickly enough to enable the controller to change the response by a pilot?

Approach in current study

The approach has been to develop a mathematical model of the latencies and to implement the model as software that runs on a PC. The latencies are decomposed into individual sequential delays whose distributions are defined and sampled stochastically to build up an overall statistical distribution of the latencies. The distributions of the latencies were then analysed to extract statistics indicating the answers to the questions given above.

The computer model was used to investigate the latencies associated with RA downlink for a number of scenarios:

- both of the candidate technologies for achieving RA downlink were considered (Mode S Report and 1090 MHz Extended Squitter);
- both typical coverage and minimum acceptable coverage by Mode S SSR and Extended Squitter ground stations were considered;
- both terminal manoeuvring area and en-route airspace controlling environments were considered; and
- the latencies in the situation in which ACAS RAs are essential to avert a risk of collision and the more general situation in which ACAS RAs are generated even though there might be no immediate risk of collision, were both considered.

Results

The study indicates that the downlink of RA information by either of the technologies considered here is sufficiently timely to allow a significant increase of the situational awareness of controllers in ACAS encounters.

- With the best configuration considered here, the downlink of RA information could allow the controller to be aware of 95% of RAs within 8.1s of their occurrence.
- Currently, in those encounters in which at least one of the aircraft deviates from its clearance due to an RA, the controller will be aware of the RA before any significant deviation occurs in less than 40% of cases. The downlink of RA information virtually doubles the chances that the controller will be aware of the RA before any significant deviation occurs.

In safety critical situations the downlink of RA information is sufficiently timely to allow a significant reduction in the incidence of controllers inadvertently passing instructions to aircraft that are subject to an ACAS RA.

- With the best configuration considered here, the downlink of RA information might prevent 88% of the incidents in terminal airspace in which a controller passes an instruction to an aircraft, unaware that it is subject to an RA. In en-route airspace the proportion could be as high as 91%.

The timeliness of downlinked RA information to allow the controller to change a pilot response, in safety critical situations, is less impressive.

- The long sequence of required events and the limited timeframe in which they must occur for controller intervention to have an effect, means that in terminal airspace the downlink of RA information would allow the controller to have an effect on pilot

behaviour in only 5% of encounters. In the en-route airspace the proportion is higher but only 12%.

In all the scenarios considered here the performance of Extended Squitter is better than the corresponding performance of Mode S Report. Indeed, in many cases the performance of Extended Squitter with minimum coverage is as good as the performance of Mode S Report with typical coverage. Even so, the performance of both methods is acceptable and one would not necessarily discount the use of Mode S Report method on the basis of this performance alone.

1. INTRODUCTION

1.1 General

This document has been prepared by QinetiQ and presents the results of an investigation (based on a mathematical model) of the latencies associated with the communication of ACAS RA information to an ATC centre on the ground. The document has been developed as part of the EUROCONTROL FARADS project.

1.2 Background

1.2.1 ACAS

ACAS stands for Airborne Collision Avoidance System and denotes a family of airborne avionics systems that use standard SSR transponder technology to provide a last resort safety-net against the risk of mid-air collision. There are two types of ACAS equipment:

- ACAS I issues Traffic Advisories (TAs) – alerts indicating the presence of another aircraft that might constitute a collision threat; and
- ACAS II issues TAs and in addition can issue vertical Resolution Advisories (RAs) against intruders that are diagnosed as posing a risk of imminent collision.

1.2.2 European mandate

The carriage of ACAS II by all civil fixed-wing turbine-engined aircraft having a maximum take-off mass exceeding 5,700kg or a maximum approved passenger seating configuration of more than 19 is mandatory in ECAC airspace.

ACAS II is standardised in ICAO SARPs [1] and currently the only compliant implementation is the Traffic alert and Collision Avoidance System (TCAS II) Version 7¹ defined in RTCA MOPS [2]. It is this equipment that is considered in this report.

1.2.3 Principle of operation of ACAS

ACAS interrogates the Mode C and Mode S transponders of nearby aircraft ('intruders') and from the replies tracks their altitude and range.

'Traffic advisories' (TAs) alert the pilot to the presence of an intruder that may become a threat to his own aircraft. They are accompanied by an aural

¹ Although TCAS II Version 6.04A is permitted in the USA.

annunciation and a change of symbol on a cockpit display of traffic information (CDTI) intended to aid visual acquisition.

'Resolution advisories' (RAs) are issued if a diagnosed risk of collision becomes urgent. An RA provides the pilot with advice on how to regulate or adjust his vertical speed so as to avoid a collision. RAs can be displayed in a number of different ways depending on the specific installation (e.g. red and green arcs on a vertical speed indicator, or pitch cues on the primary flight display) and are accompanied by an aural annunciation reinforcing the advice provided by the RA display.

The vertical sense of RAs are coordinated with other ACAS II equipped aircraft so that two aircraft choose complementary manoeuvres.

The nominal warning times (*i.e.* time before predicted collision) for TAs range from 20s near the ground to 48s at high altitude, and the warning times for RAs range from 15s at 1,000ft AGL to 35s at high altitude. ACAS does not have the capability to diagnose a near collision course directly, so these alerts are based on calculations of the time remaining should the aircraft be on collision courses: this necessarily implies a high proportion of alerts in encounters where there is no risk of collision.

ACAS issues a 'clear of conflict' indication (CoC) when the diagnosed risk of collision has passed. Generally the clear of conflict indication is issued when the range between own aircraft and the intruder starts to increase. However, a miss distance filter (MDF) allows the system to issue an early clear of conflict indication (*i.e.* while the aircraft are still converging) if a reliable estimate of the projected horizontal miss distance (HMD) indicates that the separation will be sufficiently large (in collision avoidance terms).² The HMD threshold ranges from 0.2NM at 1,000ft AGL to 1.1NM at high altitude. ACAS is unable to make a reliable estimate of HMD (effectively disabling the MDF) if either of the aircraft is turning.

1.2.4 ACAS studies

Safety studies, such as those comprising part of the EUROCONTROL ACASA Project [3], have demonstrated the safety benefit that can be expected as a result of the widespread equippage with ACAS.

The EUROCONTROL ASARP Project [4] has assessed the safety benefits of ACAS above FL285 following the introduction of RVSM into European airspace.

The EUROCONTROL IAPA Project [5] examined the potential implications for ACAS performance from the introduction of Airborne Separation Assistance Systems in future operations.

Outputs from all three of these studies have been used in the current study.

² The MDF can also completely suppress RAs when a reliable estimate of the predicted HMD is always in excess of the ACAS miss distance threshold.

1.2.5 RA downlink

Complying with an ACAS RA will in many instances cause an aircraft to deviate from its ATC clearance. Currently, a controller can only become aware of this deviation if informed by the pilot or by noticing an unexpected variation in the Mode C altitude displayed at his Controller Working Position (CWP). This awareness can occur many seconds after the RA is issued on board the aircraft, engendering a lack of situational awareness in the controller. Instances of controller instructions to manoeuvre in a sense contrary to an ACAS RA are frequent and can have catastrophic consequences.

There are several methods by which information about RAs could be downlinked and provided automatically to controllers in a more timely fashion. In all cases ground-based systems can detect detailed RA information transmitted from aircraft through their Mode S transponders and this information can then be processed and displayed to a controller at his CWP.

The feasibility of doing so is being addressed by the EUROCONTROL FARADS (Feasibility of ACAS RA Downlink Study) project. A crucial aspect in the assessment of any proposed method is the latency of the downlinked information, *i.e.* the time delay between an RA being presented to the pilot and the details of the RA being provided to the controller.

1.3 FARADS

1.3.1 Description

The high level European Action Group on ATM Safety (AGAS) [6] recommended a study to determine feasibility of downlinking ACAS RAs for display on controller screens. This led EUROCONTROL to instigate FARADS: the 'Feasibility of ACAS Resolution Advisory Downlink Study'.

The objective of the FARADS is to assess the technical and operational feasibility of displaying ACAS RA information on CWP. Some initial experiments have been conducted with the aim, among other things, of obtaining controller's views on different potential implementations of the RA Downlink concept. These experiments showed that the majority of controllers saw clear operational benefits, including:

- Improved air traffic controller situational awareness by helping them to anticipate aircraft manoeuvres.
- Reduced likelihood of contradictory ATC clearances to the conflict aircraft.
- Reduced risk of follow-up conflicts through better information and planning following the resolution advisory.

Whilst RA Downlink may be technically feasible, it is important that its use is carefully validated prior to implementation. Such validation should include examination of many issues, including:

- Evaluation of different technologies;
- Evaluation of different procedural options;
- Human factors assessment of different display options; and
- Safety impact.

Individual studies within FARADS are addressing all of these issues.

1.3.2 Technical study of RA downlink methods

A study [7] within FARADS identified four primary candidates as methods by which ACAS RA information could be downlinked to ATC centres. These methods are:

- ACAS Coordination messages – the interception by a passive ground network of the ACAS resolution message and its coordination reply between two aircraft coordinating their RAs;
- RA Broadcast – the downlinking of ACAS RA information using a 1030 MHz message to a passive ground network;
- Mode S RA Report – the downlinking of information from the aircraft in reply to an interrogating Mode S SSR radar station; and
- 1090 MHz Extended Squitter – the broadcast of an event driven extended squitter message containing RA information to a passive ground network.

These methods were assessed against various criteria and it was concluded that:

- in areas covered by a Mode S ground infrastructure, Mode S Report is the best method for RA downlink;
- in areas not covered by a Mode S ground infrastructure, Extended Squitter is the best method for RA downlink (assuming it can be economically implemented as part of an ADS-B system);
- deficiencies in ACAS coordination messages mean that they are not recommended; and
- RA Broadcast was not recommended based on provisional latency estimates.

The current study therefore concentrates on the Mode S Report and Extended Squitter methods. RA Broadcast would be a potential area for future study if ADS-B ground infrastructure becomes available.

1.3.3 Current study of RA downlink latency

The current study investigates the latencies associated with various aspects of the downlinking of RAs using either of the methods recommended for further consideration by the technical study [7] (*viz.* Mode S Report and Extended Squitter).

The objective was to investigate the following latencies:

- delay between the generation of the RA and the presentation of the downlinked RA at the CWP;
- presentation of the RA relative to the latest time that the information is of use;
- presentation of the RA relative to the time at which a controller might otherwise have passed an instruction to the pilot; and
- time by which advice from the controller, resulting from the presentation of a downlinked RA, can change the pilot response to an RA.

The study also investigates the same latencies in the current situation, where a controller can only become aware of an RA if informed by the pilot or by noticing an unexpected variation in the Mode C altitude.

The approach has been to develop a mathematical model of the latencies and to implement the model as software that runs on a PC. The latencies are decomposed into individual sequential delays whose distributions are defined and then sampled stochastically to build up an overall statistical distribution of the latencies.

The model contains many adjustable parameters that characterise the distributions of the individual delays. These parameters can be changed to reflect any particular controlling environment and downlinking scenario. Even so, any implementation of RA downlink would obviously require operational trials to verify any conclusions drawn from the model.

The model has been used with a number of representative scenarios to produce the results reported in this document. Whilst not necessarily characteristic of any specific controlling environment the range of scenarios encompass most of those in which RA downlink might be implemented and are therefore indicative of the typical performance that can be expected in European airspace.

2. SCOPE OF THE STUDY

2.1 Assumptions

A number of basic assumptions have been made in this preliminary study and these are described in the following paragraphs.

2.1.1 ACAS equippage

It has been assumed that ACAS equippage will be in accordance with the European ACAS mandate. Aircraft required to equip are assumed to be equipped with TCAS II Version 7 and aircraft not required to equip are assumed to be unequipped.

The possibility of RAs in both aircraft when both are ACAS equipped is taken into account.

2.1.2 False alerts

It has been assumed that all systems will work to specification. In particular, this means that there are no false alerts³ and that no errors are introduced into the RA information during the downlinking process.

2.1.3 Multiple aircraft encounters

It is assumed that only two aircraft are involved in each ACAS alert. The ACAS logic includes algorithms to handle multiple aircraft conflicts but such conflicts are relatively infrequent⁴ and go beyond the scope of the current study.

2.1.4 Controllers

It is assumed that both aircraft involved in an encounter are under the control of a single controller and further that only the relevant controller is provided with downlinked RA information.

³ *N.B.* A 'false alert' implies that some aspect of the system is not performing to specification. ACAS can routinely generate alerts when there is no risk of collision, but (provided the system is performing to specification) these are termed 'nuisance alerts'.

⁴ The ACASA project estimated that ACAS multiple encounters occur at the rate of one every 1.7×10^5 flight-hours in European airspace [3], while RAs occur at a rate of about one every 300 flight-hours (see e.g. [21]). On this basis it is estimated that less than 0.2% of RAs occur in multiple aircraft encounters.

2.1.5 Multiple alerts

It is assumed that there will be only one encounter generating RAs at any one time. RAs are not so infrequent that occasionally an RA will occur amongst those aircraft under the control of a single controller when an RA in a separate conflict is still active. However, such occurrences will be in the minority and the assumption is valid for a preliminary study.

2.1.6 Radar coverage

It is assumed that the controller's display of traffic is based solely on Mode S SSR coverage. Furthermore it is assumed that in the case of RA downlink via Mode S Report the same network of Mode S radars is used.

2.1.7 Traffic advisories

The effect of TAs is not considered. TA information is not available from the aircraft transponder and cannot therefore be downlinked. Also operating procedures do not require pilots to report TAs. However, the occurrence of a TA could prompt a pilot to contact the controller and in this way have an effect on the matters considered here. The effect is expected to be of minor importance and goes beyond the scope of the current study.

2.1.8 Short term conflict alert

It is assumed that there will be no STCA alerts during the duration of any RA. This will obviously be the case in sectors where STCA is not deployed, and even where STCA is used there will not necessarily be an interaction since STCA and ACAS employ differing nominal warning times and diagnose conflicts in different ways. However, there can be an interaction between the two systems with alerts being generated at similar times. Nevertheless the inclusion of STCA goes beyond the scope of this preliminary study not least because STCA is not a standardised system and different designs and/or different thresholds are employed in different areas.

2.2 Defining moment of an encounter

Downlinked RA information is of use for a limited time (from the controllers' perspective). The time after which the information is of no further use is referred to as the 'defining moment' of an encounter.

The defining moment is determined differently in two different encounter scenarios discussed below.

2.2.1 Safety scenario

ACAS is intended as a last resort safety net in encounters where there is a genuine risk of collision, *i.e.* those encounters with a negligible HMD in which manoeuvres in the vertical plane (as a response to an RA and/or a controller

instruction) will determine whether a collision is averted or not. This scenario is referred to here as the 'safety scenario'.

In these circumstances the defining moment is the closest point of approach (CPA). After this instant the range between the aircraft increases and the immediate risk of collision has passed (although there will generally still be a serious loss of separation).

