

CATS: A complete Air Traffic Simulator

Jean-Marc Alliot, Jean-François Bosc, Nicolas Durand, Lionel Maugis

► **To cite this version:**

Jean-Marc Alliot, Jean-François Bosc, Nicolas Durand, Lionel Maugis. CATS: A complete Air Traffic Simulator . 16th AIAA/IEEE Digital Avionics Systems Conference (DASC), Oct 1997, Irvine, CA, United States. IEEE, 2, pp.8.2-30 - 8.2-37, Reflections to the Future. Proceedings <10.1109/DASC.1997.637271>. <hal-01291406>

HAL Id: hal-01291406

<https://hal-enac.archives-ouvertes.fr/hal-01291406>

Submitted on 21 Mar 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

CATS: A COMPLETE AIR TRAFFIC SIMULATOR

J.M. Alliot

J.F. Bosc

N. Durand*

L. Maugis[†]

1 INTRODUCTION

In this document, we present the CATS¹ system. The main ideas behind the design of CATS were to have a general purpose simulator able to be a stable basis for the development of other modules, representing different levels of Air Traffic Control and Air Traffic Management.

In the first part of this paper, the basic framework and hypothesis of the CATS simulator are introduced. The second part deals with the general mathematical concepts regarding conflict resolution and airborne collision avoidance systems. Different results obtained so far with modules regarding conflict resolution, avoidance, or workload, built on top of the CATS simulator are presented as examples in the last part.

2 CATS SYSTEM

Planning, or making decisions for the evolution of the European Air Traffic Management System (EATMS) involves a careful evaluation of alternative scenarios. This assessment is often carried out through experimentations, in which modeling and simulation contribute significantly, by reducing the turn-around time between the design and the implementation of advanced operational concepts.

Although numerous mathematical models and fast time simulation facilities are already being used to assess the performance of the European airspace system (NASPAC, EAM, TAAM, ARC2000, RAMS, SPECTRA), recent review studies ([Sof96a, Sof96b, NAS96]) highlighted also their lack of flexibility to cope with new concepts, with complex dynamics of the ATM system, with large scale uncertainty factors, and with simulation software complexity.

Flexibility is a key feature of the simulator as new filters are easily introduced, and tuned to model spe-

cific features of an operational scenario. The core of the traffic simulator is small and easy to maintain. It is less than 3,000 lines of ML code². It works with classical aircraft performances, with real airspace designs and operational constraints. Furthermore, it provides a practical framework for thorough mathematical model development, for route network and sector design, air traffic assignment, conflict resolution and airborne collision avoidance, ground holding,

2.1 Hypothesis of CATS

The core of the CATS system is an en-route traffic simulation engine. It is based on a discrete, fixed time slice execution model: the position and velocities of aircraft are computed at fixed time steps, usually every 5, 10 or 15 seconds.

Aircraft performances are in tabulated form describing ground speed, vertical speed, and fuel burn as a function of altitude, aircraft type and flight segment (cruise, climb or descent.) Two main datasets for aircraft flight performance are used:

- CAUTRA / ENAC³ aircraft performance tables, extracted from the French flight data processing system;
- Base of aircraft data (BADA) performance summary tables derived from the total energy model of EUROCONTROL. 69 different aircraft types are described. Synonym aircraft are used to model the rest of the fleet. The Airbus A320 (EA32) is used as default aircraft.

Aircraft use different navigation modes: they can fly direct routes to their destination, or follow the sequence of navigation aids described in their flight plan. Although there is no route orientation scheme explicitly defined for the air traffic network, aircraft may use one level out of two based on its heading (north/south or east/west route orientation). Every four levels can be also used based on the (NE,SE,SW,NW) aircraft's orientation (quadruple alternate).

*Laboratoire d'Optimisation Globale, CENA ENAC, 7 av Ed Belin, 31055 Toulouse Cedex France, durand@recherche.enac.fr

[†]Centre d'Etudes de la Navigation Aérienne, France

¹CATS stands either for Complete Air Traffic Simulator or for CAML Air Traffic Simulator

²CAML is a polymorphic, strong typed with inference typing language

³Ecole Nationale de l'Aviation Civile

Different separation standards can be introduced.

