Using A* Graph Traversal to Model Conflict Resolution in Air Traffic Control

Stefan Lehmann (Stefan.Lehmann@nicta.com.au)
National ICT Australia, Queensland Research Laboratory, Level 1, McElwain Building (24A), The University of Queensland
Brisbane QLD 4072 Australia

Scott Bolland (scottb@itee.uq.edu.au)
School of ITEE, General Purpose South Building (78), The University of Queensland
Brisbane QLD 4072 Australia

Roger Remington (r.remington@psy.uq.edu.au)
School of Psychology and NICTA QRL, McElwain Building (24A), The University of Queensland
Brisbane QLD 4072 Australia

Michael S. Humphreys (mh@psy.uq.edu.au)
School of Psychology and NICTA QRL, McElwain Building (24A), The University of Queensland
Brisbane QLD 4072 Australia

Andrew Neal (Andrew@psy.uq.edu.au)
School of Psychology and NICTA QRL, McElwain Building (24A), The University of Queensland
Brisbane QLD 4072 Australia

Abstract

The efficient detection and resolution of conflicts represent the key tasks of Air Traffic Controllers in enroute environments. The complexity of these tasks imposes significant challenges on the design of cognitive models that are capable of adequately simulating them. Yet, the availability of such models is crucial for a number of applications, including the evaluation of current and future Air Traffic Control concepts. In this paper, we will propose a novel modeling approach which adopts the principles of the A* graph search scheme from Artificial Intelligence to represent the cognitive decision making process of the human operator. Results of an initial version of this model will be presented, showing that the proposed approach has promise.

Keywords: Cognitive Modeling; Cognitive Systems Engineering; Artifical Intelligence; Decision Making; Air Traffic Control.

Introduction

In most western economies, the volume of air traffic is currently growing at a rate of 4 to 6 percent per annum. According to its 2006 annual report, the US Federal Aviation Administration (FAA) acknowledges that air traffic controllers will not be able to handle traffic at 25 percent above today's level, and that traffic may increase this much by 2016 (ICAO, 2004). In response to this problem, the United States Federal Aviation Administration and Eurocontrol are currently pursuing programs to greatly increase airspace capacity (FAA, 2010; Eurocontrol, 2008), without raising either the workload or number of air traffic controllers.

Cognitive modeling could provide an important vehicle for the evaluation of new operational concepts if it is possible to simulate performance on challenging air traffic control operations. For example, models making reasonable estimates of sector workload could inform evaluations of safety and staffing. One of the more cognitively complex tasks of controllers is the detection and resolution of conflicts (Lehmann, Bolland, Remington, Humphreys, Fothergill, Hasenbosch, & Neal, 2010). The *n*-aircraft conflict resolution problem is highly combinatorial and cannot be optimally solved using classical mathematical optimization techniques. This inherent complexity imposes significant challenges on the design of corresponding models.

This paper will propose a new method that simplifies the task of modeling expert decision making in Air Traffic Control (ATC) environments by relying on domain-specific simple heuristics that humans deploy to produce accurate decisions (Todd & Gigerenzer, 2007). The conflict resolution mechanism adopts the principles of the A* search algorithm (Felner, Stern, Ben-Yair, Kraus, & Netanyahu, 2004; Lee, Osman, & Sabudin, 2009; Leigh, Louis, & Miles, 2007). The resulting scheme implements a search through a space of conflict solutions. System states are evaluated using optimization criteria encapsulating the controller's goals. Each optimization criterion is associated with a number of individual cost functions that penalize deviations of the system states from the goal states. The focus on psychologically plausible strategies, rather than representative psychological processing mechanisms, was in part a response to the complexity of decision making in ATC and the large number of unobservable factors that would need to be incorporated (*e.g.*, memories for previous or typical solutions). Moreover, the strategies we use were elicited from highly experienced controllers and thus encapsulate experts' insights and knowledge. Our working hypothesis is that the use of psychologically plausible solution heuristics and optimization criteria in conjunction with the constraints imposed by the environment will produce human like behavior.