In such encounters the CoC generally occurs a few seconds after CPA.

2.2.2 Operational scenario

For the vast majority of encounters in which RAs are generated there is no risk of collision – indeed, RAs can even be generated when there is no loss of separation.⁵ This scenario, in which encounters have a significant (from a collision avoidance perspective) HMD, is referred to here as the 'operational scenario'.

In these circumstances the defining moment is the end of the RA when the CoC indication is issued to the pilot (or the later of the two CoC indications in encounters where both aircraft receive RAs).

The operation of the MDF means that the CoC can be issued significantly before CPA in some encounters.

2.3 Controlling regimes

2.3.1 Altitude dependence

Differing performance of RA downlink can be expected under different controlling regimes. These differences are captured within the model by considering operations as a function of altitude.

The altitude of operations has three principal effects:

- ACAS parameters – the thresholds used by the ACAS logic vary with altitude and affect the timing and nature of RAs;
- encounter geometry – the trajectories of the two aircraft and their juxtaposition will affect the timing and nature of any RAs that are generated; and
- radar coverage – coverage differs at different altitudes and this will affect the downlinking of RAs via Mode S Report and the detection, by the controller, of altitude deviations.

⁵ In this respect the majority of RAs that occur during routine operations are nuisance RAs.

The model takes account of these effects by considering operations within two broad altitude bands corresponding to terminal and en-route operations.⁶

2.3.2 Terminal regime

Terminal operations are limited in geographical extent being centred on a particular airport or group of airports. Aircraft typically fly at comparatively slow speeds, frequently executing turns and changes in speed. A large proportion of aircraft are climbing and descending and level-offs at intermediate levels are common.

Terminal operations are modelled by considering encounters in the altitude range of 1,000ft AGL to FL135.

2.3.3 En-route regime

En-route operations cover a larger geographical extent. Aircraft typically fly at comparatively high speeds, and the majority of flight profiles are 'straight and level'.

En-route operations are modelled by considering encounters above FL135.

2.4 Radar and RA downlink coverage

2.4.1 Radar coverage

The radar coverage used to provide the data on the controller's display determines the ability of the controller to become aware of an RA by an unexpected change in the Mode C altitude read-out. In this respect the radar coverage is relevant to the study whatever method of RA downlink is employed.

In addition the radar coverage determines the characteristics of RA downlink via Mode S Report since it is through the Mode S radars that the RA information is obtained from the aircraft.

Differing radar coverage can be expected in terminal and en-route controlling regimes due to the differing applicable geographical extent and altitude ranges.

Typical and minimum coverage scenarios have been considered in the study for both the terminal and en-route regimes. The same coverage details as those used in the technical study [7] have been adopted.

In the terminal regime it is assumed that relatively short range radars, with a rotation period of 4s, are used. It is assumed that coverage will typically be

⁶ The specialised operations at low-level associated with approach and departure are beyond the scope of this study.

provided by two such radars. The minimum coverage is taken to be a single such radar.

In the en-route regime it is assumed that long range radars, with a rotation period of 8s, are used. It is assumed that coverage will typically be provided by three such radars. The minimum coverage is taken to be two such radars.

In both the terminal and en-route regimes it is assumed that there is typically a 99.6% probability of successfully extracting data from an aircraft transponder on each rotation of the radar. For minimum coverage this probability is assumed to be only to 90% to allow for a worst case scenario.⁷ In each case the same probability is taken as applying to both the extraction of altitude data and the extraction of RA information.

The details of the radar coverage scenarios are summarised in Table 1.

regime	coverage	number of radars	rotation period	data extraction probability
terminal	typical	2	4s	0.996
	minimum	1	4s	0.9
en-route	typical	3	8s	0.996
	minimum	2	8s	0.9

Table 1: Radar and Mode S Report coverage.

2.4.2 Extended Squitter coverage

Extended Squitter information will be obtained independently of Mode S radars through a separate network of fixed antennas. The coverage is assumed to be independent of the controlling regime.

Typically it is assumed that Extended Squitter coverage will be provided by two ground stations. The minimum coverage is taken to be a single ground station.

It is assumed that Extended Squitters will be generated by the aircraft with a repetition cycle of 1s and that there is a 50% probability that the Extended Squitter will be detected at a given ground station on each cycle, for both typical and minimum coverage scenario.⁸

The details of the Extended Squitter ground station coverage scenarios are summarised in Table 2.

⁷ These values are taken from Appendix E of the technical study [7] as described in section A.2.28. The 90% figure represents the probability of data extraction in a single interrogation; the 99.6% figure represents the probability of data extraction assuming a typical value of up to 4 interrogations per beam dwell.

⁸ These values are taken from Appendix E of the technical study [7] as described in section A.2.30.

regime	coverage	number of ground stations	repetition period	data extraction probability
terminal	typical	2	1s	0.5
	minimum	1	1s	0.5
en-route	typical	2	1s	0.5
	minimum	1	1s	0.5

Table 2: Extended Squitter ground station coverage.

2.5 Presentation of downlinked RA information

The precise details of which information is presented to the controller and in what manner it is presented are the subject of separate studies within FARADS [8].

Here it is assumed that at least sufficient information will be available to the controller for him to be aware of the occurrence of the RA, the vertical sense of the RA, and the identity of the subject aircraft.

Ideally the downlinked RA information will be presented to the controller as soon as it is available at the ATC centre. However, at some centres it may be desirable to present the information on the next routine update of the information on the controller's display if this arrangement is found to be acceptable.

Both possibilities have been considered in the current study. It has been assumed that the refresh rate with which information on the controller's display is updated is 4.5s and that this is not correlated with the rotation of any of the individual radars providing the information.

3. LATENCY MODEL

3.1 Overall approach

3.1.1 Latencies

Latency is simply the time delay between two specific events. The approach here has been to identify the events of interest and then determine the processes that must occur for the specified event to occur after the initiating event (*viz.* an RA occurs on one of the aircraft in an encounter).

3.1.2 Timelines

The processes are decomposed into individual delays whose duration can be readily calculated. The individual delays are grouped together into families within which the delays can be arranged in a chronological sequence along a timeline. Along a timeline each delay occurs only after the previous delay has occurred.

The latency between any two events on the same timeline is simply the sum of the intermediate delays.

The delays along a number of timelines are calculated in parallel but the timing of all events is referenced to a single event within the encounter (which is taken as the instant of closest approach). The latency between events on two parallel timelines is then simply the difference in the timing of those events relative to the reference event (in fact, the same approach can be used for events on the same timeline).

When both aircraft in an encounter generate an RA a separate set of concurrent timelines is calculated for each RA.

3.1.3 Individual delays

None of the individual delays is of fixed duration but rather each of them can take a different value in each encounter. The statistical distribution of each delay is characterised so that values for the duration of the delay can be sampled stochastically from the distribution.

3.1.4 Sampling the latencies

Having established the structure of the timelines each of the individual delays is sampled once to build up one instance of the timing of each of the events of interest.

The latencies of the events of interest are then calculated to give a single instance of each latency. This procedure is repeated many times to build up a statistically valid distribution of the latencies of interest.

3.2 Timelines in the model

The model employs five separate timelines:

- RA timeline – the encounter geometry determines the instant at which an RA is generated and when the CoC occurs;
- Downlink timeline – after an RA is generated on board an aircraft it is processed by ACAS and made available for downlink via the Mode S transponder. The RA is then downlinked by the appropriate method, processed and transmitted through a ground network to the ATC centre, before being presented to the controller;
- Current means timeline – without RA downlink a controller can only become aware of an RA if informed directly by the pilot, or by contacting the pilot after noticing an unexpected variation in the aircraft's altitude;
- Controller spontaneous communication timeline – if unaware of an RA the controller might pass an instruction to the aircraft; and
- Controller downlink response timeline – if aware of an RA the controller might be able to remind a pilot who erroneously follows an earlier controller instruction that the RA takes precedence.

These timelines are described in detail in [9] and are summarised in the following sections.

3.2.1 RA timeline

The RA timeline is illustrated in Figure 1.

The geometry of the encounter determines the instant at which the criteria for the generation of an RA are satisfied, relative to the instant of closest approach.

'Noise' in the values of the variables tracked by the ACAS algorithms and the nominal 1s clock cycle of ACAS then determine the precise instant at which an RA is actually generated.

Finally the geometry of the encounter again determines the instant at which the 'clear of conflict' indication is issued (which can be before CPA).

3.2.2 Downlink timeline

The RA downlink timeline is illustrated in Figure 2.

After an RA has been generated the ACAS system makes the details of the RA available for downlink via the Mode S transponder.

Whichever downlink method is used, the presence of RA information must be detected and extracted by a ground station. The RA information will then be processed at the ground station before being transmitted onwards to the ATC centre. When there is more than one ground station this process takes place in parallel at each ground station.

When RA information first arrives at the ATC centre it must be processed before being presented at the CWP. This presentation may occur immediately or in some cases the display might be synchronised to the normal refresh cycle of the information on the controller's display.

When the RA information is presented at the CWP there will be a delay before the controller notices the alert. Having noticed that there is information available there will be a further delay while the controller focuses attention on the downlinked RA and processes the data that is being presented.

3.2.3 Current means timeline

The current means timeline is illustrated Figure 3.

After the RA is generated it must be processed and presented to the pilot, and there will then be a delay before the pilot responds to the RA.

At this point there are two means by which the controller may become aware of the RA:

- directly as a result of a communication from the pilot; or
- indirectly after noticing an unexpected variation in the aircraft's altitude and subsequently communicating with the pilot.

Having become aware of an RA the pilot may after some delay decide to inform the controller. There may be a delay until the RT frequency becomes available and then the pilot can communicate with the controller. At the end of the message there will then be a delay until the controller comprehends the message that has been passed.

Alternatively the controller may become aware of the RA indirectly. If the RA requires a deviation from the cleared altitude then the aircraft will start to change altitude after a delay determined by the precise pilot response. When a sufficiently large altitude change has occurred the corresponding altitude reports can be detected by the radar coverage and transmitted to the ATC centre. After processing at the ATC centre the data displayed at the CWP.

After some delay the controller may notice the unexpected variation in the aircraft's altitude and when he comprehends this he will attempt to communicate with the pilot. There may be a delay until the RT frequency becomes available and then the controller can communicate with the pilot.

After the message has been passed to the pilot there will be a delay until he responds and informs the controller of the RA. Again the RT frequency may be

busy and when the message is finally passed to the controller there will be a delay until the controller comprehends and is aware of the RA.

When both aircraft receive an RA there is a comparatively high probability that each pilot will try to communicate with the controller at the same time as the other pilot. This can lead to garbling of the messages that are passed and will incur further delay while the controller requests that the message is passed again. The possibility of garbling is taken into account by considering interaction between the parallel timelines for the two aircraft.

3.2.4 Controller spontaneous communication timeline

The controller spontaneous communication timeline is illustrated in Figure 4.

In the normal course of events in the absence of any RA (or equivalently when the controller is not aware of any RA) the controller might pass an instruction to one or other of the pilots. Such a 'spontaneous' communication (*i.e.* a communication not prompted by the RA) will occur some time after the RA is generated.

If the pilot were to react to such a message there would be a delay until he does so. Having initiated a response to the instruction there would then be a further delay until a significant change in altitude was achieved. It is when any significant change in altitude is achieved that the spontaneous communication, or its suppression by downlinked RA information, changes the outcome of the encounter.

3.2.5 Controller downlink response timeline

The controller downlink response timeline is illustrated in Figure 5.

In contrast to the spontaneous communication considered above, a controller may communicate with a pilot when he (the controller) is aware of the RA. This can have an effect on the outcome of the encounter when the pilot is erroneously following a previous controller instruction that was contrary to the RA. If the controller notices that the pilot response to the RA is incorrect he may be able to remind the pilot that the RA takes precedence.

For this to occur the controller must be already aware of the RA and notice the unexpected change in altitude. When this occurs there will be a delay until the controller identifies the inconsistency between the aircraft behaviour and the RA and attempts to communicate with the pilot. There may be a delay until the frequency becomes available and then the controller can communicate with the pilot.

After the message has been passed to the pilot there will be a delay until the pilot has determined what he must do to comply with the message and initiates a change in the aircraft's vertical rate. Once the manoeuvre starts there will then be a delay whilst the aircraft accelerates before a significant

change in altitude is achieved (so that the controller's communication can have an effect on the outcome of the encounter).

3.3 Distributions

The precise forms of the distributions of the individual delays and the parameters used to characterise them are described in [10] which is reproduced in Appendix A.

Some aspects of the distributions are discussed in this section.

3.3.1 Encounter geometry

The geometry of an individual encounter determines the time at which an RA is generated relative to CPA.

The encounter geometry can be characterised by a number of factors including the approach angle (the difference in track heading of the two aircraft), horizontal miss distance, the speed of each aircraft and whether each aircraft executes a horizontal manoeuvre.

Knowing these factors, and with an understanding of the threat detection algorithms of the ACAS logic, the model makes an estimate of the time at which an RA will be generated. The estimate is not based on a full simulation of the encounter and the performance of ACAS and will not necessarily reproduce the exact timing of an RA in any individual encounter. However, when estimates over a large number of encounters are made (as in this study) a reliable estimate of the distribution of RA timing is built up.

The distributions of encounter parameters differ between the safety scenario and the operational scenario (most obviously for the HMD which is, by definition, negligible in the safety scenario). This is reflected in the model which uses separate distributions for the two scenarios. These are derived from encounter models developed in other EUROCONTROL projects (see section 1.2.4).

The distributions in the safety scenario are taken from the safety encounter models developed in the ACASA project [18] and the ASARP project [16]. The distributions in the operational scenario are taken from the ATM encounter model developed in the IAPA project [19]. In each case the distributions in the encounter models have been derived from an analysis of real encounters observed in various radar data collection exercises from the core ECAC area over the period 1998 to 2005.

These encounter models include a dependence on altitude which has been used to produce encounters in either the terminal or en-route regimes.

3.3.2 Pilot response

The manner in which a pilot responds to a corrective RA will determine the delay until a significant change in altitude is achieved. The response is characterised by the delay until the pilot responds, the vertical rate he aims to achieve and the acceleration with which he achieves this.

There are standard values for these parameters which are implicitly assumed by the ACAS algorithms. However, in practice there is considerable variation in the way in which individual pilots respond (to the extent that sometimes a pilot even ignores an RA).

These variations are captured in a realistic model of pilot response that has been developed in the ASARP project [15]. This pilot response model is based on observed data collected from onboard recordings.

Both the standard and the realistic pilot response model are incorporated as options in the model, but only the realistic pilot response has been used in the study results presented in Section 4.

3.3.3 Form of distributions

Many of the delays in the model are modelled as uniform distributions in which the delay is selected randomly from a fixed range with all values in this range being equally likely.

In some cases this is the correct form of the distribution but in other cases, notably those delays associated with human factors, we would not expect the delay to be of this form. For example, the distributions of reaction time tend to exhibit a peak at a short time interval followed by a long tail.

