Preprint Proceedings 8th Innovative Research Workshop and Exhibition, EUROCONTROL Experimental Centre, Brétigny-sur-Orge, France, 1st - 3rd December, 2009.

Comparison of pair-wise priority-based resolution schemes through fast-time simulation

Richard Irvine
EUROCONTROL Experimental Centre,
Bretigny-sur-Orge, France.
richard.irvine@eurocontrol.int

***Abstract*— Self-separation concepts in air traffic management sometimes propose the use of priority rules to designate an aircraft which must modify its trajectory in order to resolve a predicted proximity with another aircraft. Fast-time simulations incorporating a resolution algorithm illustrate situations in which a designated aircraft is not able to resolve a predicted proximity. A resolution strategy is less likely to fail if the designation of the aircraft responsible for manoeuvring can take into account the feasibility of finding a conflict-free trajectory for that aircraft. One simple way of doing this is to allow priorities to be reversed if the first designated aircraft is unable to find a conflict-free trajectory. Further simulations allowing priority reversal only revealed irresolvable situations with short air-air datalink ranges.**

Keywords-simulator; resolution algorithm; self-separation; priority rule; probability; separation loss; iFly

I. Introduction

To resolve multiple aircraft situations in which two or more aircraft could lose separation various problems arise, including:

1. Which aircraft (one or more) should manoeuvre to prevent loss of separation?
2. Which new trajectory or trajectories will avoid loss of separation?
3. If two or more aircraft must modify their trajectories, how can one ensure that their new trajectories are compatible so that new potential losses of separation are not created?

In the current air traffic control system these questions are answered by an air traffic controller on the ground. He decides to which aircraft he will give instructions and he decides what instructions to give them in order to modify their trajectories. In the event that two or more aircraft must modify their trajectories, the controller ensures that the new trajectories are compatible and that they do not lead to new losses of separation. Within a sector the control of air traffic is centralised in the mind of the air traffic controller.

The current air traffic control system has various capacity bottlenecks. These include the physical capacity of runways and airports, but also the workload limits of air traffic controllers.

To avoid workload limitations the air traffic management research and development community has proposed possible solutions, some of which involve moving tasks from the air traffic controller to the aircraft. One of the more radical approaches is self-separation or autonomous aircraft, in which essentially all separation tasks are moved from the air traffic controller to the aircraft. In common with other candidate systems, self-separation systems must answer the questions posed above.

The EUROCONTROL FREER concept proposed the use of Extended Flight Rules in order to designate one aircraft which should manoeuvre in order to avoid a predicted proximity between itself and another aircraft [FREER, 1997]. The calculation of priority uses parameters which are available to both aircraft involved in a predicted proximity, so that they both reach the same result. By designating only one aircraft which must modify its trajectory the problem of ensuring the compatibility of the trajectories of the aircraft involved in a predicted proximity is simplified: it is responsibility of the designated aircraft to find a new trajectory which avoids loss of separation with the other aircraft involved in the proximity and with all other aircraft. The FREER Extended Flight Rules do not take into account the feasibility of finding a conflict-free trajectory for the designated aircraft. Simultaneous resolutions for different proximities could result in a new predicted proximity. Resolution of one predicted proximity may interfere with and delay the resolution of a nearby proximity.

In [FACES, 2000] a token passing scheme is used to develop a resolution order amongst a group of aircraft. The aircraft then modify their trajectories in sequence. This approach offers a solution to the third of the above problems, namely avoiding the creation of new losses of separation (due to concurrent resolutions).

In the self-separation concept described in [NASA, 2003], priority rules are used for "staggered alerting": a lower priority aircraft is alerted to a potential conflict some time before alerting a higher priority aircraft. In this way it is likely that the lower priority aircraft will resolve the conflict and the higher priority aircraft may not even need to be alerted to a predicted conflict. If the lower priority aircraft does not resolve the conflict then at some point the higher priority aircraft will be alerted and it may also manoeuvre.

The European Commission's iFly project is developing a self-separation operational concept called the Autonomous Aircraft Advanced (A3) ConOps [iFly, 2009]. In addition to Long Term Area Conflict Detection, this concept envisages Medium and Short-Term Conflict Detection and Resolution. The use of priority rules to designate an aircraft which should manoeuvre is proposed for the Medium-Term Conflict Resolution. In the event that the designated aircraft does not resolve the predicted conflict in the medium term, it should be solved by co-operative manoeuvring in the short-term.

As part of an innovative study, a fast-time simulator has been developed which incorporates a conflict resolution algorithm. This simulator can be used to test the feasibility of separating aircraft under a variety of conditions, and could be of general use within ATM research. The approach is not sufficient to demonstrate or assess safety, but it can be used to detect conditions under which the algorithm cannot separate aircraft. Conditions under which aircraft cannot be separated are, prima facie, unsafe, and highlight the need for improved system design or improved performance of system components.