Uncertainties on speed (either vertical or horizontal) can be modified, and a convex conflict detection (detailed in part 3.1) is used.

2.2 Outputs of CATS

The simulation system records and computes the following output variables:

- instantaneous aircraft count per sector;
- aircraft flow rates through sectors;
- conflict statistics (geometry, aircraft maneuvers, duration, etc);
- conflict resolution statistics (number of maneuvers, number of clusters, clusters sizes, maneuvers duration, delays due to maneuvers, type of maneuvers)
- airborne separation and collision avoidance system statistics (number, duration, type of maneuvers);
- statistics from other filters such as ground delays, runway capacity utilization are also available.

3 CONFLICT SOLVER MODELING

Among the problems that can be tested on real traffic with CATS, two approaches of automatic conflict resolution are detailed as examples in this part. The first one deals with a centralized control system for high density areas and medium term (10 to 15 minutes ahead) control. The second one deals with an on board reactive short term (3 to 10 minutes ahead) automatic conflict solver.

As both solvers use a conflict detection that is robust to uncertainties, a convex modeling of uncertainties is first introduced.

3.1 Convex modeling of uncertainties

We assume that there is an error about the aircraft's future location because of ground speed prediction uncertainties⁴. The uncertainties on climbing and descending rates are even more important. As the conflict detection must be robust regarding these and

⁴Uncertainties on ground track will not be considered, as they do not increase with time and will be usually included in the separation standard




Figure 1: Modeling of speed uncertainties.

many other uncertainties, an aircraft is represented by a point at the initial time. The point becomes a line segment in the uncertainty direction (figure 1). When changing direction because of a maneuver order or a beacon ($t = 4$), the segment becomes a parallelogram that increases in the speed direction. When changing a second time direction ($t = 7$), the parallelogram becomes an hexagon that increases in the new speed direction, and so on. To check a conflict at time t , the distance between the two polygons modeling the aircraft positions is compared to the separation standard.

In the vertical plane, a cylindrical modeling is used (figure 1). Each aircraft has a maximal altitude and a minimal altitude. To check if two aircraft are in conflict, the minimal altitude of the higher aircraft is compared to the maximal altitude of the lower aircraft.

This model guarantees that if aircraft can handle such uncertainties than they will always fly in their boxes. The uncertainty rate can be different for every aircraft.

3.2 Centralized problem solver

The centralized problem solver developed by Durand [Dur96] can handle conflicts involving more than 2 aircraft. These conflicts are called clusters as they are obtained by merging 1-to-1 conflicts involving common aircraft during the next T_w minutes (T_w is called the time window, for the applications, $T_w = 12$ mn).

3.2.1 Specifications

The main idea guiding the design of the solver, is to be as close as possible to the current ATC system:

- Conflict free trajectories must respect both aircraft and pilot performances. Considering the evolution of ATC toward automation [DAM93], trajectories must remain simple for controllers to describe as well as for pilots to understand and follow.
- Maneuver orders must be given with an advance notice to the pilot. When a maneuver has begun, it must not be called into question.

The function to optimize is described in part 3.2.6.

3.2.2 Modeling

The current aircraft positions and flight plans are sent by CATS to a process DC (Detection Clustering) that builds trajectories forecast for T_w minutes, does conflict detection by pairs and transforms 1-to-1 conflicts in n -aircraft conflict called clusters. Then, the problem solver solves in parallel each cluster, as aircraft in each cluster are independent from aircraft in the other clusters. The problem solver sends to DC new orders and DC builds new trajectories forecast based on these orders. Once again DC runs a conflict detection process to check that modified trajectories for aircraft do not interfere with aircraft in another cluster, or with new aircraft. If no interference is found, new flight orders are sent to CATS. If there are interferences, interfering clusters are joined and the problem solver is used again on that (these) cluster(s). The process is iterated until no interference between clusters remains, or no new aircraft is concerned by modified trajectories. The new orders are sent back to the traffic simulator.