We first describe the conflict detection mechanism, then detail the manner in which the model selects solutions using the optimization criteria to find a path in the search tree. Finally, we present empirical tests of an initial implementation of the model showing good but not perfect fits to data from human controllers.

Conflict Detection Scheme

The current implementation of the conflict detection scheme is based on the model proposed in Loft et al. (2009). It detects pairs of conflicting aircraft in a hierarchical fashion. Its decomposition into three operational stages allows for a run-time efficient implementation. Potential conflicts are verified by extrapolating the flight paths of all aircraft that are present in the given scenario, and by subsequently identifying violations of separation standards between the flight paths. Positional aircraft uncertainty is accounted for in this process. The three stages proceed as follows:

<u>Stage 1: Coarse check of vertical separation</u>

A coarse check is performed to verify the vertical separation between aircraft. This stage checks if the vertical corridors of any two aircraft of interest are separated by more than 1000 ft, where the vertical corridors are defined by the aircraft's target altitude and cleared altitude respectively.

<u>Stage 2: Lateral separation check</u>

If the first stage (coarse check) reveals the existence of a possible vertical conflict between two aircraft, the model deploys the so-called *Trajectory Modeller* to check for a lateral conflict. At any given time *t,* the Trajectory Modeller extrapolates the flight paths up to time $t + 10\ min$ in discrete $\Delta T = 5\ sec$ steps. The aircraft positions at each time step are subject to positional uncertainty, where the uncertainty increases successively over time based on a step function. More specifically, the extrapolated aircraft position at a discrete time step $t_k = k\Delta T,\ k=0, 1, 2, 3, \ldots$ is associated with a discrete uncertainty interval $[a_k\Delta T,\ b_k\Delta T]$, where the coefficients a_k and b_k associated with the lower and upper limits of the interval are:

$$a_k = \mathrm{trunc}\,(0.98 \cdot k) \qquad \text{Equation 1}$$

$$b_k = \mathrm{trunc}\,(1.02 \cdot [k+1]) \qquad \text{Equation 2}$$

<u>Stage 3: Final vertical separation check</u>

If the second stage (lateral separation check) verifies a potential lateral conflict between two aircraft of interest, a third stage will be deployed to check for vertical conflicts. For this purpose, the respective flight paths are vertically extrapolated based on the maximum and minimum climb or descent rates of the aircraft. Response times of the aircraft are currently not considered. That is, the aircraft are assumed to instantaneously initiate the actions associated with the controller's interventions.

Decision Making Model

The proposed decision making model adopts the principles of the A* graph search algorithm (Felner, Stern, Ben-Yair, Kraus, & Netanyahu, 2004; Lee, Osman, & Sabudin, 2009; Leigh, Louis, & Miles, 2007). This algorithm relies on a state-space search engine to evaluate the decision alternatives in a hierarchical fashion. Hierarchical search has been shown to produce good modeling solutions to complex aeronautical problems in the past (Nason & Laird, 2005; Rosbe, Chong & Kieras, 2001).

A* finds the minimum cost path in a decision tree through a partial search in the solution space. The avoidance of an exhaustive search presents a significant advantage for its application in the ATC domain, where the decision making process poses a complex problem that typically leads to an extensive search tree in general traffic scenarios. That is, the topology of the search structure does not need to be known *a-priori.* In our model, the search space consists of *solution types*, each representing an action that could be taken to resolve the conflict. The solution types are based on simple heuristics that have been obtained from experts (using interviews and controlled experiments), and from data mining (using radar track data).

