

# An Implementation of a Reactive Distributed Air Conflict Resolution System\*

Jean-François Bosc  
ENAC<sup>†</sup> Département MI  
31055 Toulouse CEDEX 4  
FRANCE

Garfield Dean  
Eurocontrol, B.P. 15  
91222 Bretigny sur Orge CEDEX  
FRANCE

## Abstract

*The fast growth of air traffic that began during the 80's raises some questions about what future Air Traffic Control (ATC) systems will need to be. A major change will probably be required in Europe within a 30 years timeframe. Fortunately, recent developments in air navigation systems provide new opportunities for efficient autonomous navigation. Several studies have been conducted in order to develop automation in ATC, based either on centralized or distributed systems. An important part of the ATC task is to avoid separation losses (which are called "conflicts") between aircraft. This article presents an implementation of a reactive distributed conflict resolution method on a traffic simulator providing a realistic traffic sample. The method is described, and some results regarding its efficiency and its degradation with traffic increase are given. In particular, a critical traffic level appears, above which efficiency drops dramatically. This level depends on the method used; the airspace itself is not saturated.*

## 1 Introduction

Since air travel began to increase rapidly in the late 50's, Air Traffic Control (ATC) has been performed in a similar manner. Aircraft follow airways, which consist in a succession of route segments linking crossing points defined by navigational aids (radio-electrical beacons providing distance or bearing information). Air traffic controllers provide instructions to pilots in order to ensure fluidity and safety of the traffic. In particular, minimum horizontal or vertical separations have to be maintained between aircraft in order to avoid collisions. A potential violation of these separations is called a conflict. One important part of the control task consists in detecting conflict situations

and modifying aircraft trajectories to prevent any loss of separation from happening.

Airspace is divided in "control sectors". A sector is a volume of space defined by lateral, upper and lower limits, which comprises a few route segments and crossing points. Generally the size of a sector corresponds to 15 mn flight, or about 100 NM. Each sector is placed under the responsibility of two controllers (one planning controller and one tactical (or radar) controller). The main limitation of the system is the number of aircraft that a controller can handle simultaneously, which is about 15. Therefore there must never be more than 15 aircraft in the same sector at the same time. The division of airspace into sectors is intended to maximize the flow of traffic that can be handled. Reducing the size of sectors beyond a certain limit can't be a solution. A minimal transit time in sectors is required so that the controller can grasp the situation and achieve the necessary maneuvers. Moreover, the transfer of a flight between two sectors generates a specific workload ("coordination" workload).

The whole ATC system can be viewed as a succession of filters with decreasing time horizons, designed to make sure that a controller will never be overloaded. These filters are :

- Airspace Management (sector and route design)
- Air Traffic Flow Management (ATFM) (attribution of routes and take-off times, ...)
- Medium term (strategic) planning (performed by planning controller)
- Radar control (performed by tactical controller)
- Short-term collision avoidance systems (ground or airborne), which provide warnings and in some cases suggest avoidance maneuvers.

\*Supported by Eurocontrol Experimental Center

<sup>†</sup>Ecole Nationale de l'Aviation Civile / National School for Civil Aviation

Some recent technological developments in navigation system, e.g Flight Management Systems (FMS) and Global Positioning Satellite (GPS), offer the possibility of suppressing some restrictions in navigation, in particular the requirement of following pre-defined route paths. Modern FMS can provide precise navigation to any GPS coordinate, without any need for radio beacons. Apart from ATC requirements and restricted areas (military in particular), each equipped aircraft could now fly its own optimal route to its destination airport. These capabilities are not used in the current system. Allowing free navigation to flights (which is sometimes referred to as the "free route" concept) would provide significant benefits in flight time and fuel consumption. A simulation conducted with a 1-day traffic sample including all controlled flights in the French airspace indicated a 6.4% reduction in flight time.

Moreover, the recent increase of air traffic, which is expected to go on during the next 10 or 20 years, has raised some questions. In particular, many people believe that an ATC system based on control sectors where control is performed by human beings will not be able to handle the traffic density expected around 2015-2020 ([2], which considers scenarios ranging from 50% to 500% increase for the 1990-2020 period).

From that perspective, several attempts to develop automation in ATC have been made since the early 80's ([6, 7, 5, 1]). A fully automated ATC system could still be centralized, or totally distributed. In a centralized ground system, aircraft follow instructions, with possibly some freedom for negotiation. In a distributed system each aircraft would have to negotiate its trajectory with surrounding aircraft in order to ensure avoidance. The present article describes an implementation of a distributed conflict resolution system.