The above process is iterated and all trajectories are optimized each δ minutes. However, during the computation time, aircraft are flying and must know if they must change their route or not. δ should be large enough to compute a solution, send it to the pilot and let him time enough to begin the maneuver. Consequently, for each aircraft, at the beginning of the current optimization, trajectories are determined by the previous run of the problem solver and cannot be changed for the next δ mn (3 mn in the applications).

3.2.3 Clustering

After pair detection, DC process does a clustering which is a transitive closing on all pairs. Each equiv-




Figure 2: Horizontal maneuver modeling.

valence class for the relation “is in conflict with”, is a cluster.

The process will always converge: in the worst case, the problem solver will have to solve a very large cluster including all aircraft present in the next T_w minutes. However, this technique is usually efficient as a very large number of clusters can be solved very quickly in parallel.

3.2.4 Complexity of the problem

The complexity of the problem is exposed by Medioni, Durand and Alliot in [Dur96]. Let’s consider a conflict between two aircraft. We can easily prove that the minimized function is convex, but the set of conflict free trajectories is not. It is not even connected. If trajectories don’t loop, the set of conflict free trajectories has two connected components. For a conflict involving n aircraft, there may be $2^{\frac{n(n-1)}{2}}$ connected components in the free trajectory space which strongly suggests that any method which requires exploring every connected component is NP⁵. It is important to note that this complexity is independent of the modeling chosen (see [Dur96]).

3.2.5 Maneuver modeling

In the horizontal plane, aircraft are given turning point maneuvers (figure 2). The turning point angles are 10, 20 or 30 degrees.

A maneuver is determined by the maneuver starting time t_0 , the turning point time t_1 and the deviation angle s .

In the vertical plane, aircraft trajectories are divided in 4 periods (figure 3). During the climbing period, aircraft can be leveled at a lower than requested flight level during a moment to resolve a conflict. When aircraft have reached their desired flight level, they may be moved to the nearest lower level to resolve a conflict. When aircraft are about 50 nautic miles from beginning their descent to destination, they

⁵A Non deterministic Polynomial (NP) problem belongs to a class of problem for which there are no polynomial-time algorithm known to solve the problem.




Figure 3: Vertical maneuver modeling.

may be moved to a lower level to resolve a conflict. During the descent no vertical maneuver is possible.

No maneuver is simultaneously done in the horizontal and vertical plane. This model has the great advantage of reducing the size of the problem. In order to solve conflict due to aircraft taking off or entering the airspace simultaneously at the same point, a variable of delay t_d is introduced.

For a conflict involving n aircraft, the dimension of the search space is $4n$. This will allow us to solve very difficult conflicts with many aircraft without investigating a large solution space.

3.2.6 The function to optimize

One of the principal algorithm design challenges is to define a suitable function to optimize. A multiple-criteria function is required that simultaneously attempts to:

- minimize the delay due to deviations imposed on aircraft.
- minimize the total number of resolution maneuvers required and the total number of aircraft that will be moved⁶.
- minimize the maneuver duration so that aircraft are freed as soon as possible for maneuvers that may be necessary subsequently.
- enforce all separation constraints between aircraft.

3.2.7 Optimization tool

Use of local methods, such as gradient for example, is useless here, because these methods rely on the arbitrary choice of a starting point. Each connected component may contain one or several local optima, and

⁶Thus, instead of sharing the global delay on all the aircraft, some aircraft will support a part of the delay and others will not.

we can easily understand that the choice of the starting point in one of these components cannot lead by a local method to an optimum in another component. We can thus expect only a local optimum.

The problem solver uses Genetic Algorithms (GAs) to find the set of variables (t_d, t_0, t_1, s) that optimizes the above criteria. GAs are global stochastic optimization techniques that mimic natural evolution. They were initially developed by John Holland [Hol75] in the sixties. The subject of this paper is not GAs and the interested reader should read the appropriate literature on the subject [Gol89].

Genetic algorithms are a very powerful tool, because they do not require much information and are able to find many different optima that can be presented to a human operator.

Moreover, we know much about the function to optimize and this information can be used to create adapted crossover [DAN96] and mutation operators, an other advantage of GAs over other optimization techniques.