When seeking to reproduce the form of such individual delays distributions such as the exponential-Gaussian are often used. Such forms require at least three parameters to specify them (e.g. mean, variance, and a shape factor) and in some cases (see section 3.3.4) even the first two of these parameters are not known with as much accuracy as we would like.

Using more sophisticated forms would complicate the model without necessarily adding benefit to a preliminary study such as this. The model combines many individual delays and as more of these are combined it is the mean and variance of the individual delays, rather than their precise form, that determine the overall delays.⁹

Nevertheless if the form of distributions and the parameters characterising them can be determined with sufficient accuracy there is no reason why a model should not use them. This is one area in which the utility of the model could be readily enhanced if further work was undertaken to better characterise the appropriate distributions.

⁹ When very many independent delays are combined the Central Limit Theorem tells us that the overall delay will tend towards a Gaussian distribution regardless of the form of the individual delays.

3.3.4 Human factors delays

A number of individual delays relate to the information processing and reaction times of the pilots and controller to various inputs.

In these cases the pilot or controller realises that information is available and initiates some action based on that information. For simplicity this is modelled as two distinct sequential processes:

- first the pilot or controller must realise that information of some sort is available and direct his attention toward it ('notices' the information);
- then he must determine what the information is, and decide what action he will perform on the basis of that information ('comprehends' the information).

The actual process will be considerably more complex than this, with tasks, to some degree, performed in parallel. However, the requirement here is not to precisely model the individual processes but rather to model the overall delay within the context of the timeline.

An extensive literature search was conducted in an attempt to find data, from studies relevant to ATM tasks, that would enable the determination of precise values for these delays. However, it was discovered that such data are not readily available, even when the scope was widened to include studies in related fields.

As a fall-back position, general data relating to human reactions and the informed opinions of private pilots and former controllers were used to arrive at suitable values. These values were reviewed by a cognitive psychologist with human factors expertise in the ATM domain. These values remain, to some extent, subjective.

This observation does not invalidate the study as it will be seen later that the results based on latencies that involve these factors are so clear cut that even a comparatively large variation in the parameters would not affect the conclusions.

A sensitivity analysis was conducted to determine which variations in the values of human factors delays had the greatest effect on key latencies calculated by the model, and the results are presented in Appendix B.

The parameters to which latencies were most sensitive were:

- the duration of RT messages;
- the probability that the pilot would inform the controller of an RA;
- the delay in a controller contacting the pilot after noticing an unexpected altitude deviation; and
- the delay in a controller noticing downlinked RA information.

3.4 Implementation of the model

The mathematical model has been implemented as a software application written in Java and known as 'Kairos'.¹⁰ The detailed specification of Kairos is contained in [11].

Kairos uses the default values for the parameters of the individual delays as described in Appendix A. However, the facility exists to change virtually any of these parameters, should the user wish to.

The individual delays are sampled stochastically using reproducible pseudo-random numbers. Five million sampled encounters are used to build up the distribution of the latencies. It is found that times obtained from the model are generally accurate to 0.1s and that proportions are generally accurate to 0.1ppt.

Kairos builds the distributions of 35 separate latencies and statistically analyses them. Histograms of these latencies can be saved in a form that is readily imported into an Excel spreadsheet to perform further analysis 'off-line'.

¹⁰ From the Greek *καιρός* = 'instant of time' or 'critical moment'.

4. RESULTS

4.1 General

4.1.1 Use of the model

The Kairos model has been used to investigate the latencies associated with the downlinking of ACAS RA information to ATC centres.

The model has been run 32 times (one run for each of the downlink scenarios described in section 4.1.2). For each run the model sampled five million separate encounters as described in section 3.4. In each encounter the individual delays described in Appendix A were sampled stochastically. Thus the distributions of the 35 separate latencies in each run were determined.

The results were exported to an Excel spreadsheet and analysed 'off-line' to produce the statistics presented in the rest of this section.

In the tables presented in this results section all values are either percentages (and written as such) or times in seconds (written without any units for the sake of clarity). When a value is followed by another value in parentheses, the main value relates to the case where downlinked RAs are presented at the CWP immediately they are available, whereas the value in parentheses relates to the case where the presentation is synchronised to the routine refresh cycle of the data on the display (assumed here to be 4.5s).

4.1.2 Downlink scenarios

All combinations of the following pairs of conditions have been investigated:

- safety / operational encounter scenario;
- Mode S Report / Extended Squitter downlink method;
- typical / minimum downlink coverage;¹¹
- terminal / en-route controlling regime; and
- immediate / synchronised display of downlinked RAs at the CWP.

The combinations of five pairs lead to a total of 32 separate downlink scenarios that have been investigated.

¹¹ Typical radar coverage is assumed in both the typical and minimum coverage scenarios of Extended Squitter.

4.1.3 Areas investigated

Latencies, in the various downlink scenarios, have been investigated to assess potential benefits of downlinking RAs in three key areas:

- Situational awareness – how quickly can downlinked RA information be presented to the controller and how quickly will the controller become aware of the RA?
- Uninformed controller involvement – can RA information be downlinked quickly enough to prevent the controller (without knowledge of the RA) becoming involved in an encounter?
- Informed controller intervention – can RA information be downlinked quickly enough to enable the controller to change the response by a pilot?

4.2 Situational awareness

Downlinking RA information can increase the controller's situational awareness and enable him to make decisions on an informed basis. Enhanced situational awareness might improve the controller's conduct of not only the RA encounter but also allow him to manage third party aircraft in the vicinity in an optimal manner.

4.2.1 Controller aware without RA downlink

4.2.1.1 *Latency relative to time of RA*

In the current environment (without the downlink of RA information) the controller can only become aware of an RA if informed directly by one of the pilots involved in the encounter, or indirectly after a dialogue with the pilot prompted by an unexpected change in the aircraft's altitude.

Even under these circumstance the controller's awareness of the RA may be incomplete. Depending on the information passed, the controller will not necessarily know the vertical sense of the RA, nor the identity of both aircraft involved in the encounter.

The first column of data in Table 3 shows the proportion of encounters in the current environment in which the controller is never aware of the RA. The controller may be unaware of an RA because neither pilot informs him; because the aircraft are not required to deviate from their clearance to comply with the RA; or because the controller does not notice any deviation. Currently the controller is not aware of approximately 15% of RAs (for the typical scenarios considered here).

For those encounters in which the controller does become aware of the RA Table 3 also shows the mean delay, after the earliest RA¹² in the encounter,

¹² In an encounter between two ACAS equipped aircraft, RAs can be generated at different times.

until the controller will be aware (to some degree) of the RA. In all cases the mean delay is approximately 29s. The mean delay is slightly longer when radar coverage is reduced from the typical coverage network to the minimum coverage network but not by a large amount, indicating that the principal means by which controllers become aware of RAs is by being informed by one or other of the pilots.

	controller never aware	typical coverage	minimum coverage
terminal	16.8%	28.7	28.8
en-route	14.6%	29.0	29.1

Table 3: Proportion of encounters in which controller is never aware of an RA and delay until a controller becomes aware of an RA encounter in the current environment (in seconds).

4.2.1.2 *Latency relative to time of significant deviation*

Another measure of interest is the latency of the controller becoming aware of the RA, relative to the time at which either pilot achieves a significant deviation in response to the RA. This is a measure of the extent to which the controller can be forewarned of any deviations that either aircraft might make.

Considering those encounters in the safety scenario,¹³ in which there is a deviation in response to an RA by at least one of the aircraft, we find that the controller is aware of the RA before either aircraft has made a significant deviation¹⁴ in only 37.2% of encounters in the terminal regime, and in 39.8% of encounters in the en-route.

Naturally, such awareness must be due to a communication by one of the pilots rather than the controller observing the unexpected deviation in altitude.

4.2.2 Latency of RA downlink

4.2.2.1 *Latency relative to time of RA*

Table 4 shows the basic RA downlink latency: the delay between an RA first being generated on one of the aircraft in an encounter and RA information being presented at the CWP.

In the first column of data we see the mean time between an RA and its presentation at the CWP with typical downlink coverage.

With immediate presentation of the information the mean time for RA downlink via Mode S Report is 3.4s in the terminal regime and 4.1s in the en-route regime. Even though coverage in the en-route regime is provided by more radars than in the terminal regime, the latency is slightly greater because

¹³ Results for the operational scenario are very similar.

¹⁴ A significant deviation is taken to be 100ft or greater (see section A.2.24).

these radars have a longer rotation period (*viz.* 8s in en-route and 4s in the terminal regime).

		typical coverage		minimum coverage	
		mean	95%-ile	mean	95%-ile
Mode S Report	terminal	3.4 (5.7)	5.2 (8.4)	4.1 (6.3)	6.0 (9.3)
	en-route	4.1 (6.3)	7.1 (10.0)	4.7 (7.0)	8.3 (11.1)
Extended Squitter	terminal	2.3 (4.6)	3.1 (6.8)	2.5 (4.7)	3.3 (6.9)
	en-route				

Table 4: Latency of RA downlink relative to time of earliest RA (in seconds).

With immediate presentation of the information the mean time for RA downlink via Extended Squitter is 2.3s. The latency is the same in both the terminal and en-route regimes because the same ground station coverage is assumed.

The values in parentheses indicate the corresponding latencies when the presentation of downlinked RA information is synchronised to the update cycle of the controller's display. As one would expect the mean latency is greater by an amount equal to half the update interval ($\frac{1}{2} \times 4.5\text{s} = 2.25\text{s}$).

The second column of data in Table 4 shows the ninety-five percentile – the time by which there is a 95% probability that the RA information will have been presented at the CWP.¹⁵

The third and fourth columns of data Table 4 repeat the data of the first and second columns but this time with the minimum downlink coverage. It is noticeable that the degradation in performance (*i.e.* the increase in the latency) is more pronounced for Mode S Report than for Extended Squitter in both absolute and relative terms.

4.2.2.2 *Latency relative to time of significant deviation*

Table 5 considers those cases in the safety scenario in which there is a deviation in response to an RA by at least one of the aircraft. The table presents the proportion of these encounters in which downlinked RA information is presented at the CWP before either aircraft has made a significant deviation in response to the RA.

With immediate presentation of the information there is a very high probability that the downlinked RA will be presented before either aircraft makes a significant deviation. When RAs are downlinked via Extended Squitter it is virtually certain that the downlinked RA will be presented at the CWP before any significant deviation, as is also the case in the terminal regime for either the Extended Squitter or Mode S Report methods of downlinking RAs. Using

¹⁵ The 95%-ile is effectively a measure of the spread of the result: if the spread increased while the mean remained the same the 95%-ile value would move out to a larger time difference.

Mode S Report in the en-route regime the probability is lower but still better than 99%.

		typical coverage	minimum coverage
Mode S Report	terminal	100.0% (98.8%)	100.0% (97.4%)
	en-route	99.8% (97.0%)	99.2% (94.8%)
Extended Squitter	terminal	100.0% (99.9%)	100.0% (99.9%)
	en-route	100.0% (99.9%)	100.0% (99.9%)

Table 5: Proportion of downlinked RAs presented at the CWP before either aircraft makes a significant deviation.

Naturally, the probabilities are slightly lower when the presentation of downlinked RA information is synchronised to the update cycle of the controller's display. For the Extended Squitter method of downlinking RAs the probability is still consistently high at 99.9%. Only for minimum coverage using Mode S Report to downlink RAs in the en-route regime does the probability fall below 95%.

The probabilities presented in Table 5 indicate an upper bound on the performance of RA downlink. In practice not only must the information be presented at the CWP but it must also be noticed and understood by the controller. Table 6 considers the same scenario as Table 5 but allows for the delay until the controller becomes fully aware of the downlinked RA information. (It also takes account of the minority of cases in which the controller first becomes aware of the RA through current means rather than by downlinked RA information.)

		typical coverage	minimum coverage
Mode S Report	terminal	94.9% (82.0%)	91.4% (78.4%)
	en-route	92.0% (80.8%)	88.6% (78.0%)
Extended Squitter	terminal	98.8% (88.3%)	98.5% (87.5%)
	en-route	98.9% (89.5%)	98.7% (88.5%)

Table 6: Proportion of RAs of which the controller is aware before either aircraft makes a significant deviation.

With immediate presentation of the downlinked information there is still a high probability that the controller will be aware of the RA before any significant deviation in altitude. For the Extended Squitter method of downlinking RAs the probability is greater than 98%. Only for minimum coverage using Mode S Report to downlink RAs in the en-route regime does the probability fall below 90%.

Even when the presentation of downlinked RA information is synchronised to the update cycle of the controller's display the probability of the controller becoming aware of the RA before any significant altitude deviation is still

considerable higher than the current situation (when the controller can only become so aware if contacted in time by the pilot). In all cases the controller is twice as likely or almost twice as likely (for the case of minimum coverage using Mode S Report to downlink RAs in the en-route regime) to become aware of an RA before any significant altitude deviation than is currently the case (see section 4.2.1.2).

4.2.3 Latency of controller awareness

The previous section considered time by which RA information can be presented at the CWP. The controller must notice the downlinked information and comprehend it before it is of use. In this section we consider the latency of the time by which the controller is fully aware of the RA. Included in this process are the conventional means currently available which can occasionally alert the controller before the downlinked information.

4.2.3.1 Terminal regime

Table 7 shows the latency of controller awareness for the terminal regime. The first row of data repeats data for the current situation (without RA downlink) taken from Table 3. In the next two rows we see the corresponding latencies when downlinked RA information via either Mode S Report or Extended Squitter is also available as a means by which the controller can become aware of the RA.

	typical coverage		minimum coverage	
	mean	95%-ile	mean	95%-ile
currently	28.8	–	28.7	–
with Mode S Report	7.5 (9.8)	9.8 (12.9)	8.2 (10.5)	10.7 (13.7)
with Extended Squitter	6.5 (8.7)	8.1 (11.4)	6.6 (8.9)	8.3 (11.6)

Table 7: Latency of controller awareness relative to time of earliest RA (in seconds) – terminal regime.

No ninety-five percentile figure is presented for current means as we saw in Table 3 that the controller is never aware of more than 85% of RAs.

4.2.3.2 En-route regime

Table 8 shows the latency of controller awareness for the en-route regime. The values for Extended Squitter are the same as those given in Table 7 in the terminal regime as Extended Squitter coverage is the same in both cases.

Again the latencies for Mode S Report are slightly longer in the en-route regime due to the longer rotation time of the radars assumed in the en-route Mode S coverage scenario.

	typical coverage		minimum coverage	
	mean	95%-ile	mean	95%-ile
currently	29.0	–	29.1	–
with Mode S Report	8.2 (10.4)	11.5 (14.3)	8.9 (11.1)	12.7 (15.4)
with Extended Squitter	6.5 (8.7)	8.1 (11.4)	6.6 (8.9)	8.3 (11.6)

Table 8: Latency of controller awareness relative to time of earliest RA (in seconds) – en-route regime.