## 2 CATS : Caml All-purpose Traffic Simulator

The LOG<sup>1</sup> Team developed a traffic simulator in order to provide a realistic environment to test various algorithms developed by the team, that deal mainly with ATFM and medium- or short-term conflict resolution. The simulator uses flight plan data from the French ATC system (CAUTRA) archive, and a performance model coming from the ENAC ATC simulator which is used for ATC controllers training. The traffic sample used for simulation includes all controlled flights in the French airspace for a particular day. The

flights can follow their scheduled route (just as they really did) or fly straight to their destination. The simulator provides realistic control-free flight profiles.

It is possible to change several parameters, in particular vertical and horizontal separations. Typical values used here are 6 nautical miles horizontally and 1000 ft vertically. A traffic increase can be simulated by dividing all take-off times by a constant factor. During time periods where traffic is stable, there is a very strict correlation between the number of aircraft in flight and that factor, which therefore represents global traffic density (the density of the original sample being equal to 1).

Results are normally observed in a time window chosen during peak hours, at a time where traffic is reasonably stable.

The simulator has been used for testing resolution algorithms developed by the team. These algorithms are based on genetic global optimization techniques for medium-term control ([4]), and on reactive techniques ([9]) for short-term resolution. In the near future, some tests will be made with ATFM algorithms.

## 3 Reactive conflict resolution

The first resolution method implemented in the simulator was based on reactive techniques, which consist in altering aircraft trajectories in real-time (without planning) to ensure separation. The main goal was to provide a reference for comparison with genetic techniques, which are expected to be more efficient.

For a particular aircraft, trajectory alterations are derived from the computation of forces induced by surrounding aircraft that may cause a conflict. Karim Zeghal provided a complete theory of interaction between mobiles, for the purpose of avoidance or interception ([9, 8]). In particular, he introduced a tangential force to ensure bypassing of obstacles (either fixed or mobile). The main advantage of this method is simplicity. The minimal set of information concerning an intruder aircraft is distance, altitude, azimuth, relative speed, and the intruder's type of behaviour (avoidance with an identical or different method, interception). A force can then be computed for each intruder aircraft. Forces corresponding to different intruders are summed, and the results define the course change. Some additional information, in particular planned changes in course or vertical speed, can help to distinguish intruders that constitute a real threat.

It is clear that this kind of method can be distributed. Necessary information can easily be exchanged via an interrogation-response process similar to that of secondary surveillance radar (SSR)<sup>2</sup>. Each

---

<sup>1</sup>Laboratoire d'Optimisation Globale / Global Optimization Laboratory

---

<sup>2</sup>An SSR radar sends an interrogation message that is replied

aircraft can then compute its own force. With a time horizon of 5 minutes, the number of intruder aircraft is less than 10 in almost all cases at current traffic densities. Therefore the computing power required for each aircraft is very limited. On the other hand, computing maneuvers on the ground would require the assignment of a specific portion of airspace to each computing center, and therefore a special treatment near the borders (just like with current control sectors).

#### 4 Implementation

The reactive resolution method used here has been adapted from Karim Zeghal. In a first step some additional limitations have been put on aircraft maneuvers. Normally, forces can alter aircraft motion in three dimensions. However, commercial aircraft performance in cruise flight is optimized for a particular speed and altitude. Flying at a lower altitude is less efficient, and a higher altitude may be unreachable because of engine limitations. Moreover, a climb maneuver near the maximum altitude is costly in time and fuel. Regarding speed, changes are limited to 2 to 3% of normal cruise speed, and must be planned a very long time ahead in order to produce some effect. Karim Zeghal described a way to integrate performance constraints in the computation of forces. However, because of the reasons listed above, altitude changes are not used by ATC services for cruise flights, and speed changes are rarely used. Since our studies mainly deal with en-route traffic, we decided to only allow lateral maneuvers. Even though the main reason was simplicity, this restriction is very close to reality.

However, this raised some difficulties in the case of aircraft flying parallel track, eg for an overtaking. When the two aircraft fly side by side, the forces are longitudinal and tend to increase the speed of the faster one and slow down the slower one. To avoid disruptive reactions, an additional repulsive component has been added to the force vector. The repartition between the two components depends on the conflict geometry. The resulting force for a particular aircraft  $A$  and a particular intruder  $I$  (see figure 1) is :

$$\vec{F} = F \cdot \{\cos(i) \cdot \vec{u_T} + (1 - \cos(i)) \cdot \vec{u_R}\}$$

where  $F$  is the intensity of the force,  $\vec{u_T}$  and  $\vec{u_R}$  unary vectors in the tangential and repulsive directions respectively, and  $i$  the angle between the intruder's azimuth and the relative speed vector  $\vec{V_r}$ .