3.3 On board conflict solver

Reactive resolution methods similar to those developed by Karim Zeghal ([Zeg93, Zeg94]) are used in this part. Their main advantage is simplicity of development and use, since no extrapolation of trajectories is necessary. 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). Some additional information, in particular planned changes in course or vertical speed, can be helpful.

Limitations have been put on aircraft maneuvers. No speed change is allowed and only lateral maneuvers are given. To improve conflict resolution between aircraft following quasi-parallel tracks (pass-over or convergence with a small angle), a repulsive component has been added to the force vector. The repartition between the two components depends on the conflict geometry. The resulting force is purely tangential if both aircraft are on a colliding course and purely repulsive if they fly on parallel tracks.

The intensity of the avoidance force is proportional to the inverse of the time remaining until normal separation will be lost. 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, with a limit of $3^\circ/s$ on the turning rate.

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 1⁷.

Two different implementations 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 is not 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) and the uncertainty approximation described in part 3.1 are used to build clusters. The computation of reactive forces is then only applied to aircraft of the same clusters. The time of computation is approximately divided by 6.

4 RESULTS

The results presented below are simulations using different modules regarding the centralized problem solver and the on board conflict solver presented above, and the impact of reduced vertical separation minima (RVSM) on conflict workload, built on top of the CATS simulator. These simulation are not exhaustive and are examples of the possibilities offered by CATS.

4.1 Centralized problem solver

A complete experiment done with unregulated flight plans of the 21th of June 1996 is described here. It involves 6388 aircraft over France. Uncertainties on climbing rate and ground speed are respectively set to 20% and 5%, and standard separations are set to 6 nm and 1000 feet. The experiment is run under the Direct Route hypothesis (aircraft are allowed to go directly to their destination). We only detect and solve conflicts above 10000 feet, as we are only interested in En Route conflicts. Aircraft entering Paris TMA control

⁷These cases are known to be the most difficult to solve [DAN94]

area are sequenced on the TMA entry points, but no control is done inside the TMA. The simulation was performed in real time on 5 pentium 200.

When running this one day test with a very basic conflict detection algorithm (only actual conflicts are detected, with no uncertainty on speed) and with no conflict resolution, 1649 conflicts are detected.

When running the complete simulation with detection and resolution, the DC process detects 5064 1-to-1 different conflicts (This means that a detected conflict has $\frac{1649}{5064} = 32.6\%$ chance to really occur. The problem solver resolves 9130 clusters of different sizes (from 2 to 32 aircraft). There are 4155 clusters including different sets of aircraft. There is no unsolved cluster and consequently no conflict remains.

Only 1687 aircraft are given 2200 maneuvers which represents 1.3 maneuver per aircraft. The mean duration of a maneuver is $2mn23s$.

1337 aircraft are delayed before taking off or entering the French airspace. For these aircraft, the mean take off or entering delay is $2mn56s$ (maximum $6mn$). The global mean take off or entering delay is $37s$.

The mean maneuver duration expectation per aircraft is $49s$ which represents 1.79% of the flight duration.

The mean flight duration is $45mn54s$ before resolution and $45mn58s$ after resolution. The mean delay caused by maneuvers is $4s$. Only 934 aircraft are delayed because of maneuvers (most of the aircraft moved in the vertical plane are not delayed). The maximum delay is $4mn$ and the mean delay (for aircraft delayed) is $29s$.

This simulation gives a good idea of the ability of the problem solver to solve en route conflicts with actual traffic. More details can be found in [DA97].

4.2 On board conflict solver

When using the on board conflict solver on similar datas, 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). The maneuvering time is approximately 10%.




Figure 4: Number of separation losses vs density




Figure 5: Percentage of unsolved conflicts vs density

4.2.1 Results with knowledge of future trajectories

Figure 4 shows the evolution of the number of separation losses versus traffic density, both without (almost straight line⁸) and with 3 other conflict resolution. Resolution is performed with pre-detection of conflicts, with different conditions for each curve : respectively every 5 mn with twice the normal separation, every minute with twice the normal separation, and every 5 mn with one normal separation. Figure 5 shows the percentage of unsolved conflicts relative to the number of separation losses observed without resolution.