As part of the Experimental Centre's contribution to the iFly project, this simulator has been used to investigate the feasibility of finding conflict-free trajectories for an aircraft designated by a priority rule. This is compared with a resolution strategy which allows priorities to be reversed when conflict-free trajectories cannot be found for the first designated aircraft.

II. Simulator Design

A. *Basic features*

The simulator is a discrete-time simulator. Aircraft performance, including ceilings, horizontal and vertical speeds, is taken from lookup tables which are part of the BADA aircraft performance model [BADA, 2004]. Flights are created at departure airports in accordance with flight plans taken from a traffic sample. Flights navigate according to a routing scheme: direct routing was used in the simulations described here. Target cruise levels are assigned to flights taking into account aircraft performance and are constrained by a flight level allocation scheme. Turns are arcs of circles, assuming a standard bank angle. The simulator includes proximity detection, trajectory prediction, proximity prediction and proximity resolution. Proximity detection allows failures of the resolution strategy to be recorded.

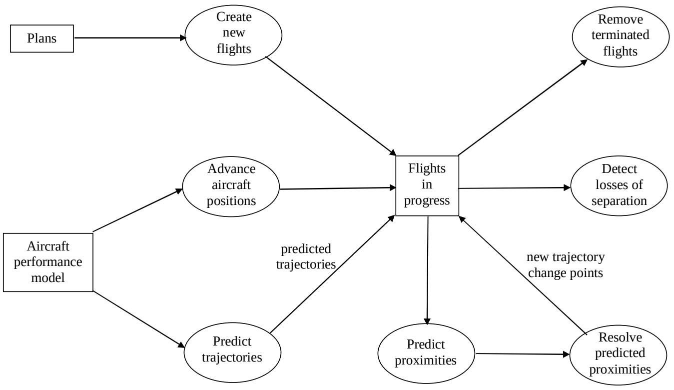

Figure 1. Main simulator functions, performed on each time step

B. *Resolution strategy*

At each time step, predicted proximities are used to initiate proximity resolution (or, more generally, trajectory replanning). A priority rule is used to designate one of the flights involved in the proximity, and a new trajectory is generated for that flight. The resolution strategy may look for solutions only for the designated (low priority) flight, or, if priority reversal is permitted, it may also consider solutions for the high priority flight.

Lateral and vertical resolution for the low priority aircraft are performed using the GEARS conflict resolution algorithm [GEARS, 1998]. This is a one-against-many algorithm which requires that the trajectories of obstacle aircraft be known. The trajectories of the obstacle aircraft may include turns, vertical movements and changes of horizontal and vertical speed. The algorithm constructs a tree of manoeuvres (of a single type) for a manoeuvred aircraft which avoid the obstacle aircraft. In lateral resolution (see figure) a manoeuvre consists of a turn to a new track angle and continuation on that track (great circle segment) once the turn is complete. In vertical resolution a manoeuvre consists of the movement to a new target flight level and continuation at that flight level when it is reached. Each path through the tree consists of a sequence of manoeuvres which constitute a resolution trajectory. One of the trajectories is selected in accordance with some criteria, e.g. proximity to destination at the end of the resolution period.

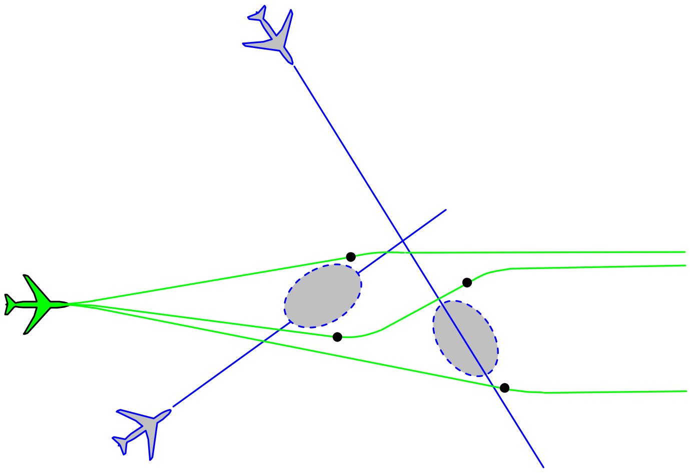

Figure 2. Tree of (lateral) manoeuvres for a manoeuvred aircraft (green) which avoid the obstacle aircraft (blue).

The algorithm yields "first" and "second" class resolutions. In both cases obstacle aircraft are avoided. A first class lateral resolution is one in which the designated or manoeuvred aircraft is heading towards its destination at the end of the resolution. In a first class vertical resolution the aircraft has reached or is moving towards its target flight level at the end of the resolution. When first class resolutions are available they are chosen in preference to second class resolutions. At present, lateral resolutions are chosen in preference to vertical resolutions (see limitations below). In principle, the algorithm finds manoeuvres which begin at a given starting time. In cases where no solutions can be found at the given starting time, solutions may be found considering later starting times.

This algorithm was chosen partly because the author is familiar with its implementation, but also because it can incorporate a wide range of proximity definitions and aircraft behaviour models, and can generate resolution trajectories which avoid reserved areas.