Solution Types

The current implementation of the conflict resolution model provides a set of three different solution types which may be applied to the aircraft involved in potential conflicts. Before a solution can be considered for exploration, one or more conditions of applicability must be satisfied. Each solution has a particular weight. A smaller weight corresponds to a more favourable solution. The effective weight of a solution is the sum of a base weight and a penalty value. The purpose of the penalty values is to impede the selection of solutions that would severely disturb an aircraft's intended flight path. The individual solution types and their weights are:

<u>*A. Assign closest level below or above conflict zone*</u>

The principle of this solution type is to ensure sufficient vertical separation by assigning one of the two aircraft of the conflict pair a safe altitude either beneath (*low solution*) or above (*high solution*) the other aircraft whilst they are in the region of the airspace where a loss of lateral separation is possible. More specifically, assuming two conflicting aircraft *A* and *B*, the low solution is applicable if A is not already descending through the low solution. Alternatively, the high solution is applicable if *A* is not already climbing through the high solution. This avoids direct transitions from a descent into a climb or from a climb into descent respectively.

Figure 1 illustrates an example where both aircraft *A* and *B* are on climb from Flight Level *(FL) 110* to *FL150* and from *FL120* to *FL160* respectively.

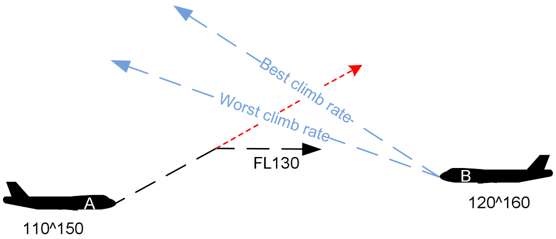

Figure 1: Assign closest level below

The climb of aircraft *A* is halted below aircraft *B* by assigning *FL130* to aircraft *A*.

The base weight of this solution type is (-0.5). Penalty values in the amount of +0.1 are additionally applied if the solution applied to A falls outside the transitional altitude band defined by A's current and cleared altitudes.

B. Assign separated levels

The second solution type involves modifying the levels of both aircraft, assuming a pair of conflicting aircraft where one aircraft is climbing and the other aircraft is descending. Figure 2 illustrates the basic concept of this solution, once again using a conflict pair of aircraft *A* and *B*. In this example, aircraft *A* is climbing from *FL110* to *FL150*, while aircraft *B* is descending from *FL150* to *FL110*.

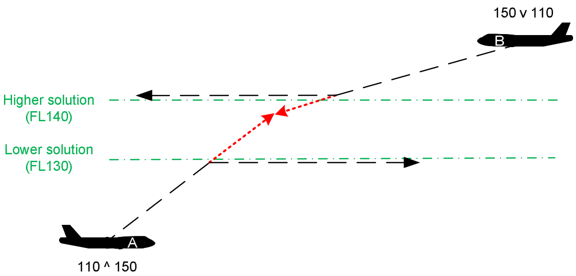

Figure 2: Assign separated levels

In this case, the applicable solution is to interrupt both the climb of aircraft *A* and the descent of aircraft *B* by assigning *FL130* to aircraft *A* and *FL140* to aircraft *B*, thereby ensuring that sufficient vertical separation between the aircraft is maintained.

The base weight of this solution type is (-0.5). Penalty values in the amount of +0.1 are added to the weight for any reverse climb or reverse descent intervention.

C. Vector behind solution

The *vector behind* solution proceeds as follows: A circle with a radius of *6nm* (nautical miles) is placed around aircraft *B* at its current position. Aircraft *A* is pointed behind aircraft *B* by vectoring it to the heading that establishes a tangent to this circle, thereby ensuring sufficient lateral separation between the two aircraft.

This solution is generally applicable to all conflicting aircraft. Its base weight is *(-0.5)*. There are no additional penalties.

Adaptation of A* to the ATC decision making task

The search space of the A* algorithm can be graphically represented by a decision tree. An example graph is shown in Figure 3. Each node in the decision tree represents a system state that, with the exception of the start node (*S*), results from the path of previous actions leading to it. The edges between the nodes represent the path of actions. Each edge has a value (shown as an integer in Figure 3) representing the cost incurred by traversing that edge. It is worthwhile to note that apart from the goal node (*G*), each node has at least one decision alternative associated with it, leading to a so-called *child node*.