---

to by an on-board device called a "transponder". Current reply messages (mode-C) include aircraft identification and altitude, but in the near future so-called Mode S messages will be much longer.

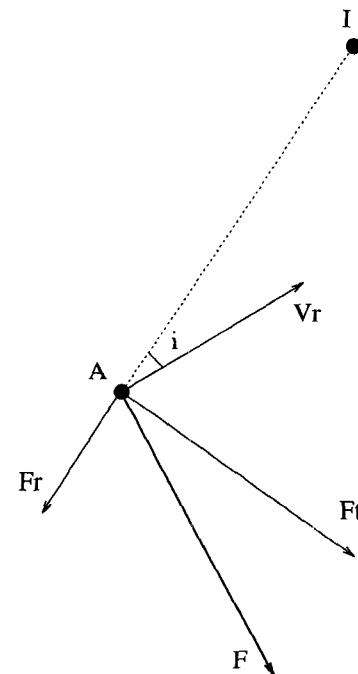


Figure 1: Force vector

Therefore  $\vec{F}$  is purely tangential if both aircraft are on a colliding course and purely repulsive if they fly on parallel tracks. The tangential component, which determines the characteristics of the maneuver (which aircraft passes ahead of the other), complies with the initial geometry of the conflict, ie the position of the relative speed compared to the intruder's azimuth.

Moreover, risk criteria have been suppressed. The intensity of the avoidance force is proportional to the inverse of the time remaining until normal separation will be lost, assuming that the closing speed will remain constant.

An attraction towards the aircraft's destination is added to the total force-vector (the current implementation only allows straight route navigation). The aircraft then tries to reach the resulting direction. There's a limit of  $3^\circ/s$  on the turning rate. It is moreover assumed that all aircraft in the simulation behave exactly the same way.

In most cases this method ensures avoidance and generates smooth trajectories (which however are not flyable by a human pilot since during maneuvers heading changes slightly at every step of computation). Some problems still occur in cases where the angle between the two trajectories is small and the ratio of the speeds is close to 1. These cases are known to be the most difficult to solve, in the sense that the

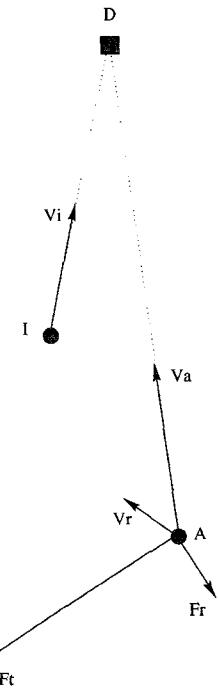


Figure 2: Improper resolution

avoidance maneuver takes more time and generates a greater increase of flight time or distance ([3]).

In figure 2, aircraft  $A$  and its intruder  $I$  (slower) are heading towards the same destination  $D$ . Speeds are  $V_a$  and  $V_i$ , relative speed is  $V_r$ . Generated forces are  $F_t$  (tangential) and  $F_r$  (repulsive). In this case, the crossing geometry given by  $V_r$  ( $A$  passes behind) is incorrect. After the maneuver, the two aircraft will find themselves in a symmetric situation ( $A$  on the left,  $I$  on the right). Another (similar) maneuver will begin, and so on, until eventually separation is lost.

Two versions of the simulator have been tested. One is truly reactive, all intruder aircraft are taken into account, which includes many aircraft that don't really constitute a threat. The only parameters used to distinguish aircraft that may generate a conflict are distance and closing speed, therefore the selection process can't be very efficient.

In the other version, aircraft trajectories are first simulated (without resolution, ie aircraft fly straight to their destination) to detect future conflicts. Typically the simulation occurs every 1 to 5 minutes, and spans over the next 5 to 20 minutes. An increased horizontal separation (usually twice the normal separation) is used in order to take into account aircraft flying near a conflict, which may interfere with the resolution maneuver. Even so this method dramat-

ically reduces the number of aircraft pairs to consider. Aircraft involved in conflicts are then grouped in "clusters" by transitive closure. The computation of reactive forces is then only applied to pairs of aircraft belonging to the same cluster. Compared to the purely reactive method, the time of computation is approximately divided by 6, despite the time spent in pre-simulation of trajectories.