The resolution strategy is very limited compared to the palette of possibilities available to an air traffic controller: predicted proximities are resolved by modifying the trajectory of only one aircraft and resolution trajectories always consist of a sequence of manoeuvres of the same type, either lateral manoeuvres (without changing the vertical profile) or vertical manoeuvres (along the direct path to the destination).

C. Current limitations

The current stage of modelling has various limitations, including:

- There is no wind model - the atmosphere is still.
- There is no uncertainty modelling. Aircraft do not deviate from predicted trajectories. Effectively this corresponds to a perfect 4D flight management system.
- There are no reserved areas.
- Proximities are resolved by manoeuvring only one of the aircraft involved.
- Resolution trajectories consist of manoeuvres of a single type (lateral or vertical). This is not appropriate for long resolutions, as there is no reason why the type of manoeuvre used to avoid one aircraft should be appropriate for avoiding another. More generally resolution trajectories could consist of sequences of manoeuvres of the same or different types. Still more generally one might envisage simultaneous modifications to the path of the aircraft and to its vertical profile.
- Lateral resolutions are chosen in preference to vertical resolutions. For this reason most resolutions are lateral. However, occasionally this results in a poor choice of solution. An improvement would be to choose the resolution which minimises fuel consumption.

III. SIMULATOR CONFIGURATION

A. Goal of simulations and independent variables

The goal of the simulations was to identify conditions in which the resolution strategy fails, and to measure the frequency of occurrence of these failures. The conditions which were varied were traffic level and air-air datalink range.

B. Traffic sample

The base traffic sample used is one which has been developed for the Episode 3 project for use in initial validation of the SESAR Target Concept [SESAR]. It contains about 3 times as many flights as on the peak day in 2006. The sample contains flights in the ECAC area over a period of 24 hours. It also contains those flights which begin on the preceding day but which are aloft within the ECAC area at the start of the nominal day of the traffic sample. The traffic sample is effectively an estimation of airline demand without modification of departure times to avoid runway or airspace congestion.

C. Departure sequencing

Within the simulation most flights are allowed to depart at the departure time in the traffic sample flight plan, i.e. no runway capacity constraints are applied. However, an exception is made for flights departing in similar directions. In this case, a delay is introduced, if needed, to avoid the creation of in-trail losses of separation. The traffic sample contains many same direction flights which depart at the same time. Furthermore, where flights have been generated through cloning, the aircraft types and cruise levels are the same. Introducing a delay between successive same direction flights effectively creates trains. The size of the delay has an impact on the encounters which occur in the simulation. If the delay is small, trains will be short, and it will often be possible to find trajectories for a manoeuvred aircraft which pass behind or in front of a train. If the delay is larger, the trains will be relatively long, making it harder for a manoeuvred aircraft to pass behind or in front of them. On the other hand a larger delay creates gaps between successive aircraft in the train which will often be big enough for a manoeuvred aircraft to pass through. If the delay is too large, then it may result in aircraft being backed up at an airport, preventing flights from entering the simulation. Through a tuning exercise (using the 3x traffic sample), it was found that a delay between successive same-direction flights of 3 minutes resulted in 'permeable' trains without significantly preventing the entry of aircraft into the simulation. The sequencing of same direction departures from airports in order to create short or permeable trains may be an important operational issue as traffic levels increase. Furthermore, reserved areas can be expected to merge streams from many airports.

D. Volume of interest

Proximities are predicted and resolved for flights which are within or which are predicted to enter a measured "volume of interest". This volume was a square with a side of 500 nautical miles, including many of the major European airports, including London, Amsterdam, Frankfurt and Paris. The floor of the volume was set arbitrarily at FL245, as this is often taken as the division between the lower and upper airspace. There was no ceiling. See figure below.

The volume of interest lies within a trajectory prediction area. Trajectories are predicted for all aircraft within the trajectory prediction area which will enter the volume of interest. In this way proximities within the volume of interest can be predicted, including those involving aircraft which are currently outside of the volume of interest.

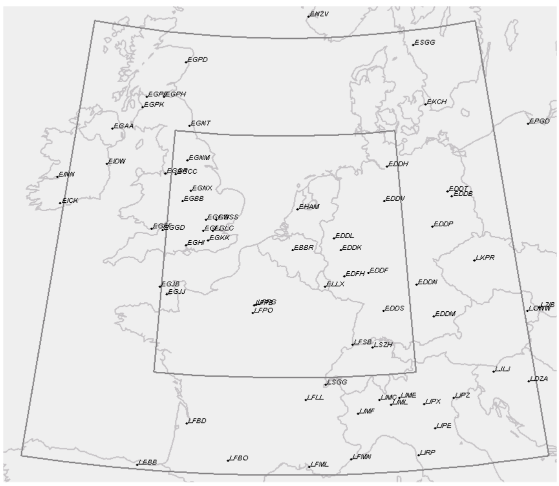

Figure 3. Volume of interest and surrounding trajectory prediction area

E. Routing and flight level allocation

Flights fly direct to their destinations as reserved areas have not yet been modelled. Trajectory re-planning to avoid other aircraft introduces deviations following which aircraft fly direct again. Target cruise levels are constrained by a semi-circular rule.