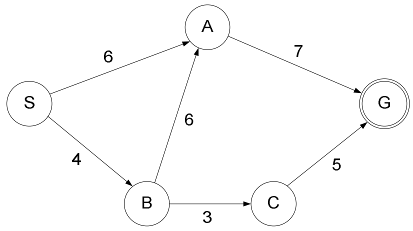

Figure 3: A* example graph

The decision making process is effectively driven by the cost function *f(x)*. That is, A* ranks each path currently under consideration based on *f(x)* to find the path with the lowest traversal cost. *f(x)* is decomposed into a so-called path-cost function *g(x)* reflecting the cost from the starting node to the node of interest, and a "heuristic estimate" *h(x)* of the distance to the goal node.

$$f(x) = g(x) + h(x), \qquad \text{Equation 3}$$

where *x* denotes some partial path. In other words, *f(x)* represents the estimated final cost of the path leading to the goal and including *x*. Under the right conditions, A* guarantees to find the path with the lowest traversal cost (Leigh, Louis, & Miles, 2007). The performance of A* relies heavily upon the heuristic estimate *h(x)*. A necessary condition for A* to find the shortest path is that the heuristic must underestimate the remaining distance.

One of the key aims in adopting the A* search scheme to the ATC conflict resolution task consists in achieving a model behavior that is closely aligned to the behavior of human controllers. For this purpose, the concept of optimization criteria was introduced. Each optimization criterion C_n encapsulates the n^{th} goal of the controller. Table 1 shows three examples for possible optimization criteria:

Table 1: Three exemplary optimization criteria

n	Optimization criterion C_n
1	Minimization of total number of aircraft interventions
2	Minimization of disruption to aircraft flow
3	Minimization of the controller's workload

Each optimization criterion C_n is associated with a set of descriptive attributes, A_{nk}. These attributes are represented by corresponding cost functions

$$f_{nk} = g_{nk} + h_{nk}. \qquad \text{Equation 4}$$

Summing up all the cost contributions across the individual attributes yields the final cost function for the individual criterion C_n:

$$f_n = \sum_k (g_{nk} + h_{nk}) \qquad \text{Equation 5}$$

Our initial version mainly aims at the implementation of optimization criterion C_1 from Table 1. That is, it tries to resolve all conflicts given in the scenario with the fewest aircraft interventions. However, the second criterion listed in Table 1, C_2, was additionally integrated into the model, to account for the attempts of controllers to minimize unfavorable flight maneuvers. Table 2 shows the individual attributes for C_1 and C_2.

Table 2: Attributes of the optimization criteria as per the current model implementation

C_n	k	Attribute A_{nk}
C_1	1	Preference of graph nodes of lower depth level
C_1	2	Preference of nodes showing fewer remaining conflicts
C_1	3	Number of conflicts of the aircraft subject to intervention
C_1	4	Number of occurrences of the solution of interest
C_2	1	Obstruction of unfavorable flight maneuvers

As Table 2 shows, C_1 is represented by four attributes and C_2 by one attribute respectively. The aim of the attribute A_{11} in Table 2 is to prioritize the selection of solutions that belong to graph nodes at low depth levels. The depth level of a node is determined by the number of subsequent nodes lying in the decision path, that is, by the number of actions leading to it. Therefore, the node depth defining the corresponding cost function $g_{11}(x)$ represents the number of interventions that have already occurred in the path of interest *x*, and that have consequently already imposed a penalty on the achievement of optimization criterion C_1.