The information necessary (knowledge of planned future trajectories) is not available in the current system. However, it may be approximated over a small period of time (eg 5 minutes), provided that planned changes in vertical speed or heading are transmitted between aircraft. In the future, aircraft may have a 4-dimension flight plan valid for the next 20 to 30 minutes of flight. It will be desirable to have descriptions of trajectories defined and transmitted between aircraft and/or ground systems. For example, the ARC2000 simulator ([7]) is based on the planning of aircraft trajectories described by 4-dimensional tubes, which are transmitted and negotiated between aircraft and the control system.

Both versions (i.e., with and without conflict pre-detection) give similar results regarding the number of unsolved conflicts. The knowledge of future trajectories reduces the amount of processing (but the volume of data to transmit is increased), and also the number of useless avoidance maneuvers. The average increase in flight time induced by maneuvers is 0.3% (compared to 1.5% with the purely reactive version), which is very low.

## 5 Results

Only true "en-route" conflicts are taken into account. There's a lower bound on altitude (6000 ft), and the first and last 5 mn of each flight are not considered either. This is to suppress problems due to aircraft entering or leaving the simulated airspace at the same point and time. In the first case separation is lost as soon as the planes are created in the simulator, in the second the proximity of the exit point hinders avoidance maneuvers.

Figure 3 shows the evolution of the number of separation losses versus traffic density, both without (first curve) and with (3 others) conflict resolution. Here resolution is performed with pre-detection of conflicts, with different conditions for each curve : respectively every 5 mn with twice the normal separation, every mn with twice the normal separation, and every 5 mn with one normal separation. A few corresponding values are presented in table 1. Figure 4 shows the percentage of unsolved conflicts relative to the number of separation losses observed without resolution. The

Traffic	std	+50%	+100%	+150%	+300%	+500%
no resol	508	748	962	1266	2071	2953
2 sep, 5 mn	7 (1.4%)	24 (3.2%)	38 (4.0%)	89 (7%)	1487 (72%)	
2 sep, 1 mn	10 (2.0%)		22 (2.3%)		112 (5.4%)	739 (25%)
1 sep, 5 mn	18 (3.5%)		72 (7.5%)		197 (9.5%)	545 (19%)

Table 1: Remaining conflicts.

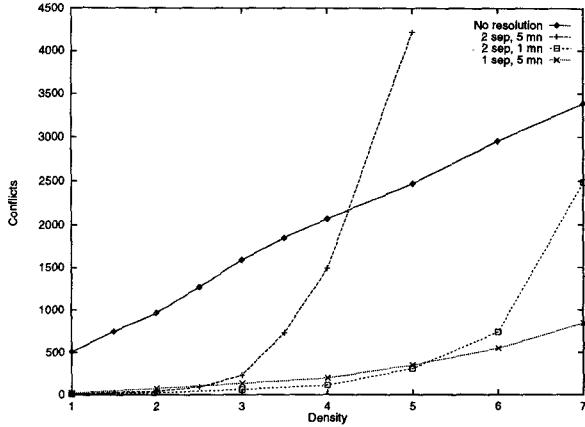


Figure 3: Number of separation losses vs density

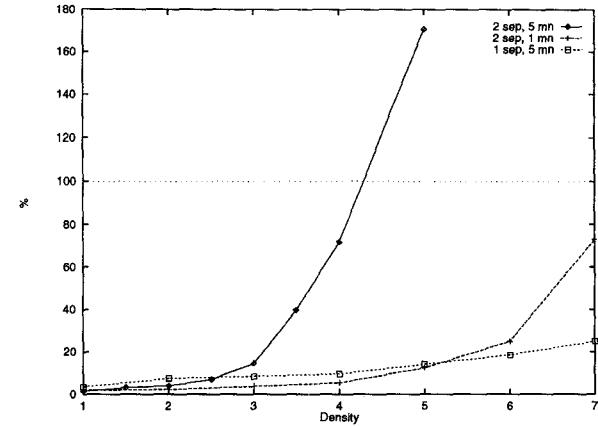


Figure 4: Percentage of unsolved conflicts vs density

results here are observed on the full traffic sample (1 day), which explains the linear shape of the conflict number without resolution. If traffic was stable during the whole simulation, this would be quadratic.

At normal traffic density, all the 7 remaining conflicts either have the same geometry described above (small angle and speed ratio), or occurred near the entry- or exit-point of the flights involved. It appears that density itself is not a problem.