4.3 Uninformed controller involvement

4.3.1 Background

Controller involvement in an encounter, when he is not aware that an RA is in progress, can be catastrophic: although flight crew should be aware that RAs takes precedence over any controller instruction, situations have occurred where the pilot followed a controller instruction that was contrary to an active RA. In such circumstances it will be beneficial if downlinked RA information can be provided to the controller sufficiently early to prevent uninformed controller involvement.¹⁶

The timeliness of downlinked RAs in this scenario can be assessed by considering two separate latencies:

- the time at which the controller becomes aware of an RA relative to the defining moment of the encounter (recall that after the defining moment the RA information is of no immediate use); and
- the time at which the controller becomes aware of an RA relative to the time at which he might otherwise have issued an instruction to one of the aircraft – if, during the RA, the controller becomes aware of the RA before he would otherwise have issued an instruction then the downlinked information has prevented uninformed involvement.

These measures are reported below for both the terminal and en-route regimes.

Preventing uninformed controller involvement is most important in the safety scenario. In the operational scenario there is a significant horizontal separation and in these encounters uninformed controller involvement is not an immediate safety issue. However, results are presented for both the safety and operational scenarios for completeness.

¹⁶ There is a similar timeline for the potential disbenefit where a controller is prompted by RA downlink to inappropriately send instructions during an RA.

4.3.2 Terminal regime

Table 9 shows the mean time at which the controller becomes aware of an RA relative to the defining moment of the encounter. Negative values indicate that the controller becomes aware before the defining moment – the desirable situation. Also shown is the proportion of encounters in which the controller becomes aware before the defining moment.¹⁷

It is noticeable that in the current situation the controller becomes aware of the RA in a timely fashion in less than one quarter of encounters in either of the operational or safety scenarios. On average the controller becomes aware 11s or more after the defining moment.

		typical coverage		minimum coverage	
		mean	proportion	mean	proportion
operational scenario	current	+20.7	14.3%	+20.8	14.3%
	Mode S	-2.1 (+0.2)	40.1% (39.8%)	-1.5 (+0.8)	40.0% (39.7%)
	ES	-3.2 (-1.0)	40.3% (39.9%)	-3.0 (-0.8)	40.2% (39.9%)
safety scenario	current	+11.0	24.4%	+11.1	24.4%
	Mode S	-14.5 (-12.3)	99.9% (99.9%)	-13.9 (-11.6)	99.9% (99.7%)
	ES	-15.6 (-13.4)	99.9% (99.9%)	-15.5 (-13.2)	99.9% (99.9%)

Table 9: Latency of controller awareness relative to the defining moment of the encounter (in seconds) in the terminal regime – negative values indicate that the controller is aware before the defining moment.

In the operational scenario downlinked RA information allows the controller to become aware, on average, a few seconds before the end of the RA given immediate presentation of the RA information. When the presentation of RA information is synchronised with the display update the controller becomes aware, on average, at about the end of the RA. In all cases downlinked RA information allows the controller to become aware before the end of the RA in approximately 40% of RAs.

In the safety scenario all methods perform better. This is because there is no early CoC and the RAs are, on average, persisting for the full nominal warning time until the instant of closest approach. For both RA downlink methods the controller becomes aware, on average, about 15s before closest approach; Extended Squitter performing slightly better than Mode S Report. Both methods allow the controller to become aware before closest approach in virtually all encounters.

¹⁷ This proportion is effectively a measure of the spread of the result: if the spread increased while the mean remained the same the proportion of encounters in which the controller becomes aware before the defining moment would decrease.

In Table 10 we turn our attention to the proportion of encounters in which, given that the controller would otherwise have passed an instruction to one of the aircraft, the downlinked RA was in time to prevent this.

		typical coverage	minimum coverage
operational scenario	Mode S Report	33.8% (29.5%)	32.4% (28.3%)
	Extended Squitter	36.3% (31.5%)	35.9% (31.2%)
safety scenario	Mode S Report	83.2% (72.4%)	80.0% (69.2%)
	Extended Squitter	88.4% (77.6%)	87.7% (76.8%)

Table 10: Proportion of controller involvement that could be prevented by RA downlink – terminal regime.

In the operational scenario we see that both downlink methods are in time to allow the prevention of approximately one third of undesirable controller involvements, with Extended Squitter performing slightly better than Mode S Report.

In the safety scenario we see that the longer warning time before the defining moment could allow RA downlink to prevent between 80% and 89% of undesirable controller involvement. Again we see that Extended Squitter performs better than Mode S Report.

4.3.3 En-route regime

The ACAS RA nominal warning times are greater in the en-route regime than they are in the terminal regime and so there is generally more time for the RA information to be successfully downlinked. This is reflected in the better performance of RA downlink that is observed in this section.

Table 11 shows the mean time at which the controller becomes aware of an RA relative to the defining moment of the encounter. Negative values indicate that the controller becomes aware before the defining moment – the desirable situation. Also shown is the proportion of encounters in which the controller becomes aware before the defining moment.

With the current situation the controller becomes aware of the RA in a timely fashion in no more than 54% of encounters. On average the controller becomes aware 3s or more after the defining moment.

In the operational scenario downlinked RA information allows the controller to become aware, on average, 6s or more before the end of the RA given immediate presentation of the RA information. When the presentation of RA information is synchronised with the display update the controller becomes aware, on average, 4s or more before the end of the RA. In all cases downlinked RA information allows the controller to become aware before the end of the RA more than half of encounters.

		typical coverage		minimum coverage	
		mean	proportion	mean	proportion
operational scenario	current	+15.9	27.6%	+16.0	27.5%
	Mode S	-7.1 (-4.9)	51.7% (51.6%)	-6.5 (-4.2)	51.7% (51.5%)
	ES	-8.5 (-6.6)	51.9% (51.6%)	-8.7 (-6.5)	51.8% (51.6%)
safety scenario	current	+3.3	53.8%	+3.4	53.7%
	Mode S	-22.8 (-20.6)	99.8% (99.8%)	-22.1 (-19.9)	99.8% (99.9%)
	ES	-24.5 (-22.3)	99.8% (99.8%)	-24.4 (-22.1)	99.8% (99.8%)

Table 11: Latency of controller awareness relative to the defining moment of the encounter (in seconds) in the en-route regime – negative values indicate that the controller is aware before the defining moment.

In the safety scenario all methods perform better as was similarly observed in the terminal regime. For both RA downlink methods the controller becomes aware, on average, about 23s before closest approach; Extended Squitter performing slightly better than Mode S Report. Both methods allow the controller to become aware before closest approach in virtually all encounters.

In Table 12 we again turn our attention to the proportion of encounters in which, given that the controller would otherwise have passed an instruction to one of the aircraft, the downlinked RA is in time to prevent this.

		typical coverage	minimum coverage
		operational scenario	Mode S Report
	Extended Squitter	47.8% (47.8%)	47.4% (42.6%)
safety scenario	Mode S Report	86.1% (78.6%)	83.9% (76.4%)
	Extended Squitter	91.9% (84.4%)	91.4% (83.9%)

Table 12: Proportion of controller involvement that could be prevented by RA downlink – en-route regime.

In the operational scenario we see that both downlink methods are in time to allow the prevention of more than 40% of undesirable controller involvements, with Extended Squitter performing better than Mode S Report.

In the safety scenario we see that the longer warning time before the defining moment could allow RA downlink to prevent between 83% and 92% of undesirable controller involvement. Again we see that Extended Squitter performs better than Mode S Report.

4.4 Informed controller intervention

4.4.1 Introduction

We now turn to one of the more complicated of the possible benefits of the downlink of RA information.

Through downlinked information the controller can become aware of the RA or RAs that are in effect for a pair of aircraft. The controller can then form an expectation of the subsequent behaviour of the aircraft in response to the ACAS alert.

In some cases the controller will previously (before he was aware of the RA and possibly before the RA was generated) have given an instruction to the aircraft which it so happens is contrary to the sense of the RA.

In these circumstances the flight crew should give precedence to the RA but the behaviour of pilots is not always ideal and sometimes the (now invalid) controller instruction will erroneously be followed.

If the controller monitors the encounter he may notice the erroneous pilot behaviour and be able to remind the pilot that the RA takes precedence.

For this mechanism to be of benefit the, now corrected, behaviour of the pilot must produce a significant change in the aircraft's altitude before the defining moment of the encounter.

The possible benefit of downlinked RAs can be assessed by considering the time at which a significant change in the aircraft altitude can be achieved, through this mechanism, relative to the defining moment of the encounter, and also by considering the proportion of encounters in which such an effect is achieved in a timely manner.

These measures are reported below for both the terminal and en-route regimes.

The prompting of informed controller intervention is most significant in the safety scenario. In the operational scenario there is a significant horizontal separation and in these encounters the precise response of the pilot is not an immediate safety issue. However, results are presented for both the safety and operational scenarios for completeness.

It is found that there is a negligible difference whether the display of downlinked RA information is immediate or synchronised to the display update (and so only one set of data is presented in this section). This is because the limiting process in the mechanism considered here is the display of altitude data that indicates that the RA is not being followed and this almost always occurs after the downlinked RA information has been presented at the CWP.

4.4.2 Terminal regime

Table 13 shows the average time at which an informed controller intervention could first have a significant effect on the course of an encounter in the terminal regime. In all cases the mean time is 42s or more after the defining moment. In less than 6% of those encounters in which the pilot does not comply with the RA will downlinked RA information allow the controller to intervene and rectify the situation.

		typical coverage		minimum coverage	
		mean	proportion	mean	proportion
operational scenario	Mode S	55.0	5.4%	55.6	5.4%
	ES	55.0	5.4%	55.1	5.4%
safety scenario	Mode S	42.5	5.4%	43.2	5.2%
	ES	42.5	5.4%	42.6	5.4%

Table 13: Latency of significant controller intervention relative to the defining moment (in seconds) – terminal regime.

The performance of Mode S Report and Extended Squitter is virtual the same because, as explained above, the limiting process in this mechanism is the display of altitude data that indicates that the RA is not being followed.

4.4.3 En-route regime

Table 14 shows the average time at which an informed controller intervention could first have a significant effect on the course of an encounter in the en-route regime. In all cases the mean time is 35s or more after the defining moment. In less than 13% of those encounters in which the pilot does not comply with the RA will downlinked RA information allow the controller to intervene and rectify the situation.

		typical coverage		minimum coverage	
		mean	proportion	mean	proportion
operational scenario	Mode S	50.5	8.0%	51.0	7.7%
	ES	50.5	8.0%	50.5	8.0%
safety scenario	Mode S	35.6	12.4%	36.3	11.4%
	ES	35.6	12.3%	36.3	12.4%

Table 14: Latency of significant controller intervention relative to the defining moment (in seconds) – en-route regime.

The prospect of a successful intervention is higher in the en-route regime than in the terminal regime due to the longer ACAS RA warning times at higher altitude. However, the performance in neither regime is particularly striking.

5. CONCLUSIONS

5.1 RA downlink latency model

A powerful and flexible mathematical model of the latencies associated with the downlink of ACAS RA information to ATC centres has been developed and implemented as a software application.

The model has been used to investigate the likely performance of two alternative methods by which the downlink of ACAS RA information could be achieved:

- Mode S RA Report – the downlinking of information from the aircraft in reply to an interrogating Mode S SSR radar station; and
- 1090 MHz Extended Squitter – the broadcast of an event driven extended squitter message containing RA information to a passive ground network.

The performance of RA downlink has been investigated in both terminal and en-route controlling regimes, for scenarios that are representative of those that can be expected in European airspace.

5.2 Summary of results

5.2.1 Situational awareness

The downlink of RA information is sufficiently timely to allow a significant increase of the situational awareness of controllers in ACAS encounters.

Currently a controller may be unaware of up to 15% of ACAS RAs and in the remainder he will typically become aware of the RA 30s after the RA has occurred.

With the best configuration considered here, the downlink of RA information could make the controller aware of 95% of RAs within 8.1s of their occurrence.

Currently, in those encounters in which at least one of the aircraft deviates from its clearance due to an RA, the controller will be aware of the RA before any significant deviation occurs in less than 40% of cases.

The downlink of RA information virtually doubles the chances that the controller will be aware of the RA before any significant deviation occurs.

5.2.2 Uninformed controller involvement

In safety critical situations the downlink of RA information is sufficiently timely to allow significant reduction in the incidence of controllers inadvertently passing instructions to aircraft that are subject to an ACAS RA.

With the best configuration considered here, the downlink of RA information might prevent 88.4% of the terminal airspace incidents in which a controller passes an instruction to an aircraft, unaware that it is subject to an RA.

In en-route airspace the proportion could be as high as 91.9%.

5.2.3 Informed controller intervention

The timeliness of downlinked RA information to allow the controller to change a pilot response, in safety critical situations, is less impressive.

The long sequence of required events and the limited timeframe in which they must occur for controller intervention to have an effect, means that in terminal airspace the downlink of RA information would allow the controller to have an effect on only 5.4% of pilot responses.

In the en-route airspace the proportion is higher but only 12.3%.

5.2.4 Extended Squitter vs. Mode S RA Report

In all the scenarios considered here the performance of Extended Squitter is better than the corresponding performance of Mode S report. Indeed, in many cases the performance of Extended Squitter with minimum coverage is as good as the performance of Mode S Report with typical coverage.

Even so, the performance of both methods appears acceptable and one would not necessarily discount the use of Mode S Report method on the basis of this performance alone.

5.2.5 Immediate vs. synchronised display of RA downlink

The immediate presentation of downlinked RA information naturally performs better than the synchronised presentation (which must wait for the next update of the controller's display). However, the difference in performance is not so great that one would discount the use of synchronised presentation out of hand (except perhaps in the one case mentioned below).

A suggested criterion by which downlink performance could be judged is that 95% of downlinked RA information should be displayed within 10s of the occurrence of an RA.

With immediate presentation at the CWP both the methods considered here meet the 10s criterion in both terminal and en-route regimes.

However, with synchronised presentation at the CWP Mode S Report in the en-route regime only just satisfies the 10s criterion with typical coverage and fails to satisfy the criterion with minimum coverage.

5.3 Further work

The latency results presented here are encouraging and suggest that it is worthwhile to continue with the evaluation of RA downlink.

As a preliminary study, the model developed here employed a number of assumptions and simplifications. These could usefully be relaxed to enhance the utility of the model:

- In a number of places the model uses a uniform distribution of a delay (for want of better information) where a skewed distribution such as the ex-Gaussian would be more appropriate (if the relevant parameters were known). If the appropriate form of these distributions can be determined (either from expert advice or purpose designed experiments) they can be readily incorporated into the model, allowing latencies to be determined with greater confidence.
- A principal simplifying assumption was that short term conflict alert (STCA) was not deployed in the controlling environment. The interaction between downlinked RAs and STCA could be incorporated into the model (although this would require a substantial redesign of certain areas) to enhance the utility of the model.