Generally, the number of remaining conflicts in a given node establishes a good indicator for the expected number of remaining interventions. Consequently, this measure was taken to define the cost component $h_{12}(x)$ for the corresponding attribute A_{12} in Table 2. The metric was encapsulated in the heuristic component *h* of the cost function *f* as it represents a *predictive* cost estimate. The number of conflicts that the aircraft the solution acts upon is involved in represents an additional indicator for the efficiency of the solution with respect to achieving criterion C_1 in the remaining path to the goal. The number of remaining conflicts therefore forms the cost component $h_{12}(x)$ corresponding to attribute A_{12}. The underlying idea is that in comparison to solutions that are applied to aircraft that are involved in a single conflict only, solutions applied to an aircraft involved in multiple conflicts have a greater than zero probability of resolving multiple conflicts this aircraft is subject to in one go. This likelihood of efficiently minimizing the intervention count is further increased if in addition to acting on aircraft involved in multiple conflicts, the particular solution is suggested multiple times by the solution logics for resolving different conflicts. The number of total occurrences of the solution under consideration was therefore taken to define cost component $h_{13}(x)$ corresponding to attribute A_{13}.

The cost function for attribute A_{21} is simply the sum of the base weights of the solutions and the respective penalties as described in the subsection entitled *Solution Types*. While the base weights for the individual solutions are identical for all solution types in the current implementation, the additional penalties depend on the situational context. Their purpose is to prevent the selection of solutions yielding unfavorable aircraft maneuvers, such as reverse climbs and reverse descents.

Based on this set of individual cost components, the cost functions $f_0(x)$ and $f_1(x)$ are computed using Equation 5. The final cost function $f(x)$ is then just formed by adding $f_0(x)$, $f_1(x)$, and a Gaussian noise term that accounts for the probabilistic nature of the human decision maker. This noise term is characterized by a relatively small variance and therefore predominantly influences the selection of solutions belonging to the same search tree level. Impacts of this noise on solutions belonging to different tree levels are very unlikely. All parameters required for the formulation of the cost functions, including the variance of the noise, were empirically chosen in the current implementation. The effective cost $f(x)$ establishes the basis for the decision making process in the ATC search tree. This process will be discussed in the following subsection.

ATC Search Tree

An example of the resulting ATC search tree is depicted in Figure 4.

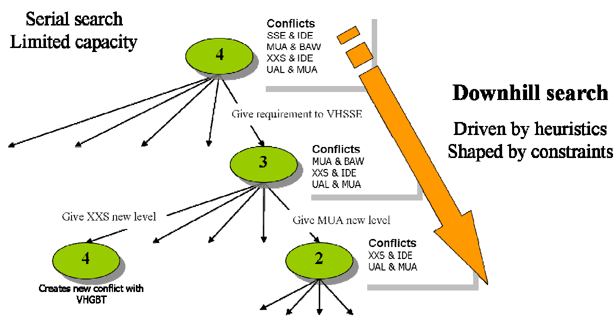

Figure 4: ATC search tree

In this example, the conflict detection model initially detects four potential conflicts between aircraft pairs in the scenario, as depicted in the root node within Figure 4. A set of potential solutions is then constructed for each of the potential conflicts present in this node. The entire set of potential solutions is then evaluated by assigning individual cost values $f_{i,j}$ to the solutions, where i ($i = 0$ for root node) and j denote the indices of the current node and the solution under consideration respectively. The solution having the smallest cost value will finally be selected and applied, creating a new child node with an associated set of conflicts. In the example in Figure 4, the solution selected in the root node resolves one of the four problems, leaving the respective child node with three remaining problem pairs. The process applied to the root node is then repeated for the child node in a recursive fashion. Figure 4 also demonstrates that solutions selected via *a-priori* evaluation may be deemed to be inefficient via *a-posteriori* evaluation. For example, the solution entitled '*Give XXS new level*' creates a new conflict, which leads to back-tracking behavior in the search process. That is, the subsequent search evaluation step may select a solution associated with the parent node, rather than propagating further down from the child node produced by the previous, inefficient solution. The overall optimization scheme effectively leads to a downhill search which is driven by the available set of solution types (heuristics) and shaped by the situational context (constraints).

Experiments

Aim and Methodology

To compare the model's behavior against the behavior of controllers, we simulated performance on a set of four different scenarios of varying complexity that were also presented to *14* En-Route, radar endorsed air traffic controllers from Brisbane Centre. Figure 5 shows the scenario with the highest complexity.