F. Aircraft performance

The BADA aircraft performance model includes data for low, nominal and high mass. In these simulations all aircraft were assumed to have nominal mass. Turn radii are calculated assuming a bank angle of 15 degrees.

G. Proximity detection

The proximity detector records losses of separation. A loss of separation is defined to occur when the horizontal distance between any two aircraft is less than 5 nautical miles and the vertical distance between them is less than 1000 feet.

H. Trajectory prediction look-ahead

For how long should future trajectories be predicted? Ideally, the future trajectories of aircraft would be available all the way to the destination, once aircraft are within air-air datalink range. Within the simulator, trajectory prediction look-ahead time directly affects the computation needed for trajectory prediction, proximity prediction and proximity resolution. To limit the simulation time to a manageable length, the trajectory prediction look-ahead was set to 20 minutes.

I. Proximity resolution

A priority rule designates a low priority flight in a predicted proximity. New trajectories are found for the low priority flight. Two sets of simulations were performed. In the first set of simulations, new trajectories were not sought for the high priority flight, in other words, priority reversal was not considered. In the second set of simulations, if resolutions could not be found for the low priority flight then they were sought for the high priority flight.

The latest time at which a resolution can be performed is 60 seconds prior to loss of separation. If a resolution is not provided before or at this time, then a loss of separation will occur.

Proximity resolution was only performed taking into account aircraft which would be within range via the air-air datalink. The air-air datalink is assumed to be perfectly reliable.

In the simulator the sequencing of proximity resolutions is centralised, which, from the perspective of self-separation, is effectively equivalent to the favourable assumption that the problem of distributed coordination of trajectories of aircraft can be solved.

IV. RESULTS

A. Separation loss counts without proximity resolution

From the base 3x 2006 traffic sample, traffic can be decreased to a given level by omitting flight plans. Traffic can be increased to a given level by cloning plans. With proximity resolution disabled, the simulator was run to count the number of separation losses which occurred in the measured volume as a function of traffic level.

Traffic level / 2006 traffic	Number of flights entering the measured volume	Total flight time in measured volume / hours	Separation losses in the measured volume	Separation losses per flight hour
1	12247	5997	2759	0.46
2	24509	11973	10912	0.91
3	36775	17905	24484	1.36
4	49018	23900	45165	1.88

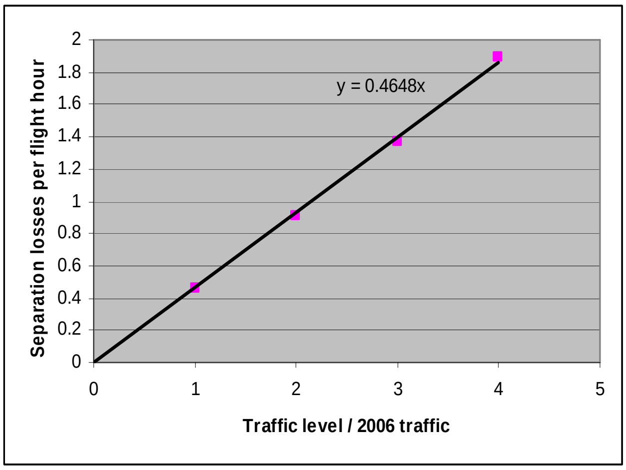

Figure 4. Separation losses per flight hour as a function of traffic level, with no proximity resolution

As may be expected, the number of separation losses per flight hour increases linearly with increasing traffic level.

B. Proximity resolution without priority reversal

In the first set of simulations priority reversal was not allowed.

1) Counts of separation loss varying air-air datalink range

The traffic sample used is one which was developed for the Episode 3 project for use in initial validation of the SESAR Target Concept. It contains about 3 times as many flights as on the peak day in 2006. The air-air data link range was varied from 24 to 88 nautical miles in steps of 8 nautical miles. The number of separation losses occurring in a period of 24 hours for each of these values is shown in the figure below.

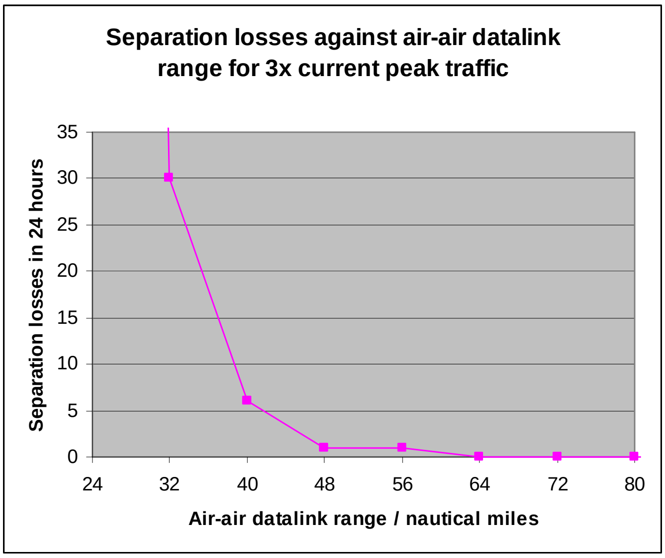

Figure 5. Separation losses against air-air datalink range for 3x current peak traffic