6. REFERENCES

1. *'Standards and Recommended Practices and Guidance Material for Airborne Collision Avoidance Systems'*, Convention on International Civil Aviation, Annex 10, Volume IV, 3rd Edition, ICAO, July 2002.
2. *'Minimum operational performance standards for Traffic alert and Collision Avoidance System II airborne equipment'*, DO-185A, RTCA Inc., 1997.
3. *'ACASA Work Package 1: Studies on the safety of ACAS II in Europe'*, ACASA/WP1/210D, EUROCONTROL, March 2002.
4. *'Final report on the safety of ACAS II in the European RVSM environment'*, ASARP/WP9/72D, EUROCONTROL, March 2006.
5. *'IAPA Project Final Report – Synthesis and Guidelines'*, IAPA/WP11/114D, EUROCONTROL, October 2005.
6. *'AGAS High Level European Action Group for ATM Safety – One safe sky for Europe – a strategic safety action plan for enhanced ATM safety in a single pan-European sky'*, EUROCONTROL web report: <http://www.eurocontrol.int/activities/safety/AGASSAP.pdf>
7. *'FARADS – Technical Study of RA downlink Methods'*, EUROCONTROL, version 1.2, January 2005.
8. *'FARADS – RADE 1: RA Downlink Experiment 1, Final Report'*, EUROCONTROL, May 2004.
9. *'Sources of delay in latency study of RA downlink'*, QinetiQ/KI/C&IS/ADC/ATC/FARADS/02/05, version 2.3, Dec 2005.
10. *'Parameter values for model of latency of RA downlink'*, QinetiQ/KI/C&IS/ADC/ATC/FARADS/02/07, version 1.4, Dec 2005.
11. *'Specification of RA downlink latency modelling tool ("Kairos")'*, QinetiQ/KI/C&IS/ADC/ATC/FARADS/02/06, version 3.3, Dec 2005.
12. *'An Analysis of En Route Domain Air Traffic Control Voice Transcripts'*, B. Hung, MITRE, April 2005.
13. *'Main lessons learnt from the VOCALISE 2000 study'*, L. Graglia, CENA, April 2002.
14. *'Final Communications Operating Concept and Requirements for the Future Radio System'*, EUROCONTROL/FAA, version 0.2, June 2005.

15. *'Final report on the analysis of airborne recorded data and pilot model specification'*, ASARP WP4/19D, version 1.0, EUROCONTROL, April 2005.
16. *'Final Report on post-RVSM European safety encounter model'*, ASARP WP2/34D, version 1.2, EUROCONTROL, January 2006.
17. *'Event tree probabilities'*, ACASA WP1/213W, version 3.0, EUROCONTROL, November 2001.
18. *'European Encounter Model: specifications and probability tables'*, ACASA/WP1.1/186D, version 1.1, EUROCONTROL, December 2001.
19. *'ATM encounter model specification – description and specification'*, IAPA/WP05/030D, version 1.10, EUROCONTROL, January 2005.
20. *'TCAS RA downlink progress meeting presentation'*, Helios Technology, P505D004, version 1, October 2004.
21. *'Operational implications of TCAS II Version 7 in the UK'*, DERA/LSB(ATC)/ACAS/CR/PAP/10, March 1999.
22. *'Minimum operational performance standards for ATRBS/Mode S airborne equipment'*, DO-181A, RTCA Inc., 1992.
23. *'A literature review on reaction time'*, R. Kosinski, Clemson University, <http://biae.clemson.edu/bpc/bp/Lab/110/reaction.htm>
24. *'Human reaction time'*, <http://www.gecdsb.on.ca/sub/projects/psl/senior/24/reaction.htm>
25. *'ACAS RA downlink'*, S. Drozdowski, IFATCA Conference, Hong Kong, EUROCONTROL, March 2004.
26. *'Commentary on the Aircraft Accident of 1/2 July 2002 (Mid-air Collision of Flights DHX611 (DHL) + BTC2937 (BAL) at Überlingen (Lake Constance) Germany'*, International Advisory Group Air Navigation Services (ANSA & AIRADIO), January 2004.

7. GLOSSARY

ACAS	Airborne Collision Avoidance System – ACAS II provides resolution advisories in the vertical plane advising the pilot how to regulate or adjust his vertical speed so as to avoid a collision.
ACASA	ACAS Analysis project – a EUROCONTROL study within which was developed a Safety Encounter Model.
AGAS	European Action Group on ATM Safety.
AGL	Above Ground Level.
ARTAS	ATM Surveillance Tracker and Server.
ASARP	ACAS Safety Analysis post-RVSM Project – a EUROCONTROL study within which the ACASA Safety Encounter Model has been modified to reflect the introduction of RVSM.
ATC	Air Traffic Control.
ATM	Air Traffic Management.
CDTI	Cockpit Display of Traffic Information.
CPA	Closest Point of Approach – the instant in an encounter at which the slant range between the two aircraft is a minimum.
CoC	Clear of conflict – the indication given by ACAS that an RA has ended.
CWP	Controller Working Position.
ECAC	European Civil Aviation Conference.
ES	Extended Squitter
EUROCONTROL	European Organisation for the Safety of Air Navigation
Extended Squitter	The broadcast of an event driven extended squitter message containing RA information to a passive ground network
FARADS	Feasibility of ACAS RA Downlink Study

HMD	Horizontal Miss Distance – the horizontal separation between two aircraft at CPA.
IAPA	Implications on ACAS Performances due to ASAS implementation project – a EUROCONTROL study within which was developed an ATM Encounter Model.
ICAO	International Civil Aviation Organization.
MDF	Miss Distance Filter – a feature of ACAS that can suppress RAs or issue an early CoC when it can be reliably diagnosed that a significant HMD will exist.
Mode S RA Report	The downlinking of information from the aircraft in reply to an interrogating Mode S SSR radar station.
MOPS	Minimum Operational Performance Standard.
NMAC	Near Mid-air Collision – an encounter in which at some point the horizontal separation between the aircraft is less than 500ft and simultaneously the vertical separation is less than 100ft.
RA	Resolution Advisory – an ACAS alert advising the pilot how to regulate or adjust his vertical speed so as to avoid a collision.
RTCA	An independent USA body including representatives of interested parties from the aviation community. The TCAS II MOPS are prepared under the supervision of RTCA's Special Committee 147.
RVSM	Reduced Vertical Separation Minimum.
SARPs	Standards and Recommended Practices.
SSR	Secondary Surveillance Radar.
STCA	Short Term Conflict Alert – a ground based system alerting controllers to potential conflicts.
TA	Traffic Advisory – an ACAS alert warning the pilot of the presence of another aircraft that might become the subject of an RA.
TCAS	Traffic alert and Collision Avoidance System – a specific implementation of the ACAS concept. TCAS II Version 7 is currently the only available equipment that is fully compliant with the ACAS SARPs.

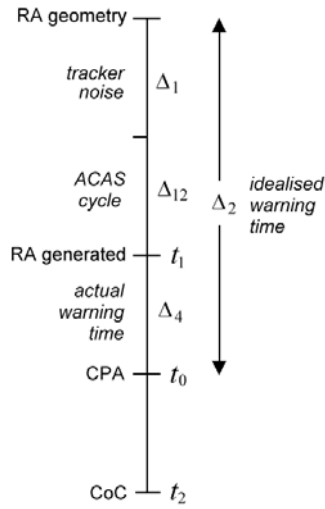


Figure 1: RA timeline

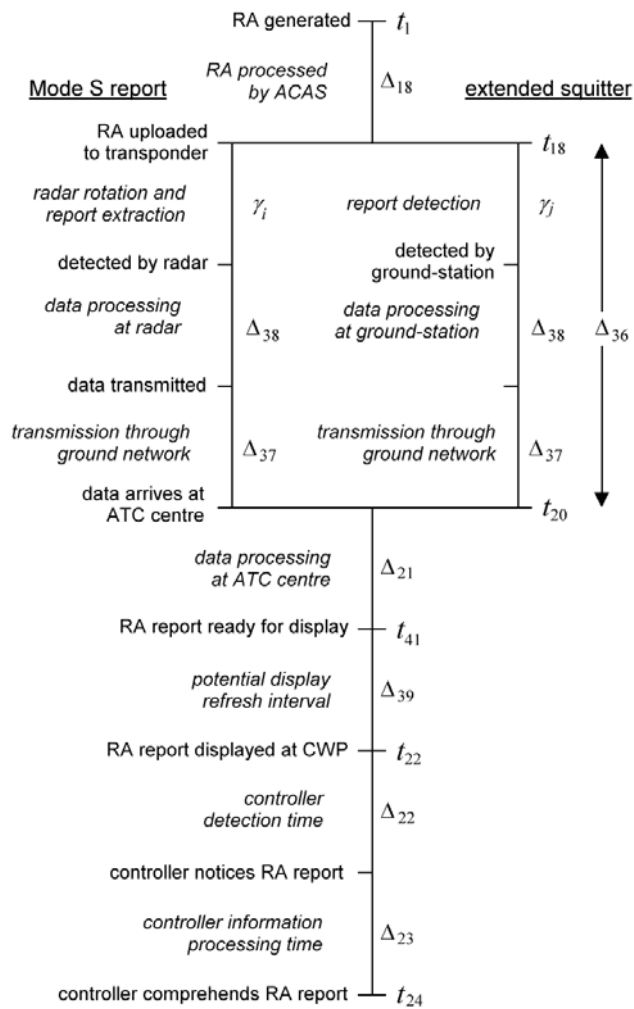


Figure 2: Downlink timeline

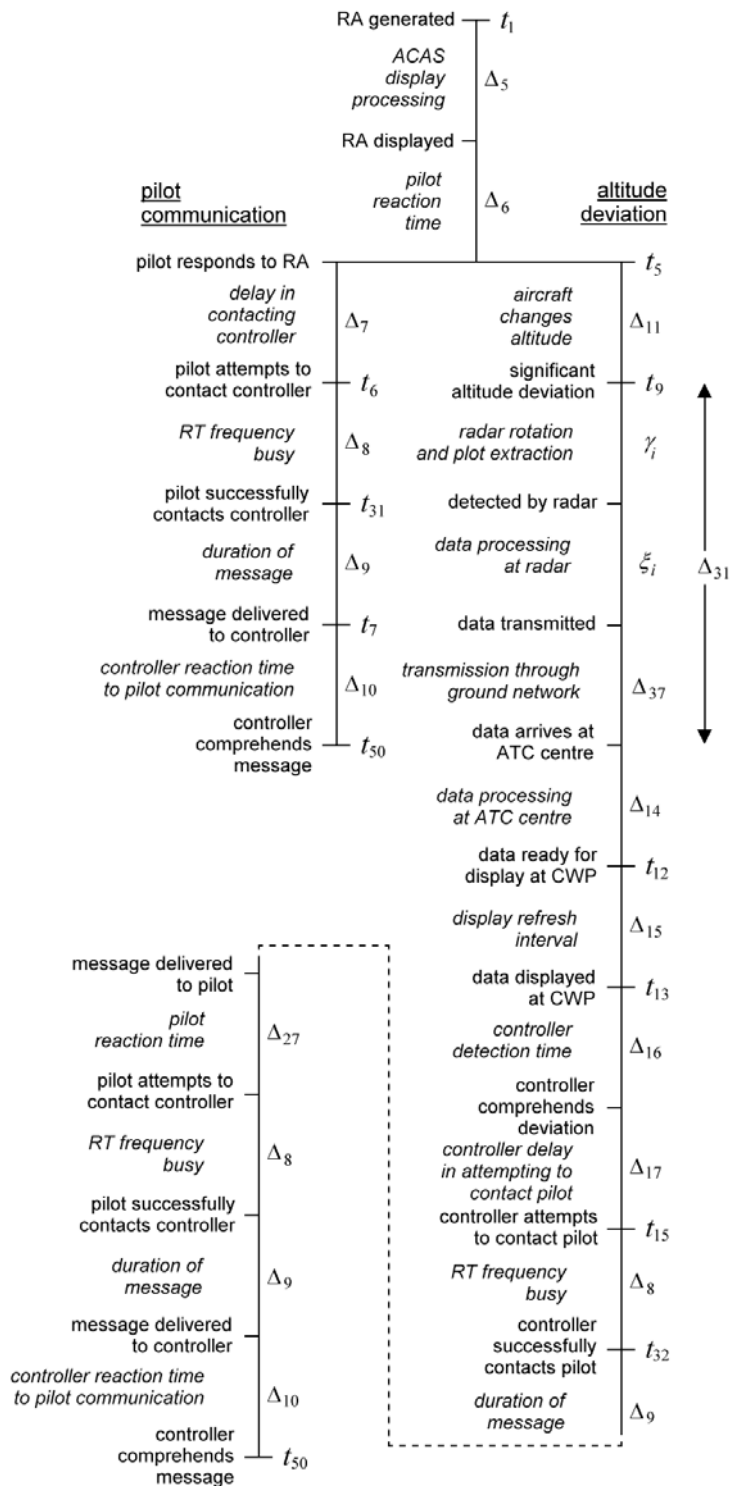


Figure 3: Current means timeline




Figure 4: Controller spontaneous communication timeline

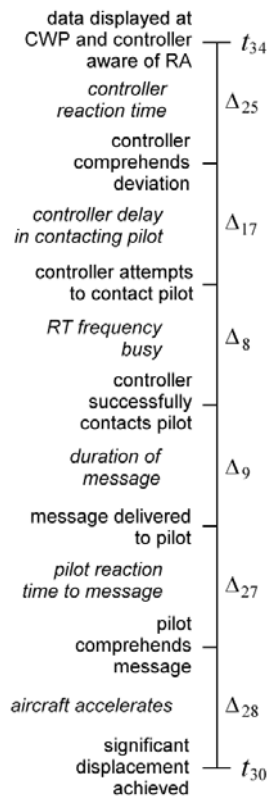


Figure 5: Controller downlink response timeline

A. PARAMETER VALUES USED IN THIS STUDY

A.1 Introduction

A.1.1 Background

A mathematical model of latency in the downlink of RA information is specified in [11]. The model employs many parameters that characterise probabilities and the distributions of individual delays. The default values for these parameters are described in this appendix.

The nomenclature is the same as that employed in Figure 1 to Figure 5.

A.1.2 Defining moment of encounter

Two scenarios are considered; the operational scenario corresponding to the majority of RAs that are routinely generated with a significant horizontal miss distance; and the safety scenario corresponding to RAs with a negligible horizontal miss distance and a consequent risk of collision.

In the operational scenario the defining moment in the encounter, t_{40} , corresponds to the 'clear of conflict' indication (CoC) at t_2 – the end of the RA. In the safety scenario the defining moment corresponds to the closest point of approach (CPA) at t_0 after which the aircraft diverge in range and the immediate risk of collision has passed.

