Computational Modeling of Air Traffic Control: Terminal Area Case Study

Todd J. Callantine
San Jose State University/
NASA Ames Research Center
Moffett Field, CA, USA
tcallantine@mail.arc.nasa.gov

Abstract – *This paper presents an air traffic control modeling case study in which agent performance on a simulated terminal-area air traffic control task was compared with human air traffic controller performance. The paper first provides background on the simulation with human air traffic controllers, and on prior air traffic control modeling research. It then presents an agent model that enables various air traffic control strategies to be evaluated. Results for four different control strategy models show that agent performance compares more favorably to human performance when clearance alternatives and resultant aircraft behaviors are more constrained. The paper discusses these results and describes proposed future enhancements. The research was conducted with support from the NASA Aviation Safety Program.*

Keywords: Air traffic control, agents, control strategy, coordination, flight management system.

1 Introduction

New air traffic management (ATM) concepts are needed to remedy the capacity limits, inefficiencies, and adverse environmental impacts of the current ATM system. The research and development process is challenging because of the complexity of the ATM system and the broad range of issues—including automation functionality and operator interaction, operational scenarios, and training—it must address. Computational agents that can control simulated air traffic are an attractive complement to traditional human-in-the-loop simulation methods, with significant roles to play in ATM concept development and safety/risk assessment.

Recently, modeling research has addressed air traffic control (ATC) from several perspectives. For example, researchers have used a computational model to study how controllers construct a 'picture' of the traffic situation [15]. Others have investigated control strategies [16] and conflict detection and resolution rules [14]. Models have been developed to assess control techniques [12], enable decision support [11], and produce predictive performance measures to support safety and benefits analyses [13].

Monte Carlo-style safety and efficiency studies are a particularly attractive application, because various traffic conditions can be simulated inexpensively in less time than with real-time human-in-the-loop simulations. Researchers have examined the impacts of new alerting systems [17], and of proposed changes to practitioner roles and responsibilities through simulations with computational agents [8, 10]. Agents can also help investigate error effects, scenario design, and interaction design for new ATM support tools [2, 7].

Different ATM environments and simulations alter the requirements for computational ATC agents, because each emphasizes and provides tools that support certain ATC capabilities differently. Thus, each affords an opportunity to examine the effectiveness of ATC agent models and improve them. This paper presents a case study of agents developed to control traffic in a terminal-area environment in which air traffic controllers are responsible for descending and slowing aircraft for the final approach to landing. Sequencing, spacing, and merging aircraft transitioning to the final approach requires some of the same capabilities as en route ATC, but aircraft behavior is more difficult to predict than when aircraft are in level flight and the airspace is more confined, which constrains how controllers can safely separate aircraft.

1.1 DAG-TM Terminal-Area Simulation

This case study specifically examines agent control of traffic in a terminal-area ATM environment investigated in the NASA Advanced Air Transportation Technologies project Distributed Air Ground Traffic Management (DAG-TM) element. A recent DAG-TM terminal-area simulation at NASA Ames Research Center evaluated the feasibility and potential benefits of using pilot and controller decision support tools for time-based airborne spacing [6]. The baseline condition resembled current-day terminal-area operations, except that all aircraft were equipped with Flight Management System (FMS) automation and Automatic Dependent Surveillance-Broadcast (ADS-B) for providing additional state information. Aircraft entered the terminal area on FMS routes that merged prior to the final approach fix. Eight traffic scenarios each began with a traffic flow that was

well coordinated for merging and spacing and ended with an uncoordinated flow. Two professional air traffic controllers staffed 'feeder' and 'final' control positions. The feeder controller was envisioned to perform initial spacing adjustments; the final controller would then merge the flows and 'lock in' the required spacing. A key metric in DAG-TM study was spacing accuracy at the final approach fix.

1.2 En route Air Traffic Controller Agents

The terminal-area controller agents presented here extend and improve agents developed earlier for en route air traffic control. In that research, agents used a Crew Activity Tracking System (CATS) activity model to represent the high-level structure of the air traffic control task [5], and controlled traffic simulated in real time. A set of static priorities governed how the agents selected aircraft to focus on; these priorities gave rise to a characteristic 'flow of control.' The agents maintained beliefs about the current task context and current traffic situation. Agents transformed their belief set by performing activities. A collection of heuristics determined the specific control strategies that agents applied to address separation and in-trail spacing requirements. Strategies that entailed planning were supported by a collection of plans and plan adaptation conditions. Agents issued a clearance, then some time later either issued a planned clearance (e.g., to return the aircraft to its flight plan route), or adapted the planned clearance based on an updated 'picture' of the traffic situation. The agents performed reasonably well in terms of their ability to keep aircraft safely separated. However, results suggested that performance could be improved if agents were capable of dynamically focusing attention to address the most critical control needs.

1.3 TCSim Agents

The agents presented here are also the most recent to be integrated with the Trajectory-Centered Simulator (TCSim), a fast-time simulation tool implemented in Java™ that uses models of aircraft FMS trajectories to simulate aircraft [4]. A variety of elementary air traffic control agents have been implemented in TCSim to examine the effectiveness of specific control strategies in terminal-area operations [3, 7]. Some of these agents simply tried to issue a particular type of clearance at one of several pre-specified 'control points.' As each aircraft in an arrival flow crossed a control point, the agent would evaluate the applicability of potential clearances and issue a suitable one. This approach is useful for comparative

examinations of various control strategies, but it does not address important task management and aircraft ownership issues. The remainder of this paper describes agents integrated within TCSim for the DAG-TM baseline terminal area environment that draw from lessons learned from the agents developed for en route operations.