As the air-air datalink range is increased, the number of separation losses decreases. In view of the limited nature of the resolution strategy (designation of a single aircraft which must resolved the predicted proximity, resolution manoeuvres of a single type, i.e. lateral or vertical), it is surprising that, for air-air datalink ranges of 64 nautical miles and greater, there were no losses of separation.

2) Counts of separation loss, varying traffic level and air-air datalink range

Counts of separation losses in 24 hours for a range of traffic levels and air-air datalink ranges are shown in the table and graph below.

	Air-air range	24	32	40	48	56	64	72	80	88
Traffic level / 2006 traffic										
1		16	0	0	0	0	0	0	0	0
2		116	8	1	0	0	0	0	0	0
3		450	30	6	1	1	0	0	0	0
4		-	188	33	10	1	1	0	1	0

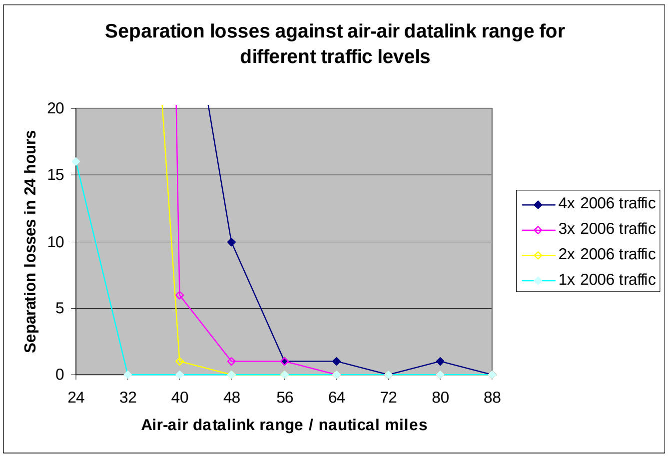

Figure 6. Counts of separation losses as a function of air-air datalink range for different levels of traffic

3) Relative frequency of separation loss per encounter (without priority reversal)

The number of encounters occurring in a random traffic sample increases as the square of the number of flights. In the simulation with 1 x 2006 traffic there are about 2600 encounters, whereas in the simulation with 4 x 2006 traffic there are about 43 000. To avoid reaching conclusions which may be related to the size of the traffic sample, it is helpful to divide by the number of encounters which occurred. In the case of separation losses, this effectively gives a relative frequency of separation loss (or failure of the resolution strategy) per encounter. As the number of encounters in a sample tends to infinity, the relative frequency (of separation loss) tends to the probability (of separation loss). (If a sample contains n encounters, it is unlikely that one will see an instance of an unresolved encounter if its probability of occurrence is less than $1/n$. For this reason, if no unresolved encounters are observed in a simulation this is not indicated as a relative frequency of zero in the following graph.)

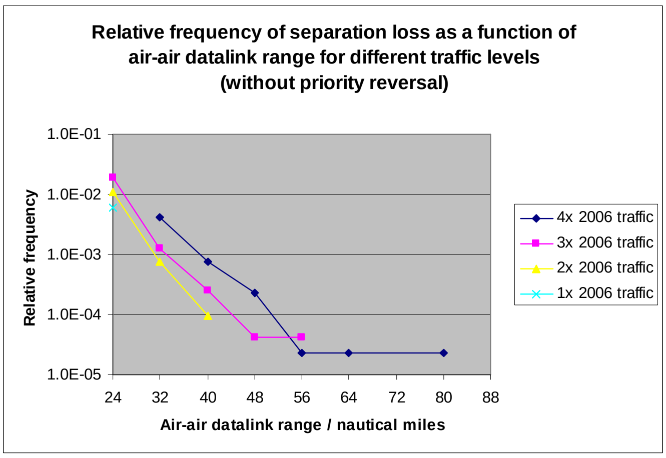

Figure 7. Relative frequency of separation loss as a function of air-air datalink range for different traffic levels (without priority reversal)

4) Estimate of an upper bound on the probability of separation loss per encounter

Under those conditions for which separation losses do not occur in the simulation it is not possible to estimate the probability of separation loss. An upper bound can be estimated from the number of encounters in the traffic sample. The 4x sample contains about 43 000 encounters and no separation losses occurred for air-air ranges of 88 nautical miles or more. There is a small probability that such a result could occur by chance. We can choose a level of statistical significance, i.e. the probability of the result occurring by chance, and then calculate the probability of separation loss for which the level of significance would be achieved.