When both aircraft are ACAS equipped and receive an RA the time of the CoC indication in each aircraft is calculated separately. The defining moment for the encounter in the operational scenario is taken as the later of the two CoC times:

$$t_{40} = \max(t_2(1), t_2(2))$$

A.2 Delays

A.2.1 ACAS tracker noise

Delay Δ_1 is due to the tracking algorithms in the ACAS logic. The variables defining the relative positions of aircraft are estimated by the ACAS logic by tracking the measured values. The tracked values will inevitably differ from the true values. Consequently the time at which an alert is declared by the ACAS logic can differ from the time at which the alert would be declared given perfect surveillance.

The ACAS surveillance is conducted on a nominal one-second cycle [1] and we cannot expect the noise due to the tracking algorithms to be smaller than

this. The tracked variables may lead the real world or may lag behind it (*i.e.* an alert may be declared earlier or later than it would with perfect surveillance). Consequently we set the tracker noise to be uniformly distributed from $-1s$ to $+1s$.

$$\Delta_1 = U[-1s, 1s]$$

A.2.2 Idealised warning time

Not strictly a delay, Δ_2 is the time (relative to CPA) when the geometry of the encounter is such that an RA would be declared given perfect surveillance.

We assume that all encounters can be approximated by those in which horizontal flight is linear, and that the RA is equally likely to be generated at any time between an earliest possible time and a latest possible time. The earliest possible time (the upper limit) is the time at which the range test is passed. The latest possible time (the lower limit) is the time at which the test on vertical closure passes. This still requires assumptions concerning statistical distribution of encounter geometries, and for these the ACASA encounter model [18], ASARP encounter model [16] and IAPA encounter model [19] have been used.

Upper limit

For a 'linear' encounter (two aircraft each flying at constant velocity) with a closing speed v , and a miss distance m , a simple Bramson criterion 'time to collision' test, with time threshold T_R , and distance modifier d_m , gives an alert with time τ_0 to collision where

$$\tau_0 = \frac{1}{2}T_R + \sqrt{\left(\frac{1}{2}T_R\right)^2 + \frac{d_m^2 - m^2}{v^2}}$$

This is effectively the range test criterion of ACAS and can be taken as the upper limit of the warning time. (When the expression inside the square root is negative there is no alert.)

The encounter altitude is set according to the proportions of the ASARP encounter model. This in turn determines the appropriate values of the altitude dependent thresholds T_R and d_m from the ACAS logic.

In the safety scenario we set $m = 0$; in the operational scenario m is chosen from the distributions defined in the IAPA encounter model. Altitude is set according to the proportions of the ASARP encounter model. Details of these distribution can be found in the specification [11].

The relative speed v is determined from the individual speeds of the two aircraft, v_1 and v_2 , and the approach angle θ . These variables are chosen from distributions defined in the ASARP encounter model and detailed in the specification [11]. The relative speed is then calculated from the cosine rule:

$$v = \sqrt{v_1^2 + v_2^2 - 2v_1v_2 \cos \theta}$$

Lower limit

Time delays built into the ACAS altitude test can reduce the warning time to values less than that upper limit estimated above. The ACAS threshold T_V is taken as the nominal lower limit of the warning time.

The 'nuisance alarm filter' in the TCAS logic prevents alerts that would otherwise be generated close to CPA when the separation between the aircraft is more than 1.5NM. The effect of the filter is to suppress alerts that would be generated, in encounters in which the miss distance is greater than $m = 1.5\text{NM}/\sqrt{2}$, with less than time τ_{NAF} to go before CPA where τ_{NAF} is calculated as:

$$\tau_{NAF} = \sqrt{\frac{m^2 - d_m^2}{v^2}}$$

The lower limit of the idealised warning time, L , is T_V or the greater of T_V and τ_{NAF} when the nuisance alarm filter is applicable.

Distribution

We assume that for each aircraft the idealised warning time is uniformly distributed between the upper and lower limits. The warning time is sampled separately for each aircraft but using the same limits.

$$\Delta_2 = U[L, \tau_0]$$

A.2.3 Actual warning time

The actual warning time, Δ_4 , is calculated from the idealised warning time, Δ_2 ; the variation due to tracker noise, Δ_1 ; and a delay due to the ACAS logic clock cycle, Δ_{12} . As indicated in [9]:

$$\Delta_4 = \Delta_2 - \Delta_1 - \Delta_{12}$$

The time at which the RA is generated, t_1 , is calculated as

$$t_1 = t_9 - \Delta_4$$

When both aircraft are ACAS equipped and receive RAs the time, t_1 , is calculated as the earlier of the times at which an RA is generated in the two aircraft.

$$t_1 = \min(t_1(1), t_1(2))$$

A.2.4 Clear of conflict

The latest time at which CoC could occur is calculated according to the algorithms in the ACAS SARPS [1] (see [11]).

In the operational scenario the functioning of the miss distance filter can allow an early CoC (*i.e.* before CPA, when the aircraft are still converging in range). This possibility is addressed in Kairos by identifying likely encounter geometries. These are encounters above 2,350ft in which the horizontal miss distance exceeds the DMOD parameter of the ACAS logic [2]. Kairos then determines the probability that neither aircraft in these encounters manoeuvres horizontally before CPA using probability p_{10} .

If there is no manoeuvre then the encounter is declared to be one in which there may be a CoC before CPA. The time at which CoC occurs is assumed to be uniformly distributed from 5s after the start of the RA (a provision from the TCAS MOPS) until the time of the originally calculated CoC.

A.2.5 ACAS display processing time

When the ACAS logic declares an alert there is a delay, Δ_5 , while the system processes the alert and determines what form the display and aural annunciation to the pilot should take.

The ACAS SARPs [1] state that the ACAS system response time from receipt of relevant SSR reply to presentation of an RA to pilot should be as short as possible and not exceed 1.5s. The ACAS clock cycle is nominally one second and so no longer than 0.5s should be taken to process the alert. A lower limit of one tenth of this value is used and the delay is assumed to be uniformly distributed between these limits:

$$\Delta_3 = U[0.05s, 0.5s]$$

A.2.6 Pilot reaction time

The pilot reaction time to an initial RA is delay Δ_6 . The ACAS logic assumes a standard response time by a pilot of 5s [1]; studies of operational data, such as that in ASARP [15], indicate some variation in this value.

Kairos permits the choice of an 'optimal' pilot response, in which case the delay is set at $\Delta_6 = 5\text{s}$; or a 'typical' pilot response, in which case the delay is sampled from a distribution taken from the ASARP study [15].¹⁸

A.2.7 Pilot delay in contacting controller

If a pilot decides to contact the controller to inform him of an RA, the delay after the pilot's reaction time to the RA (Δ_6) is Δ_7 .

Precise values are difficult to obtain, but the experience of former controllers and the opinion of private pilots suggests that the delay can vary from just a couple of seconds to more than ten seconds. We will assume that the delay is uniformly distributed between 2.5s and 12.5s:

$$\Delta_7 = U[2.5\text{s}, 12.5\text{s}]$$

A.2.8 Time until frequency becomes available

If the RT frequency is busy when a pilot or controller attempts to communicate he will have to wait for a period up to the duration of a typical message before the frequency becomes available. This delay is Δ_8 . On each occasion the possible duration of a message, f , is determined by sampling the delay Δ_9 . The delay will then be uniformly distributed between 0s and f .

$$\Delta_8 = U[0\text{s}, f]$$

A.2.9 Duration of message on the RT

Very many types of RT message are transmitted over the frequency and there is much variation in message duration. An examination of a recent study of voice transmissions in en-route airspace [12] suggests that a typical message can last anywhere between 3s and 6s. Therefore a uniform distribution of message duration between these limits is used:

$$\Delta_9 = U[3\text{s}, 6\text{s}]$$

A.2.10 Controller reaction time to a pilot communication

The delay due to the time it takes for a controller to react to a communication from a pilot and be in a position to reply is Δ_{10} . The opinion of a former controller is that this will typically be between 2s and 5s.

This is confirmed by an examination of the routine RT traffic in the transcript of the communications in the fifteen minutes before the Überlingen midair

¹⁸ 'Pilot response' consists of a set of parameters: the delay between the enunciation of the RA and the pilot initiating a vertical acceleration, the magnitude of the vertical acceleration, and the vertical rate that is achieved as a result of that acceleration.

collision [26]. Aircraft contacted the controller on 17 occasions and the delay before a reply ranged from 1s to 6s. The average delay was 3.6s.

A uniform distribution between the limits of 2s and 5s is used:

$$\Delta_{10} = U[2s, 5s]$$

A.2.11 Response to RA causes a significant apparent altitude deviation

When a pilot complies with an RA and the RA requires a deviation from his ATC clearance the deviation will be apparent at the CWP after a delay Δ_{11} .

The deviation that is required before the deviation becomes apparent at the CWP is z .

- When altitude is quantised to the nearest 25ft the first indication that the aircraft is not behaving as the controller expects will come from a discrepancy in the vertical trend indicated on the controller's display. A significant change in the vertical rate can generally be diagnosed once an aircraft has changed altitude by about 125ft: a further 25ft allows for the quantisation of the altitude. Consequently, when altitude is reported with 25-ft precision we assume that z is uniformly distributed from 125ft to 150ft.

$$z = U[125ft, 150ft]$$

- When altitude is quantised to the nearest 100ft the first indication that the aircraft is not behaving as the controller expects will come from a discrepancy in the altitude indicated on the controller's display. A controller will generally tolerate a discrepancy of two flight levels in an aircraft's altitude: a further 100ft allows for the quantisation of the altitude. Consequently, when altitude is reported with 100-ft precision we assume that z is uniformly distributed from 200ft to 300ft.

$$z = U[200ft, 300ft]$$

The pilot will respond with an acceleration, a , and aim to achieve a vertical rate of \dot{z} . We assume that the target change in vertical rate corresponds to that associated with a positive RA received when in level flight, viz. 1,500fpm.

When an optimal pilot response is selected we assume that the standard values of $a = 0.25g$ and $\dot{z} = 1,500\text{fpm}$ apply; when a typical pilot response is selected the values of a and \dot{z} are selected from a distribution derived from the ASARP study [15].

The time required to achieve a change in vertical rate of \dot{z} is \dot{z}/a . During this time the change in altitude is $\dot{z}^2/(2a)$. There are two cases depending upon whether the required change in altitude z is achieved before or after the target vertical rate is achieved:

- if $z \leq \frac{\dot{z}^2}{2a}$ then $\Delta_{11} = \sqrt{\frac{2z}{a}}$
- if $z > \frac{\dot{z}^2}{2a}$ then $\Delta_{11} = \frac{z}{\dot{z}} + \frac{\dot{z}}{2a}$

A.2.12 ACAS clock cycle

The ACAS logic evaluates possible threats and generates alerts by repeatedly proceeding through the collision avoidance algorithms on a fixed cycle. There will therefore be a delay, Δ_{12} , between the conditions for an RA being satisfied and an RA being generated by the logic. This delay is uniformly distributed between zero and the length of the ACAS cycle.

The ACAS cycle is nominally 1s but the SARPs [1] allow a tolerance of $\pm 20\%$ (*i.e.* the cycle can be between 0.8s and 1.2s). Informal conversations with TCAS manufacturers suggest that in practice the cycle is very close to 1s and so this value is used:

$$\Delta_{12} = U[0\text{s}, 1\text{s}]$$

A.2.13 Radar processing delay

SSR radar data arrives at the ATC centre and must be processed before it can be displayed at the CWP. The delay associated with this is Δ_{14} .

Advice from an SSR expert is that this will not be less than 0.25s and can be up to 1s when there are many plots. We will assume that the delay is uniformly distributed between these limits:

$$\Delta_{14} = U[0.25\text{s}, 1\text{s}]$$

A.2.14 Update of CWP display

When radar data is ready for display at the ATC centre there will be a delay, Δ_{15} , until it is displayed when the CWP display is next updated. A single instance of this delay is uniformly distributed between 0s and d .

When this delay is sampled more than once the delays are correlated. It is then necessary to consider the precise time at which the delay is sampled relative to the display update cycle.

The phase of the display update cycle, ζ , is uniformly distributed between 0s and d .

$$\zeta = U[0s, d]$$

Sampled at time t then the delay is:

$$\Delta_{15} = d - ((\zeta + t) \text{MOD } d)$$

For historical reasons (when display updates were directly connected to the radar scan) the display update rate is generally the same as a typical ATC radar scan rate, so we will set $d = 4.5s$.

The ARTAS tracker can be set to 'sympathise' with one of the source radars so that the update rate corresponds to the chosen radar scan rate. Consequently, Kairos will allow the user to select a value other than the default value.

A.2.15 Controller detection time of unexpected deviation (clearance not followed)

An aircraft may deviate from its ATC clearance as a result of an RA. If the controller is not aware of the RA through a pilot communication or downlinked RA information he may still become aware through the observation of an unexpected deviation. The delay, after the deviation is displayed at the CWP, is Δ_{16} .

During a routine scan of the aircraft on the display requiring close attention the controller takes an average time w to assess an individual aircraft. This is taken as the lower limit of Δ_{16} corresponding to the case in which the controller monitors the aircraft in question immediately after the deviation has been displayed on the CWP. It is generally accepted that for moderately complex tasks such as this a value of 0.5s is typical [23]. The default value is $w = 0.5s$.

The upper limit will correspond to the frequency with which the controller monitors those aircraft under his control requiring close attention. A former controller indicated that he would typically expect to look at each of these aircraft at least once every 10s. This value is used as the default value of the upper limit.

The delay is uniformly distributed between the upper and lower limits:

$$\Delta_{16} = U[w, 10]$$

A.2.16 Controller delay in attempting to contact pilot (unexpected deviation)

When the controller observes an aircraft not behaving as expected (either failing to follow an ATC clearance when the controller is not aware of the RA,

or following a previous ATC clearance in preference to the RA when the controller is aware of the RA) it is assumed that he will attempt to contact the pilot. The delay until the attempt is Δ_{17} .

The quickest human reaction times are generally about a quarter of a second [23], but reaction times can increase when the subject is under stress [24]. We assume that the delay is uniformly distributed between 0.25s and 1s:

$$\Delta_{17} = U[0.25s, 1s]$$

A.2.17 RA information uploaded to Mode S transponder

When an RA is generated by the ACAS system it must be uploaded to the Mode S transponder before it is available to be downlinked through either a Mode S report or an extended squitter. The delay associated with the upload is Δ_{18} .

The Mode S transponder MOPS [22] state, that while an RA is active, the content of the MB field in the RA report shall be updated at least once a second: we therefore set the upper limit of the delay at 1s.¹⁹ The lower limit is non-zero and so we use a value that is one tenth of the upper limit:

$$\Delta_{18} = U[0.1s, 1.0s]$$

A.2.18 RA data processing at ATC centre

When downlinked RA data arrives at the ATC centre it must be processed to determine what form the display should take before it can be displayed at the CWP. The delay associated with the processing is Δ_{21} .