When density is increased, the curve clearly shows the saturation of the resolution method. The percentage of unsolved conflicts is pretty good up to density 2.5, but it increases dramatically when density exceeds 3. At some point the resolution method generates more conflicts than it solves, because of “oscillations” of trajectories generated at high densities. It is interesting to note that with ARC2000 saturation appears at a traffic density that we believe to be similar. However, the ARC2000 traffic sample has been generated manually in order to maximize its complexity, and the resolution process itself is submitted to much stricter operational constraints.

It is possible to delay saturation when density increases. The previous results were obtained with conflict pre-detection performed every 5 mn on a 15 mn period, which is not optimal. A higher frequency

allows more accurate updating of clusters, while a shorter period avoids treatment of conflict that are too far ahead in the future. With a frequency of once every mn and a period of 5 mn, performance remains similar at low density, but saturation appears only at densities higher than 5.

Another possibility is to use a smaller horizontal separation for pre-detection. On both figures, curves “1 sep” are obtained with pre-detection every 5 mn, using the normal horizontal separation. This limits the number of aircraft pairs that are treated simultaneously, which is the cause of saturation. At some point the trajectories of aircraft submitted to multiple forces become so irregular that more conflicts are generated than solved (which doesn’t mean that no conflict-free trajectories exist). The curves show that saturation appears much later. However performance is slightly worse at low density (more than twice as many conflicts remain unsolved). This is because with a reduced separation for pre-detection, maneuvering aircraft may create new conflicts with surrounding aircraft not initially involved. Those conflicts won’t be treated until the next pre-detection step, which of course is unacceptable in a real system.

## 6 Conclusion

Some tests should be conducted with different traffic samples. However, the first results obtained are excellent. At low densities, unsolved conflicts are mostly due to conflict geometry (as has been explained above). It is clear that at current traffic density the resolution method is far from being overloaded. Therefore, a special treatment of "difficult" geometries and of flight exits might ensure the resolution of all conflicts in the sample.

With an additional filter to solve possible failures, a simple reactive method like the one used here may even have the capacity to reach the desired level of safety (which is around  $2.5 \times 10^{-9}$  collision per flight hour). For example, the horizontal separation required for such a filter could be only half of the normal separation. This would enable the resolution of remaining conflicts without interfering with those already solved by the previous filter.

Our next step will be to use genetic and reactive techniques in a complementary way. Genetic conflict resolution is performed every 5 mn, and it plans trajectories for the next 20 mn. Only the first 5 mn of each plan are definitive. Since uncertainty on future trajectories is taken into account, the rest of the plan can be computed again at the next step with better results because trajectories will be known more accurately. Tests have been made on particular conflict situations and gave good results with large numbers of aircraft (up to 50).

In case of a failure of the genetic resolution, the reactive method can be used as a secondary filter within a time horizon of 5 mn, and with a smaller horizontal separation. It will then be possible to evaluate the performance of the two methods used separately or complementarily. Considering the performance level of the

genetic method, some more significant data should be obtained regarding airspace saturation. Points of interest include the maximum traffic density acceptable in the airspace given some separations and aircraft performance (speed, turning rate), regardless of the resolution method, and whether saturation is created by aircraft density in a portion of airspace or by some particular conflict situations.

## References

- [1] Jean-Marc Alliot. *Techniques d'optimisation stochastique appliquées aux problèmes du trafic aérien*. Thèse d'habilitation, INPT, Mai 1996.
- [2] I. C. Berriman et al. ATLAS Phase II, Volume 2.5 : Feasibility of the Operational Requirement and Representative Operational Philosophies. Interim Report 3.0, PA Consulting Group, October 1993.
- [3] Nicolas Durand. Modélisation des trajectoires d'évitement pour la résolution de conflits en route. Technical report, CENA, Janvier 1994.
- [4] Nicolas Durand. *Résolution optimale de conflits en route*. PhD thesis, INPT, Mai 1996.
- [5] Xavier Fron, Bernard Maudry, Jean-Pierre Nicolaon, and Jean-Claude Tumelin. ARC2000 : Automated Radar Control. Technical report, EUROCONTROL.
- [6] Lawrence Goldmuntz et al. Automated En-Route Air Traffic Control (AERA) Concept. Technical Report FAA-EM-81-3, FAA, March 1981.
- [7] ARC2000 Team. ARC200 Technical Report. EEC No. 274, EUROCONTROL, May 1994.
- [8] Karim Zeghal. A Reactive Approach for Distributed Air Traffic Control. In *International Conference on Artificial Intelligence & Expert Systems*, Mai 1993.
- [9] Karim Zeghal. *Vers une théorie de la coordination d'actions - Application à la navigation aérienne*. PhD thesis, Université Paris VI, Décembre 1994.