Let p be the probability of separation loss per encounter (at the required level of significance)

q = $1-p$

n be the number of encounters in the traffic sample

ε be the required probability that the result occurred by chance

The probability of no separation losses occurring in n encounters is q^n and we require that

$$q^n = \varepsilon$$

so that

$$q = 10^{\left(\frac{\log_{10}\varepsilon}{n}\right)}$$

Taking ε = 0.01, n = 43 000, then p =1.07E-04 .

In words, applying a priority rule **without the possibility of priority reversal**, with air-air datalink ranges of 88 nautical miles or greater, the probability of separation loss per encounter in the modelled system is less than 1.1E-04 (at a level of statistical significance of 0.01).

5) Analysis of resolution failures without priority reversal

For a given traffic level the relative frequency of separation loss decreases with increasing air-air datalink range. For a given air-air datalink range, increasing the traffic increases the relative frequency of separation loss.

Because of the greater number of encounters they contain, high density traffic samples can reveal conditions which have lower frequencies of separation loss. The 3x and 4x traffic samples reveal a point of inflection in the relative frequency, beyond which the relative frequency of separation loss decreases only slowly with increasing range.

By examining encounters which cause the resolution strategy to fail a common pattern emerges:

1. An initial encounter leads to a resolution for an aircraft (aircraft A) which is conflict-free with respect to all other aircraft which are within air-air datalink range. However, the resolution effectively "boxes in" aircraft A so that it has very little room for further manoeuvre. Furthermore, this "boxing in" will persist for some time (near parallel trajectories). See figures.

2. At a later time another aircraft, whose trajectory conflicts with that of aircraft A, comes into air-air datalink range of aircraft A.

3. The priority assignment rule determines that aircraft A must resolve the predicted conflict, but since aircraft A is already "boxed in" it cannot manoeuvre to resolve the new predicted conflict.

In the example below (figure 8, 3x current traffic, air-air range 56 nm) all aircraft are cruising at FL410, except the pink aircraft at the top of the picture which is cruising at FL390. A new proximity is predicted between the blue aircraft and the pink aircraft at the bottom of the picture. The priority rule designates the blue aircraft to resolve the proximity. However, it is boxed in to the left and right by the orange aircraft and its possibilities to descend are limited by the pink aircraft at the top of the picture, which is 2000 feet below. The pink aircraft at the bottom of the picture will begin to descend near the crossing point which prevents the blue aircraft from descending by 1000 feet. The blue aircraft cannot resolve the proximity, whereas the pink aircraft has various possibilities to manoeuvre, including an earlier descent.

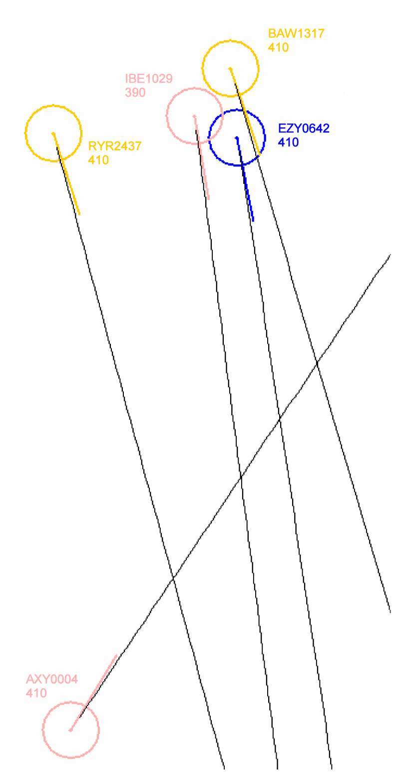

Figure 8. The aircraft (blue) designated to resolve the predicted proximity cannot find a resolution (3x current traffic, air-air range 56 nm)

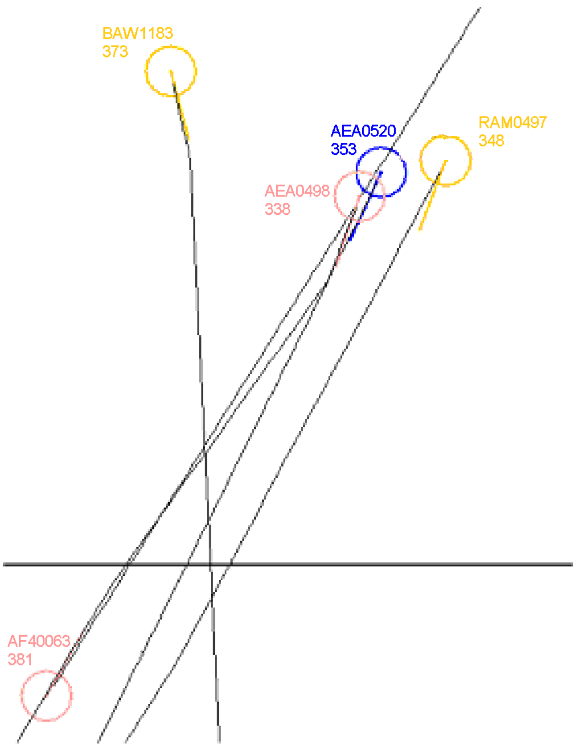

Figure 9. The aircraft (blue) designated to resolve the predicted proximity cannot find a resolution (4x current traffic, air-air range 64 nm)

6) Improving the resolution strategy

How can the resolution strategy be improved?