Discussions with an SSR expert suggest that the processing delay is likely be similar to the delay in processing other coded data such as Mode C data. The delays here are typically between 0.25s and 1s. We will assume that the delay is uniformly distributed between these limits:

$$\Delta_{21} = U[0.25s, 1s]$$

A.2.19 Controller notices downlinked RA

When downlinked RA data is displayed at the CWP there will be a delay Δ_{22} before the controller reacts to it.

A former controller, with experience of STCA, suggests that the reaction time to such an alert will be 'less than a second'; we will therefore take 1s as an

¹⁹ The technical report of RA downlink methods [7] reports that one avionics manufacturer has implemented a design where the RA report is only available after the RA has terminated. It is assumed the ambiguity in the SARPs will be resolved and that such a design need not be taken into account in the model.

upper limit. As before the lower limit is taken as the quickest typical human reactions, viz. 0.25s [23]:

$$\Delta_{22} = U[0.25s, 1s]$$

The values adopted here reflect a best case. Times could be longer, for instance, if the controller's attention was initially focussed on another display. It would be appropriate to adopt a more sophisticated form for the distribution (*i.e.* with a tail extending beyond 1s) in any work building on this preliminary study, but this would necessarily depend on the manner and form of the RA downlink presentation at the CWP.

A.2.20 Controller comprehends downlinked RA

Having noticed the downlinked RA the controller must comprehend the information that is presented. The delay associated with this is Δ_{23} .

A former controller, with experience of STCA, suggests that this will take between two seconds and five seconds:

$$\Delta_{23} = U[2s, 5s]$$

A.2.21 Controller spontaneously issues an instruction (unaware of the RA)

If the controller were unaware of the RA he might issue an instruction to the aircraft. Assuming that the controller successfully delivers an instruction during the course of an RA of duration η (the difference between the time at which the RA is generated and the defining moment of the encounter), the delay, Δ_{24} , until the instruction arrives is uniformly distributed throughout the duration of the RA.

$$\eta = t_{40} - t_1$$

$$\Delta_{24} = U[0, \eta]$$

The time at which the controller communicates with the pilot is t_{25} . When both aircraft are ACAS equipped and receive RAs the time, t_{25} , is calculated as the earlier of the times at which the controller might communicate with either pilot.

$$t_{25} = \min(t_{25}(1), t_{25}(2))$$

A.2.22 Controller notices unexpected deviation (RA not followed)

When the controller is aware of the RA he will be expecting the pilot to follow the RA. If the pilot ignores the RA following, by implication, an earlier ATC

clearance, and that clearance was contradictory to the RA there will be a delay, Δ_{25} , until the controller notices this.²⁰

We would expect the lower limit of this delay to be no lower than the corresponding limit, w , of the similar delay Δ_{16} (see section A.2.15).

The upper limit will not depend upon the total number of aircraft being controlled ($N = 20$, see section A.3.8) because the controller's attention will be focussed on the RA encounter. Rather he will concentrate on the aircraft involved in the RA encounter and perhaps up to three other nearby aircraft (five aircraft in all). The upper limit will then be $5/20 = 1/4$ of the upper limit of Δ_{16} , (*i.e.* 2.5s).

We will assume that the delay is uniformly distributed between the lower and upper limits:

$$\Delta_{25} = U[w, 2.5s]$$

A.2.23 Pilot reaction time to controller message

The delay during which the pilot reacts to the controller message and then initiates a manoeuvre of the aircraft is Δ_{27} . The process is similar to the routine read-back of a controller message – the pilot must react to the controller message and then act upon it (in this case reading back the information, in the timeline initiating a manoeuvre).

Discussions with a former controller indicate that the time it takes for a pilot to acknowledge a controller message can be very quick but can also be delayed by as much as ten seconds.

This is confirmed by an examination of the routine RT traffic in the transcript of the communications in the fifteen minutes before the Überlingen midair collision [26]. The controller communicated with aircraft on 40 occasions and the delay before a reply ranged from 2s to 13s. The average delay was 4.75s.

The situation we are considering is not routine: the pilot has made a mistake and is being told by the controller to do something counterintuitive (if it were not counterintuitive the pilot would not have made the original mistake). Under these circumstances we can expect the average delay to be greater than 4.4s.

The delay is assumed to be uniformly distributed between 1s and 10s (giving a mean value of 5.5s):

$$\Delta_{27} = U[1s, 10s]$$

²⁰ The same situation could also arise if a pilot misinterprets the RA and manoeuvres in a contrary sense (although the analysis of on-board recordings in ASARP [15] concluded that this was unlikely).

A.2.24 Change in altitude

In response to a controller instruction we wish to know the time required, Δ_{28} , to change altitude by an amount, h_0 , significant in collision avoidance terms.

If the aircraft employs an acceleration, a_0 , then the time required is:

$$\Delta_{28} = \sqrt{\frac{2h_0}{a_0}}$$

The altitude change which is considered significant in collision avoidance terms will need to be slightly greater than the typical dimensions of an aircraft. A Boeing B747 is 63ft high and an Airbus A380 is 79ft high so the NMAC²¹ vertical threshold of $h_0 = 100\text{ft}$ is appropriate.

The acceleration with which a pilot responds to an RA is a , and we expect the acceleration a_0 used here to be similar.

- In the safety scenario the controller will communicate a certain degree of urgency (e.g. the instruction may be an avoidance manoeuvre) and so we can expect a_0 to be greater than a . We will assume that the acceleration is uniformly distributed between a and $1.4a$ – for the optimal pilot response this will give limits corresponding to the standard response to an initial RA (0.25g) and the standard response to an increase rate RA or an RA reversal (0.35g):

$$a_0 = U[a, 1.4a]$$

- In the operational scenario the situation is less urgent (e.g. the instruction may be a routine ATC clearance). The lower limit needs to be similar to the acceleration with which normal aircraft manoeuvres are conducted. We will assume that the acceleration is uniformly distributed between $0.5a$ and a – for the optimal pilot response this will give limits of 0.125g and 0.25g:

$$a_0 = U[0.5a, a]$$

A.2.25 Altitude data via Mode S SSR

The altitude data from the aircraft must be detected by one or more Mode S SSR and transmitted to the ATC centre. The earliest data to arrive at the ATC centre does so with a delay Δ_{31} .

²¹ NMAC = near mid-air collision – an encounter in which the horizontal separation of two aircraft is less than 500ft while the vertical separation is simultaneously less than 100ft.

The coverage of the encounter is provided by n_1 SSRs, labelled by the subscript i . Associated with each radar are delays:

- γ_i corresponding to detection;
- ξ_i corresponding to processing; and
- Δ_{37} corresponding to transmission over the ground network.

For each SSR the total delay is:

$$\mu_i = \gamma_i + \xi_i + \Delta_{37}$$

and the delay Δ_{31} corresponds to the minimum value of μ_i .

$$\Delta_{31} = \min(\mu_1, \dots, \mu_{n_1})$$

A.2.26 RA information via Mode S SSR

The RA information data from the aircraft can be detected by one or more SSR and transmitted to the ATC centre. The earliest data to arrive at the ATC centre does so with a delay Δ_{36} .

The coverage of the encounter is provided by n_1 SSRs, labelled by the subscript i . Associated with each radar are delays:

- γ_i corresponding to detection;
- Δ_{38} corresponding to processing; and
- Δ_{37} corresponding to transmission over the ground network.

For each SSR the total delay is:

$$\mu_i = \gamma_i + \Delta_{38} + \Delta_{37}$$

and the delay Δ_{36} corresponds to the minimum value of μ_i .

$$\Delta_{36} = \min(\mu_1, \dots, \mu_{n_1})$$

A.2.27 RA information via Extended Squitter

The RA information data from the aircraft can be detected by one or more ES ground stations and transmitted to the ATC centre. The earliest data to arrive at the ATC centre does so with a delay Δ_{36} .

The coverage of the encounter is provided by n_2 ES ground stations, labelled by the subscript j . Associated with each ground station will be delays:

- γ_j corresponding to detection;

- Δ_{38} corresponding to processing; and
- Δ_{37} corresponding to transmission over the ground network.

For each SSR the total delay is:

$$\mu_j = \gamma_j + \Delta_{38} + \Delta_{37}$$

and the delay Δ_{36} corresponds to the minimum value of μ_j .

$$\Delta_{36} = \min(\mu_1, \dots, \mu_{n_2})$$

A.2.28 Detection by Mode S SSR

The delay associated with the detection of data (either altitude data or an RA report) at the i^{th} SSR is γ_i . The value will depend on the phase of the radar ε_i , the rotation period of the radar T_i , and the probability of detection on each rotation of the radar r_i .

The phase of the i^{th} radar (how far it has progressed through a revolution at time zero since the target was last illuminated) is set as:

$$\varepsilon_i = U[0, T_i]$$

The rotation period can be different for each radar.

The detection probability is considered in [7]. It was assumed that each interrogation has a 90% probability of extracting the required data.

- In the worst case scenario it was considered that there will be only one interrogation within each beam-dwell. The Kairos minimum coverage scenario will therefore use a probability of detection per rotation of $r_i = 0.9$.
- In the expected scenario it was considered that there will be up to four re-interrogations per beam-dwell. The corresponding probability of detection per rotation, used in the Kairos typical coverage scenario, is $r_i = 0.996$.

The delay associated with detection by the i^{th} SSR is γ_i . This delay will consist of two elements: the delay until the target is illuminated by the SSR, δ , and potentially a further delay, equal to a whole number of rotations ($n \geq 0$), until the scan on which the data is successfully extracted.

$$\gamma_i = \delta + nT_i$$

The target was illuminated at time $-\varepsilon_i$ (by definition). At time t the target will have been last illuminated at time $((\varepsilon_i + t) \text{ MOD } T_i)$. The time until the target is next illuminated will therefore be

$$\delta = T_i - ((\varepsilon_i + t) \text{ MOD } T_i)$$

On each rotation the probability of successful data extraction is r_i . The probability of immediate data extraction is therefore r_i ; the probability that the earliest data extraction occurs on the next rotation is $(1 - r_i) \times r_i$, and so on as illustrated in Table 15.

first detection	probability	cumulative
immediately	r_i	$1 - (1 - r_i)$
after 1 rotation	$r_i (1 - r_i)$	$1 - (1 - r_i)^2$
after 2 rotations	$r_i (1 - r_i)^2$	$1 - (1 - r_i)^3$
after 3 rotations	$r_i (1 - r_i)^3$	$1 - (1 - r_i)^4$
...
after n rotations	$r_i (1 - r_i)^n$	$1 - (1 - r_i)^{n+1}$

Table 15: detection probability on successive scans.

The probability that more than n rotations are required is found as the complement of the sum of the arithmetic progression of the probabilities:

$$1 - \sum_{k=0}^n r_i (1 - r_i)^k = (1 - r_i)^{n+1}$$

We can sample the distribution of n by selecting a random number between zero and unity:

$$R = U[0, 1]$$

and finding the value that satisfies the following inequality:

$$(1 - r_i)^n \geq R > (1 - r_i)^{n+1}$$

This is equivalent to finding the largest value of n for which:

$$(1 - r_i)^n \geq R$$

$$n \log(1 - r_i) \geq \log R$$

$$n \geq \frac{\log R}{\log(1 - r_i)}$$

The value of n is an integer so it can be found directly using the truncation function (which returns the integer part of an expression) as:

$$n = \text{TRUNC}\left(\frac{\log R}{\log(1-r_i)}\right)$$

Finally, the complete expression for γ_i is:

$$\gamma_i = T_i - ((\varepsilon_i + t) \text{MOD } T_i) + T_i \times \text{TRUNC}\left(\frac{\log R}{\log(1-r_i)}\right)$$

A.2.29 Processing of altitude data by SSR

The altitude data detected at the SSR must be processed before it is transmitted to the ATC centre. The delay associated with this is ξ_i .

Discussions with an SSR expert indicate that for a monopulse SSR an azimuth delay of between 30° and 60° (depending on the number of aircraft within the beam) can occur before the appropriate correlation and processing are complete.

We will assume that the delay is uniformly distributed between 30° and 60° (*i.e.* one twelfth of a rotation and one sixth of a rotation):

$$\xi_i = \text{U}\left[\frac{T_i}{12}, \frac{T_i}{6}\right]$$

A.2.30 Detection by ES ground station

The delay associated with the detection of RA information at the j^{th} ES ground station is γ_j . The value will depend on the periodicity with which the aircraft broadcasts the ES, k , and the probability of detection of each ES at the ground station, s_j .

As in [7] the periodicity of the ES is assumed to be $k = 1\text{s}$.

The detection probability is considered in [7]. The precise value depends upon many external factors such as the range of the target and the loading of the downlink frequency. The Helios Mode S fruit model was used and a value of $r_i = 0.5$ was adopted for both the expected and worst case scenarios. The same values are used in Kairos.

The delay associated with detection by an SSR is calculated from the number of repetitions of the ES until it is successfully detected. This number is calculated in precisely the same way as indicated in section A.2.28 and the delay is:

$$\gamma_j = k \times \text{TRUNC} \left(\frac{\log R}{\log(1-s_j)} \right)$$

where R is a random number uniformly distributed between 0 and 1.

A.2.31 Data transmitted through ground network

Data from an SSR or ES ground station must be transmitted through a ground network to reach the ATC centre. The delay associated with this is Δ_{37} .

The lower limit of the delay will depend upon the configuration of the network including such factors as the distances over which data is transmitted, number of repeater stations *etc.* SSR experts indicate that a minimum value of 0.1s is typical.

The upper limit of the delay will depend upon data traffic on the network at any one time (in extreme cases, when the band-width limit is reached data must be queued before it can be transmitted). A presentation to EUROCONTROL [20] reports a value of 0.3s provided by the agency.

We will assume that the delay is uniformly distributed between 0.1s and 0.3s.

$$\Delta_{37} = U[0.1s, 0.3s]$$

A.2.32 Processing of RA information by SSR

The RA information detected at the SSR or ES ground station must be processed before it is transmitted to the ATC centre. The delay associated with this is Δ_{38} .

Discussions with an SSR expert indicate that when there is a lot of data to process this can take up to a second. A lower limit of a tenth of this value is used.

We will assume that the delay is uniformly distributed between 0.1s and 1s.

$$\Delta_{38} = U[0.1s, 1s]$$

A.3 Probabilities

A.3.1 Probability of ACAS equippage

The probability of only one aircraft in the encounter being ACAS equipped: p_1 , and the probability of both aircraft being ACAS equipped: p_2 , are taken from the ASARP encounter model.

The probabilities are altitude dependent and are given in the specification [11].