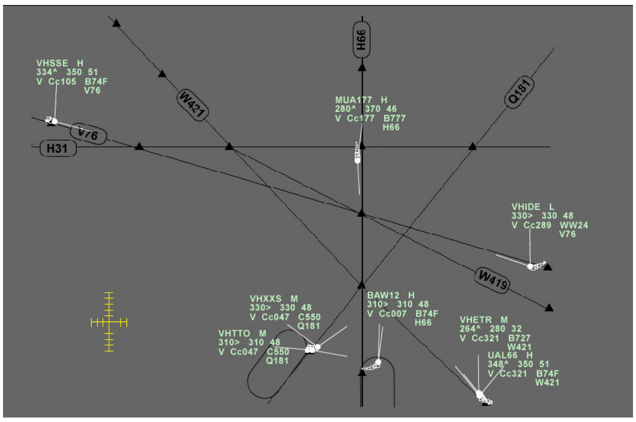

Figure 5: Scenario of highest complexity

The time participants had been endorsed as a controller ranged from *10* to *20* years. Controllers were asked to resolve the scenario by issuing restrictions to one or more of the aircraft. They were instructed to work through the scenario step by step, and to explain their actions in detail, including the evaluation of potential problems, and the processes of considering options and deciding on actions or priorities. The interviews were based on the critical decision method (Klein, Calderwood & MacGregor, 1989).

The simulation consisted of *100* runs of our decision making model for each scenario. Our interest centers on the degree to which the model used the same intervention rates and types as the human controllers. Table 3 shows the intervention types.

Table 3: Intervention types

Type	Description
H0	Intervention other than H1, H2,…, H8
H1	Vector aircraft to the left
H2	Vector aircraft to the right
H3	Issue climbing instruction
H4	Issue descent instruction
H5	Extend an existing climb
H6	Extend an existing descent
H7	Interrupt an existing climb
H8	Interrupt an existing descent

Results

The results for the scenario with the highest complexity are presented in Figures 6 and 7. Figure 6 shows the total average intervention rates for the individual aircraft for both controllers and model runs. Figure 7 shows the selection rates of the individual intervention types.

It can be seen from Figure 6 that there is a reasonable agreement between controllers and the model in selecting the aircraft that are subject to intervention. However, controllers appear to intervene with a wider range of aircraft than the model, at more variable intervention rates: Aircraft '*VHETR*' is excluded by the model in Figure 6.

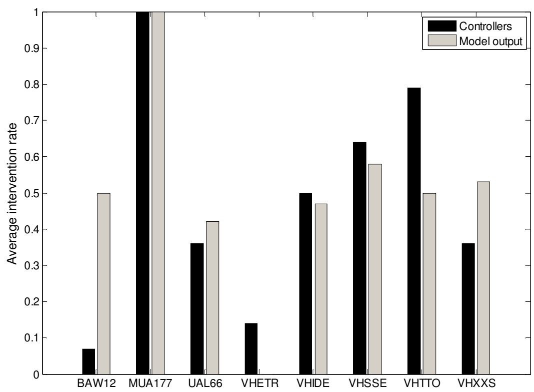

Figure 6: Total average intervention rates for the aircraft

Figure 7 demonstrates a reasonable agreement between controllers and the model in the selection of the intervention types. However, a reduced variability of the model can be observed: In contrast to controllers, the model essentially excludes the generation of intervention types *H0* (‘*Intervention other than H1, H2,…, H8*’) and *H5* (‘*Extend an existing climb*’).