1. One possibility is to have a further resolution strategy to be followed if the first fails. In the iFly concept there is a short-term cooperative strategy which will be invoked if the medium-term priority-based strategy does not resolve a predicted conflict before a certain time before loss of separation. However, it might be advantageous for the failure of the first level strategy to explicitly result in the use of the second level strategy, rather than invoking the second level strategy based on time to go to loss of separation.

2. Another possibility is to generate resolutions which provide additional space around a manoeuvred aircraft so that it is not boxed in. In other words, resolutions are generated which aim to preserve the manoeuvrability of the manoeuvred aircraft, in case it needs to resolve later predicted conflicts.

3. However, when examining cases in which the resolution fails, it is often immediately obvious to a human being that the priority rule has designated the wrong aircraft to resolve the predicted conflict. This is because the priority assignment rule only takes into account the two aircraft involved in the predicted conflict and cannot "see" the positions and trajectories of other aircraft which may prevent a designated aircraft from manoeuvring. An obvious solution to this problem is to reverse the priorities of the aircraft involved in a predicted conflict if the first designated aircraft cannot find a resolution. Situations in which a given aircraft is boxed in have a low probability of occurrence. Situations in which both of the aircraft involved in a predicted conflict are themselves boxed in have a much lower probability of occurrence. If the occurrence of "boxing in" was independent for each aircraft involved in a proximity, then the probability of both aircraft being boxed in would be the square of that for single aircraft. However, one can imagine cases where the boxing in is not independent. Implementations of a distributed resolution scheme which would allow priority reversal are not discussed here.

C. Priority reversal

1) Counts of separation loss, varying traffic level and air-air datalink range

The preceding set of simulations was repeated, allowing priority reversal when a resolution could not be found for the aircraft which was first designated to resolve a predicted conflict. The results are shown in the table below:

	Air-air range	24	32	40	48	56	64	72	80	88
Traffic level / 2006 traffic										
1		2	0	0	0	0	0	0	0	0
2		2	0	0	0	0	0	0	0	0
3		15	1	0	0	0	0	0	0	0
4		33	0	0	0	0	0	0	0	0

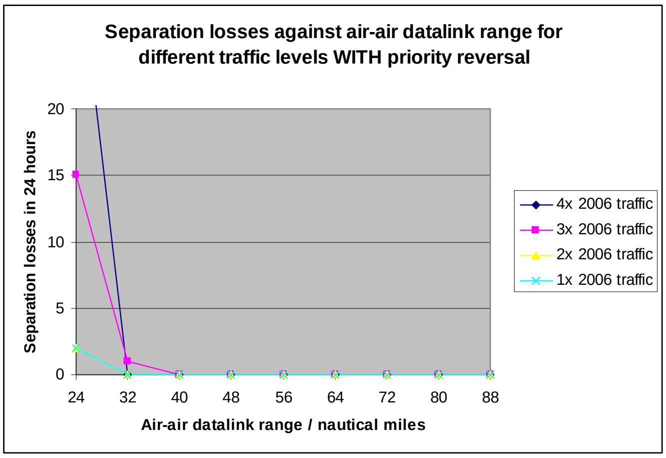

Figure 10. Counts of separation losses using a 24 hour traffic sample as a function of air-air datalink range for different levels of traffic

It can be seen that for air-air datalink ranges of 40 nautical miles or more, this simulation method was unable to illustrate cases in which the resolution strategy failed. With a range of 32 nautical miles one separation loss was generated using a traffic sample with 3x current traffic.

Referring back to the table accompanying figure 6, it can be seen that situations which cannot be resolved without priority reversal occur predominantly at low air-air datalink ranges. The greater the range the greater the number of aircraft which are visible to the resolution algorithm, and hence the greater the capacity of the algorithm to "see" later proximities. If, within the resolution process, a later proximity can be seen on a candidate trajectory then it is not a valid solution and it will not be selected. The algorithm will take account of the proximity in its search for further solutions.

At high air-air datalink ranges, situations which cannot be resolved without priority reversal are rare. Nonetheless, even with a long range, it is desirable to have a mechanism which can cope with these rare situations as early as possible, if the design is to approach a target level of safety.

In the series of simulations with 3x current traffic, 36 770 aircraft entered the volume of interest and the total flight time within this volume was about 17 900 hours. In the simulation with an air-air datalink range of 40 nautical miles, and no losses of separation, there were 23 654 resolutions, which corresponds to one resolution for every 45 minutes of flight. Since lateral resolutions were selected in preference to vertical resolutions 99.6 % of resolutions were lateral. The average route length extension within the volume of interest was 0.2 %, but it should be remembered that there were no reserved areas, and there was no uncertainty in the future positions of aircraft. The greatest deviation introduced by a lateral resolution was 52.8 nautical miles, and the greatest deviation introduced by a vertical resolution was 4000 feet. The greatest extra distance flown by any aircraft was 59.8 nautical miles.