A.3.2 Pilot decides to inform controller of RA

Here we are concerned only with those cases where the pilot voluntarily informs the controller of the RA *before* it has terminated. The experience of former controllers is confirmed by the limited sample reported in [25]: in twelve alerts the pilot informed the controller in a timely manner in six cases – we therefore set $p_3 = 0.5$.

A.3.3 RA requires a deviation from clearance

The probability that an RA is incompatible with ATC clearance, and requires a deviation, is p_4 .

The precise value will depend on the characteristics of the airspace of interest and ATM procedures employed. However a number of operational and simulation based studies indicate that between 40% and 60% of RAs in European airspace can be expected to require a deviation. A value of $p_4 = 0.5$ has been adopted in Kairos.

A.3.4 RT frequency busy when pilot attempts communication

A number of studies [12] [13] [14] indicate that the occupancy of the RT frequency can vary from between 30% to 60%.

One study [14] indicates that the occupancy is generally about 50% for terminal operations and lower for en-route operations.

Kairos adopts a value of $p_5 = 0.3$ in the terminal regime and $p_5 = 0.5$ in the en-route regime.

A.3.5 Pilot complies with RA

When optimal pilot response is selected this value is set at unity, $p_5 = 1$.

When typical pilot response is selected this value is set to the value found in the ASARP study, *viz.* $p_5 = 0.9$.

A.3.6 Aircraft altitude reports quantised to 25-ft precision

Mode C equipped aircraft report altitude to 100-ft precision. Aircraft that are equipped with Mode S (which includes ACAS equipped aircraft) can report altitude with 100-ft precision or with 25-ft precision.

The probabilities adopted in the ACASA project [17] are used in Kairos:

- for ACAS equipped aircraft, $p_7 = 0.9333$;
- for unequipped aircraft, $p_7 = 0.3213$.

A.3.7 RT frequency busy when controller attempts communication

It is assumed that the controller accounts for half of the RT traffic and that pilots account for the other half.

The probability that the frequency is busy when the controller attempts a communication will therefore be half of the probability that the frequency is busy when a pilot attempts a communication. We therefore set $p_8 = p_5/2$.

A.3.8 Controller (unaware of RA) issues an instruction

The controller might happen to issue an instruction while an aircraft is experiencing an RA. It might be thought that this is more likely in the safety scenario, compared to the operational scenario, due to the proximity of the aircraft. However, it could also be argued that the proximity of the aircraft could indicate that the controller has overlooked this encounter. On balance, it was decided that Kairos will use the same probability in both scenarios.

The RT frequency occupancy is p_5 . The controller accounts for approximately half of the RT traffic. Approximately half of this traffic will be 'spontaneous' and approximately half will be replies. The frequency occupancy by spontaneous controller instructions is therefore $p_5/4$.

During an RA of duration η we expect the frequency to be busy with spontaneous controller instructions for a time $p_5 \times \eta/4$. The average duration of a message is f given by the mean of Δ_9 , and so we expect the number of spontaneous messages during the RA to be $p_5 \times \eta/(4f)$. Finally, if the controller is controlling N aircraft we expect one N^{th} of the messages to be to any given aircraft, so we set

$$p_9 = \frac{p_5 \eta}{4Nf}$$

The number of aircraft under the control of a single controller has been taken as $N = 20$.

A.3.9 Probability of horizontal manoeuvre

In the operational scenario an early CoC may occur. To determine whether this happens we need to know the probability, p_{10} , that an aircraft manoeuvres horizontally before CPA.

The probability of a turn depends on altitude and probabilities from the IAPA encounter ATM model are employed here. Allowance is made for the fact that the encounter window in the IAPA model starts 4 minutes before closest approach by scaling the probabilities to the nominal ACAS warning times. The details of the distribution with altitude band are given in the specification [11].

B. SENSITIVITY ANALYSIS

B.1 Introduction

B.1.1 General

The Kairos model has been used to investigate the latencies associated with the communication of ACAS RA information to an ATC centre on the ground.

The model uses a large number of distributions that are sampled stochastically. The distributions are characterised by parameters whose values can be varied.

The values of the latencies determined by the model are naturally dependent on the values adopted for the characteristic parameters. As discussed in the body of the report, the values of parameters associated with the responses of the pilot and controller are not known to as high a degree of precision as would be liked (see section 3.3.4). Consequently, it is of interest to know how sensitive the latencies calculated by the model are to the values adopted for these parameters.

This appendix presents the results of a sensitivity analysis which investigated the effect of variations in various parameters on key latency values. The key results are presented in B.1.2 and the details of the analysis are presented in section B.2 onwards.

B.1.2 Key results

Two latencies were found to exhibit significant sensitivity to the inputs:

- The latency of the controller becoming aware of the RA by current means.

The duration of this latency is sensitive to the duration of RT messages and is also sensitive to the probability of the pilot directly informing the controller of any RA. In each case a variation in the mean of the input results in a variation approximately half as big again in the mean value of the latency.

- The latency of a pilot achieving a significant deviation in response to a controller instruction (when the controller is aware of the RA and is able to correct an erroneous action by the pilot).

The spread of this latency is sensitive to the spread of the delay in the controller contacting the pilot should he notice an unexpected deviation. A decrease in the spread of the input leads to a proportionally larger increase in the spread of the latency.

To summarise:

- Of those latencies investigated the most sensitive were found to be:
 - the latency of the controller becoming aware of the RA by current means (*i.e.* without RA downlink);
 - the latency of a pilot achieving a significant deviation in response to a controller instruction (when the controller is aware of the RA).
- The input parameters having the greatest affect on these latencies were found to be:
 - duration of RT messages;
 - the delay in the controller contacting the pilot should he notice an unexpected deviation;
 - the time it takes the controller to notice downlinked RA information at the CWP; and
 - the probability of the pilot informing the controller of any RA directly.

B.2 Approach

B.2.1 Parameters of interest

The following parameters were selected as being of particular interest (the identifiers are those used in the specification of the parameters [10] and in Appendix A):

- Δ_7 pilot delay in contacting controller (see section A.2.7);
- Δ_9 duration of RT message (see section A.2.9);
- Δ_{10} controller reaction time to pilot communication (see section A.2.10);
- Δ_{16} controller reaction time to unexpected deviation (clearance not followed) (see section A.2.15);
- Δ_{17} controller delay in contacting pilot (unexpected deviation) (see section A.2.16);
- Δ_{22} controller notices downlinked RA (see section A.2.19);
- Δ_{23} controller comprehends downlinked RA (see section A.2.20);
- Δ_{25} controller notices unexpected deviation (RA not followed) (see section A.2.22);
- Δ_{27} pilot reaction time to controller message (see section A.2.23);
- p_3 probability that pilot informs controller of RA (see section A.3.2).

B.2.2 Latencies of interest

The following latencies (each relative to the defining moment²²) were selected for investigation in the sensitivity study (the identifiers are those used in the specification of the model [11]):

- L_{05} the latency of the earliest time at which a pilot achieves a significant deviation in response to controller instruction (in the absence of RA information) (t_{39}) relative to the defining moment (t_{40});
- L_{06} the latency of the earliest time at which a pilot achieves a significant deviation in response to a controller instruction (with RA information) (t_{30}) relative to the defining moment (t_{40});
- L_{20} the latency of the controller becoming aware of the RA through downlinked information (t_{24}) relative to the defining moment (t_{40});
- L_{30} the latency of the controller becoming aware of the RA by current means (*i.e.* via pilot communication) (t_{50}) relative to the defining moment (t_{40}).

Not all of the delays listed in section B.1 are present on each of the timelines relevant to the latencies listed above. Consequently variations in the delay parameters will affect only some of the latencies. The latencies of interest, and the delays which affect them, are indicated in Table 16. A total of fifteen combinations are relevant to the analysis.

	Δ_7	Δ_9	Δ_{10}	Δ_{16}	Δ_{17}	Δ_{22}	Δ_{23}	Δ_{25}	Δ_{27}	p_3
L_{05}	x	x	x	x	x	x	x	x	✓	x
L_{06}	x	✓	x	x	✓	✓	✓	✓	✓	x
L_{20}	x	x	x	x	x	✓	✓	x	x	x
L_{30}	✓	✓	✓	✓	✓	x	x	x	x	✓

Table 16: Latency measures that are affected by delay parameters. The delays affect only those latencies indicated by the ticks in the shaded cells.

For each combination the sensitivity of the output (*i.e.* the mean and standard deviation of the latency values) to variations in the input (*i.e.* the mean and standard deviation of the delay/probability parameters) was investigated.

²² Here the safety scenario is used so the defining moment is CPA.

B.3 Sensitivity measures

B.3.1 Baseline scenario

Two measures of the latency have been investigated for each combination:

- the mean, μ , and
- the standard deviation, σ .

A baseline simulation was conducted to obtain reference values of the mean and standard deviation, μ_0 and σ_0 respectively, for each latency. The baseline scenario used the default parameter values indicated in Appendix A, for the safety scenario (*i.e.* the defining moment is CPA) across all altitude bands, downlinking RA information via Mode S Report, with typical en-route radar coverage, and typical pilot response.

B.3.2 Variation in delay parameters

For each delay parameter variations were applied separately first to the average (the mean was increased by 1s while leaving the variance the same), and then to the spread (the variance was halved while leaving the mean the same).

For each variation the change in the corresponding measure for each relevant latency was noted. *E.g.* the mean of an input delay was varied from μ_1 to $\mu_1 + \delta\mu_1$ and the change in the mean of the latency, from μ_0 to $\mu_0 + \delta\mu_0$, was noted. Similarly the standard deviation of an input delay was varied from σ_1 to $\sigma_1 + \delta\sigma_1$ and the change in the mean of the latency, from σ_0 to $\sigma_0 + \delta\sigma_0$, was noted.

A dimensionless measure of the sensitivity of each latency to variations in the input delay was determined by comparing the change in the output produced by a given change in the input. The sensitivity to changes in the mean was calculated as

$$S_{\mu} = \frac{\delta\mu_0}{\delta\mu_1}$$

and the sensitivity to changes in the standard deviation was calculated as

$$S_{\sigma} = \frac{\delta\sigma_0}{\delta\sigma_1}$$

B.3.3 Variation in probability

For the probability parameter, the value was varied by a certain amount (changed by 10%) and the change in the latency measures noted.

The method of defining sensitivity measures indicated in section B.3.2 is not applicable when varying probabilities. Instead a dimensionless measure of the sensitivity was produced by comparing the proportional change in the output to the proportional change in the input. The sensitivity to changes in the mean was calculated as

$$S_{\mu} = \frac{\delta\mu_0}{\mu_0} \bigg/ \frac{\delta p}{p} = \frac{p}{\mu_0} \frac{\delta\mu_0}{\delta p}$$

Similarly, the sensitivity to changes in the standard deviation was calculated as

$$S_{\sigma} = \frac{p}{\sigma_0} \frac{\delta\sigma_0}{\delta p}$$

B.3.4 Interpretation of sensitivity measures

Dimensionless sensitivity measures can be calculated as outlined in sections B.3.2 and B.3.3.

The complex nature of the interaction of the various processes on the timelines discussed in section 3.2 means that a variety of behaviours of sensitivity can be expected:

- If a sensitivity measure has a value close to unity this indicates that the output varies as the input. *E.g.* any change in the mean of a delay is matched by an equivalent change (by the same absolute amount) in the mean of the latency.
- If a sensitivity measure has a magnitude greater than unity this indicates that the output varies more than the input. *E.g.* any change in the mean of a delay causes a greater absolute change in the mean of the latency (in particular this can occur when a given delay is potentially invoked more than once in a particular latency). Conversely, if a sensitivity measure has a magnitude less than unity this indicates that the output varies less than the input.
- If a sensitivity measure has a negative value then this indicates that the output varies in the opposite sense to the input. *E.g.* a decrease in the spread of an individual delay causing an increase in the spread of the latency.

B.4 Results

B.4.1 Scatter plot

Sensitivity measures, for both the mean and the standard deviation, have been calculated for the fifteen combinations of delay and latency indicated in Table 16. The results are displayed graphically as a scatter plot in Figure 6.

B.4.2 Sensitive combinations

From the scatter plot in Figure 6 we can see that three combinations are particularly sensitive to the input values (*i.e.* at least one of the sensitivity values is greater than 1.5).




Figure 6: Scatter plot latency sensitivity measures

The combinations displaying this greatest sensitivity are:

- sensitivity of the mean of L_{30} to Δ_9 ;
- sensitivity of the mean of L_{30} to p_3 ; and
- sensitivity of the standard deviation of L_{06} to Δ_{17} .

Combinations displaying moderately sensitivity are:

- sensitivity of the mean of L_{06} to Δ_9 ;

- sensitivity of the standard deviation of L_{06} to Δ_{22} ; and
- sensitivity of the standard deviation of L_{30} to Δ_{17} .

These combinations and their sensitivity are discussed further below.

B.4.2.1 *Controller awareness achieved by current means*

L_{30} is the latency of the controller becoming aware of the RA by current means (*i.e.* through a communication with the pilot – either by information volunteered by the pilot or following a dialogue resulting from an unexpected altitude deviation observed by the controller).

The mean of this latency is sensitive to the mean of the duration of RT messages (Δ_9), and also sensitive to the probability of the pilot informing the controller of any RA directly (p_3). In each case a variation in the mean of the input results in a variation approximately half as big again in the mean value of the latency.

The standard deviation of this latency is moderately sensitive to the standard deviation of the delay in the controller contacting the pilot should he notice an unexpected deviation (Δ_{17}). An increase in the standard deviation of the input leads to a larger increase in the standard deviation of the latency.

B.4.2.2 *Pilot achieves significant deviation when action is corrected by controller*

L_{06} is the latency of a pilot achieving a significant deviation in response to a controller instruction when the controller is aware of the RA and is able to correct an erroneous action by the pilot (*i.e.* the pilot following a controller instruction when the RA should take precedence).

The standard deviation of this latency is sensitive to the standard deviation of the delay in the controller contacting the pilot should he notice an unexpected deviation (Δ_{17}). A *decrease* in the standard deviation of the input leads to a proportionally larger *increase* in the standard deviation of the latency.

This latency is also moderately sensitive to the mean of the duration of RT messages (Δ_9), and also to the standard deviation of the time it takes the controller to notice downlinked RA information at the CWP (Δ_{22}).

B.4.3 **Summary**

Certain model input parameter values, associated with the responses of pilots and air traffic controllers, may not be known with great precision. The sensitivity of a four key latencies to the mean and standard deviation of these parameters has been investigated.

The latencies most affected are:

- the latency of the controller becoming aware of the RA by current means; and

- the latency of a pilot achieving a significant deviation in response to a controller instruction (when the controller is aware of the RA).

The input parameters having the greatest affect on these latencies are:

- duration of RT messages;
- the delay in the controller contacting the pilot should he notice an unexpected deviation;
- the time it takes the controller to notice downlinked RA information at the CWP; and
- the probability of the pilot informing the controller of any RA directly.