⁸The 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.

Influence of traffic density : When density is increased, the curve clearly shows the saturation of the resolution method. However, it appears later than with the purely reactive method (no prediction of conflicts). The percentage of unsolved conflicts is pretty good up to density 2.5 (17% unsolved, compared to 25% with the purely reactive method), but it increases dramatically when density exceeds 3 (70% unsolved for density 4, compared to 50% without conflict prediction). There is also a rapid increase of the number of unpredicted conflicts.

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. 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. More details can be found in [ABNL97].

4.3 Assessment of RVSM procedures

This section describes preliminary results for assessing the introduction of reduced vertical separation minima procedures ([Int92]) in continental airspace above flight level FL290. We use the CATS simulation based planner to study three main RVSM airspace organizations, with the primary objective of estimating potential benefits on conflict work load, and a secondary objective of comparing these alternative organizations according to the density of potential conflicts.

4.3.1 Scenario description

With current ICAO vertical separation minima, aircraft must be separated by at least 2,000 feet above FL195. Furthermore, aircraft cruise at levels that are compatible with the orientation of the route network: in specific airspace areas, aircraft flying north must fly at odd levels, and aircraft flying south must fly at even levels. Under ICAO rules, FL290 is odd, and all levels above alternate: FL310 is even, FL330 is odd, FL350 is even, etc. This type of flight level orientation scheme (FLOS), which may correspond to either east-west orientation, or north-south orientation, is adapted by each air traffic control center (ACC).

Four organizations are simulated, which differ by the vertical separation minima used above FL290, and the flight level orientation schemes selected:

The geographical area of the simulation covers the the 6 French ACCs. The simulated traffic sample is extracted from French traffic archives, and comprises 353,741 filled and actual flight plans of September 1996. Airspace operational constraints were not taken into account due to their unavailability at the time of the study. RVSM scenarios use 1,000 feet of vertical separations minima, as follows:

1. ICAO - 1,000 feet below FL195, 2,000 feet above FL195 for vertical separations;
2. RVSM - single alternate: all levels are odd or even based on the second digit of the level: i.e FL330 is odd, FL320 is even.
3. RVSM - double alternate: ICAO levels keep their orientations, new intermediate levels have the same orientation of the ICAO level immediately higher.
4. RVSM - quadruple alternate: one out of four levels is used according to the north-east, north-west, south-west or south-east orientation of the route.

4.3.2 Results

To evaluate the various options of the scenario, more than 5 millions flights are evaluated: 30 days of traffic * 2 traffic demand types (filled and actual flight plans) * 6,000 flight plans per day * 14 options. These simulations were performed overnight in 10 hours of computing time on the workstation cluster at CENA, with 30 machines. We analyze below the statistics output from the simulations both at the sector level with sectors G2 and G3 of Aix-en-Provence ACC, and at the overall airspace system level.

Sector	Ref	Single	Double	Quad.
G2	4.048	2.212	2.356	1.926
G3	4.154	1.962	2.048	2.014

Figure 6: Average number of conflicts per hour computed for peak hours of September 1996

Traffic	Ref	Single	Double	Quad.
Filled	0.491	0.284	0.278	0.268
Actual	0.444	0.247	0.240	0.233

Figure 7: Conflict risk levels for proposed RVSM organizations and impact of traffic flow regulation

4.3.3 Traffic work load

Sectors G2 and G3 are defined respectively between FL275-FL320, and FL320-FL900. The analysis of sector throughput and work load show a transfer of traffic (20%) between G3 and G2 with the RVSM single alternate. In other words, with RVSM, the traffic is no longer equitably partitioned between the two sectors. They might have to be re-designed vertically.

4.3.4 Conflict work load

Conflict work load⁹ is defined as the average number of separation minima infringements per sector. For sectors G2 and G3, it is described in figure 6, for peak hours of September 1996. With RVSM, conflict rates are reduced in average by 50%. Despite traffic flow increases in G2, the conflict rate is reduced by 40%. Other conflict related indicators are in similar plight.