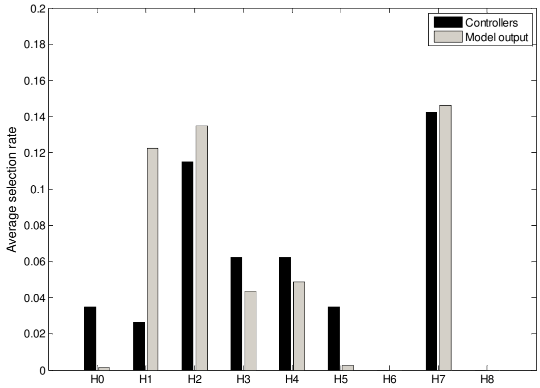

Figure 7: Average selection rates of the intervention types

Conclusions and Outlook

This paper describes a novel approach for modeling the Air Traffic Control (ATC) task using intelligent graph search. The A* algorithm was adopted to model human decision making under uncertainty and environmental constraints. This model relies on the definition of optimization criteria and associated attributes, where the attributes are represented by corresponding components of the overall cost function. The optimization criteria encapsulate properties of the situational context that influence the decision strategies of a human controller. They can consequently enable the model to alter its behavior accordingly. An initial implementation of this model is proposed that aims at minimizing the total aircraft intervention count under preservation of the realism of the generated solutions. Empirical tests demonstrate good but not perfect fits to data from human controllers. A reduced variability of the model over controllers was observed, in the selection of both the aircraft for intervention and the actual types of intervention. This variability might be induced by psychological processes that the model does not capture, such as human attention and perception.

The results suggest that the modeling concept has promise for its application to decision making in complex, dynamic task environments. We therefore plan to extend the approach in our future work by incorporating additional optimization criteria; by advancing the current decision making mechanisms; and by integrating adaptive behavior into the model.

Acknowledgments

This research project is supported by Airservices Australia. The authors acknowledge their contribution to this research in providing access to subject matter experts and facilities. The authors additionally thank Ms Selina Fothergill for providing access to the data that was generated as part of her PhD program in the School of Psychology, The University of Queensland.
NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

References

Eurocontrol (2008). *Introducing the next generation of air traffic control*. Available online: http://www.eurocontrol.int/muac/gallery/content/public/docs/Introducing%20the%20next%20generation%20of%20ATC.pdf [Last accessed: *19-April-2010]*

FAA (2010). *FAA’s NextGen Implementation Plan*. Available online: http://www.faa.gov/about/initiatives/nextgen/media/NGIP_3-2010.pdf [Last accessed: 19-April-2010]

ICAO (2004). *Outlook for Air Transport to the Year 2015*. (Circular 304). Montreal, Canada: ICAO.

Lee, W. P., Osma, M.A., & Maziani, S. (2009). Design of an Intelligent Route Planning System using an Enhanced A*-search Algorithm. *Third Asia International Conference on Modelling & Simulation* (pp. 40-44).

Lehmann, S., Bolland, S., Remington, R., Humphreys, M.S., Fothergill, S., Hasenbosch, S., & Neal, A. (2010). Evaluation of a Model of Expert Decision Making in Air Traffic Control. *Proceedings of the 9th Conference of the Australasian Society for Cognitive Science* (pp. 204-209).

Leigh, R., Louis, S.J., & Miles, C. (2007). Using a Genetic Algorithm to Explore A*-like Pathfinding Algorithms. *Proceedings of the 2007 IEEE Symposium on Computational Intelligence and Games (CIG 2007)* (pp. 72-79).

Loft, S., Bolland, S., Humphreys, M., & Neal, A. (2009). A Theory and Model of Conflict Detection in Air Traffic Control: Incorporating environmental constraints. *Journal of Experimental Psychology: Applied*, *15(2)*, 106-124.

Rosbe, J., Kieras, D., & Chong, R. (2001). *Modeling with perceptual and memory constraints: an EPIC-Soar model of a simplified enroute air traffic control task* (Tech. Rep. AFRL-HE-WP-TR-2002-0231). Ann Arbor, MI: Soar Technology Inc.

Todd, P. M., & Gigerenzer, G. (2007). Environments that make us smart: Ecological rationality. *Current Directions in Psychological Science*, *16(3)*, 167-171.