2) Estimate of an upper bound on the probability of separation loss per encounter

Under those conditions for which separation losses do not occur in the simulation it is not possible to estimate the probability of separation loss. An upper bound can be estimated from the number of encounters in the traffic sample. The 4x sample contains about 43 000 encounters, and no separation losses were seen for air-air ranges of 40 nautical miles or more. In the same way as described earlier, one can calculate that, applying a priority rule with the possibility of priority reversal, for air-air datalink ranges of 40 nautical miles or greater, the probability of separation loss per encounter in the modelled system is less than 1.1E-04 (at a level of statistical significance of 0.01). To obtain a better upper bound or an estimate of the probability of separation loss would require simulation with a greater number of encounters.

V. Future work

Because of the low probability of separation loss and the use of 24 hour traffic samples, failures of the priority reversal strategy could only be detected for very short air-air data link ranges. Two approaches are being considered in order to better quantify the probability of failure of this strategy. The first is to improve the modelling, in particular to include reserved areas. Reserved areas will tend to concentrate traffic on to certain paths, thereby creating high local traffic densities. This is likely to increase the probability of resolution failure. The second approach will be to use longer traffic samples, containing greater numbers of encounters.

VI. Conclusions

Two resolution strategies involving the use of a priority rule were simulated. Sequencing of resolutions was centralised in the simulator. In the first strategy one aircraft involved in a predicted proximity was designated to resolve the proximity. In the second strategy, the second aircraft involved in the proximity could be designated to resolve the conflict if the first aircraft was unable to find a resolution. The traffic samples used ranged from current traffic to 4 times current traffic but were limited to 24 hours duration.

The second strategy performed very much better than the first.

In the first case, it was possible to illustrate a range of conditions under which separation losses occur. In the second case, it was only possible to illustrate failures of the resolution strategy with short air-air datalink ranges.

Under those conditions for which separation losses do not occur in a simulation it is not possible to estimate the probability of separation loss. Upper bounds can be estimated from the number of encounters in the traffic sample. For the resolution strategy without priority reversal it is estimated that the probability of separation loss (in the modelled system) is less than 1.1E-04 per encounter for air-air datalink ranges of 88 nautical miles or more. For the resolution strategy with priority reversal it is estimated that the probability of separation loss (in the modelled system) is less than 1.1E-04 per encounter for air-air datalink ranges of 40 nautical miles or more.

It is hoped that this work will contribute to the refinement of the iFly Autonomous Aircraft Advanced (A3) concept.

VII. Acknowledgements

The conflict prediction and resolution features of the simulator used for this study were developed during a one year innovative study at the EUROCONTROL Experimental Centre. The simulations were performed as part of the Experimental Centre's contribution to the European Commission iFly project.

I would like to thank Jerome Borowiak for making available a machine and operating system which have enabled simulations to be conducted rapidly.

VIII. References

[FREER, 1997] Vu N. Duong, Eric Hoffman, Jean−Pierre Nicolaon, "Initial Results of Investigation into Autonomous Aircraft Concept (FREER−1)", 1st USA/EUROPE ATM R&D Seminar, Saclay, France, June 17−20, 1997.

[FACES, 2000] Nicolas Durand, Jean-Marc Alliot, Geraud Granger, "FACES: a Free flight Autonomous and Coordinated Embarked Solver", ATC Quarterly, 2000.

[NASA, 2003] David J. Wing, Dr. Karthik Krishnamurthy, Richard Barhydt, Dr. Bryan Barmore, "Pilot Interactions in an Over-constrained Conflict Scenario as Studied in a Piloted Simulation of Autonomous Aircraft Operations", 5th USA/Europe Air Traffic Management R&D Seminar, Budapest 23-27 June 2003.

[iFly, 2009] Gustavo Cuevas, Ignacio Echegoyen, José García García, Petr Cásek, Claudia Keinrath, Rosa Weber, Petr Gotthard, Frank Bussink, Aavo Luuk "Autonomous Aircraft Advanced (A3) ConOps", iFly Deliverable D1.3 version 2.7, May 2009. Available at http://iFly.nlr.nl

[BADA, 2004] A. Nuic, "User Manual for the Base of Aircraft Data (BADA) - Revision 3.6", EUROCONTROL Experimental Centre Note No. 10/04, July 2004.

[GEARS, 1998] Richard Irvine, "The GEARS Conflict Resolution Algorithm", Proceedings of the AIAA Guidance, Navigation and Control Conference, Boston, MA, August 10-12, 1998.

[SESAR, 2007] SESAR Consortium, "The ATM Target Concept", SESAR Definition Phase Deliverable 3, September 2007.