At the system level, we estimate (figure 7) the impact of RVSM organizations on conflict density with a conflict risk indicator: the number of conflict per hour of flight, for all conflicts above FL100, and for all the simulated traffic sample (353,741 flight plans).

The overall impact of RVSM on conflict work load is clearly significant, much higher than the impact of traffic flow regulation. Furthermore, there is no clear benefits between the proposed RVSM options (single or double alternate). Bias due to traffic sample estimation procedures need to be further analyzed. At this point, operational considerations such as transition steps prevail over the choice of either RVSM options.

⁹We assume that safety, defined as the probability of collision per hour of flight, is not affected by the introduction of RVSM. Safety issues are addressed with specific analytical models ([Int92]). Indeed, it is inherently difficult to estimate by simulation the distribution of mid-air collisions, due to their high improbability. We do measure however the risk of conflict as a controller work load indicator.

Traffic	Risk	Variation
	0.491	
+20%	0.571	+16%
+50%	0.698	+42%
+100%	0.901	+83%
+200%	1.295	+163%

Figure 8: Impact of density factor on risk level

4.3.5 Traffic density

We estimated the impact of traffic density on the overall risk level by taking the same traffic sample and compressing the flights' entry time in the simulation (divided by two to double the density). Standard air route network and vertical separation minima was used (figure 8).

5 CONCLUSION

The examples presented above give an idea of the testing possibilities of CATS. This paper is not exhaustive, and many other studies are in progress at different levels of ATC and ATM. This paper tried to focus on different scenarios of conflict resolution for different terms. The powerful of CATS rely on the language used (CAML) and on its small size (less than 3000 lines of code).

References

- [ABNL97] J.M. Alliot, J.F. Bosc, N.Durand, and L.Maugis. An Experimental Study of ATM Capacity. In *1 ST U.S.A/EUROPE ATM R & D Seminar*, Mai 1997.
- [DA97] N. Durand and J.M Alliot. Optimal Resolution of En Route Conflicts. In *1 ST U.S.A/EUROPE ATM R & D Seminar*, Mai 1997.
- [DAM93] Patrick Dujardin, J. M. Alliot, and Paul-Henri Mourlon. Different paths to automation. In *IFAC'93*, 1993.
- [DAN94] Nicolas Durand, J. M. Alliot, and Joseph Noailles. Algorithmes genetiques : un croisement pour les problemes partiellement separables. In *Proceedings of the Journees Evolution Artificielle Francophones*. EAF, 1994.
- [DAN96] Nicolas Durand, J. M. Alliot, and Joseph Noailles. Automatic aircraft conflict resolution using genetic algorithms. In *Proceedings of the Symposium on Applied Computing, Philadelphia*. ACM, 1996.
- [Dur96] Nicolas Durand. *Optimisation de Trajectoires pour la Résolution de Conflits en Route*. PhD thesis, ENSEEIHT, Institut National Polytechnique de Toulouse, 1996.
- [Gol89] David Goldberg. *Genetic Algorithms*. Addison Wesley, 1989. ISBN: 0-201-15767-5.
- [Hol75] J.H Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan press, 1975.
- [Int92] International Civil Aviation Organisation. Manual on implementation of a 300M (1,000) feet vertical minimum between FL290 and FL410 inclusive. Report 9574-an-934, 1992.
- [NAS96] NASA. Advanced Air Transportation Technologies. <http://www.nas.nasa.gov/AATT>, November 1996.
- [Sof96a] Sofreavia. Model Use of Fast Time Simulation. Final Technical Report MUFTIS WPR5.1, Commission of the European Communities. Directorate General VII, Brussels, February 1996.
- [Sof96b] Sofreavia. New Approaches for Air Traffic Flow Management. NOAA WP3, Speed Control, Multi-sector planning and their interfaces, Commission of the European Communities. Directorate General VII, Brussels, 1996.
- [Zeg93] Karim Zeghal. A Reactive Approach for Distributed Air Traffic Control. In *International Conference on Artificial Intelligence & Expert Systems*, Mai 1993.
- [Zeg94] Karim Zeghal. *Vers une théorie de la coordination d'actions - Application à la navigation aérienne*. PhD thesis, Université Paris VI, Décembre 1994.