2 ATCAgent Model

For this case study, a model of terminal-area ATC was implemented in an ATCAgent object that controls aircraft in TCSim. Like the CATS-based en route agents, the ATCAgents use a high-level activity model to determine the activity that the agent should perform during a given processing cycle. Each ATCAgent also maintains a 'picture' of the aircraft within their 'area of regard' (e.g, [9]). The picture is comprised of low-level models that encapsulate both what the agent knows about each aircraft, and what activities the agent could perform to control each aircraft. The following sections describe, in turn, the high level model, low-level models for different aspects of ATC, sector-specific knowledge, and control strategies of the ATCAgents.

2.1 High-Level Activities

Figure 1 shows the high-level activity model implemented in the ATCAgents. The model represents 'states of control' (a form consistent with the aircraft-specific models described later). The high-level model controls the activities the agent performs on a given processing cycle. After updating its picture of the current traffic situation, an agent formulates an agenda of activities. Unlike the en route agents, which address activities of a particular type at a given time, the ATCAgent agenda includes all the activities the agent could do. The ATCAgent then sorts, groups, and culls the activities on its agenda according to their types and specified 'urgencies' to arrive at a set to perform. If the refined agenda includes control activities, the agent evaluates alternatives to arrive at a set of specific clearances to issue. Once the agenda is complete, the ATCAgent executes each action in turn, removing each from the agenda until none remain. It then updates its picture and begins the cycle anew.

The activities in the high-level model have specified default durations used to control the timing of activities. However, in most cases the actual duration of an activity is determined as a function of the number of aircraft, agenda items, or alternatives to consider. For example, as the
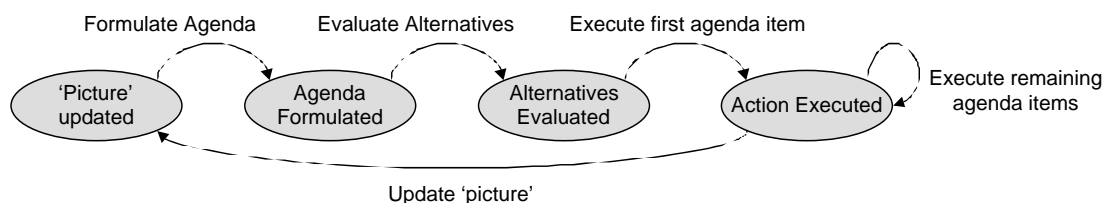

Figure 1. High-level activity model

number of aircraft in the area of regard or the number of clearance alternatives to consider increases, the time required to update the picture or evaluate the alternatives increases—enabling the agents to produce behavior that reflects variations in task load.

2.2 Flight Progress Activities

In addition to safely separating aircraft, an important aspect of ATC is delivering aircraft where they want to go, in states acceptable to the downstream air traffic controller. For each aircraft in their picture, an ATCAgent therefore maintains a model of activities required for flight progress. For the present case study, the model is depicted in Figure 2. The model is expressed in terms of progress states. The boxes in Figure 2 indicate which ATCAgent should perform the activity that effects the transition to that state. Aircraft in the case study enter the terminal area on FMS routes, but must be cleared for an 'FMS descent transition' in order to continue descending toward the final approach. The Feeder controller is responsible for ensuring aircraft are cleared to continue descending. Aircraft similarly require an approach clearance to intercept the final approach that is the responsibility of the Final controller.

2.3 Ownership Activities

Ownership status is another aircraft attribute that is crucial for ATM in which a controller is responsible for a specific sector of airspace. The ownership status determines which controller 'owns' an aircraft and is therefore eligible to issue clearances to the aircraft. Each aircraft in an ATCAgent's picture therefore includes a model of ownership (Figure 3). Ownership models for a given aircraft are synchronized between agents. For example, an aircraft in the 'handed off by me' state for one agent will be in the 'handed off to me' state for the downstream agent. When an agent's ownership model for an aircraft is in one of the states outlined in bold in Figure 3, the agent can issue clearances to the aircraft. Handoff initiation and acceptance activities are performed via cursor actions on controller's 'scope.' Initiating a handoff causes it to flash on the downstream controller's scope. Frequency transfers are performed via simulated voice clearances to the aircraft (e.g., "aircraft 123, contact regional approach on 118.1").
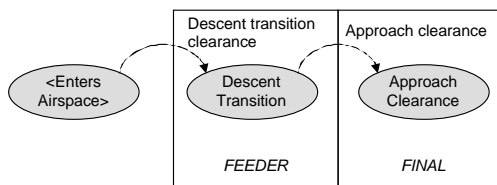
The model shown in Figure 3 represents the general case where other agents are controlling traffic both upstream and downstream of a given controller. For this case study with two controllers, there is no ATCAgent upstream of the Feeder controller, nor downstream from the Final controller. Thus, reduced models are implemented to ensure the Feeder gains control of an aircraft upon accepting the handoff, and the Final relinquishes control upon initiating the handoff.

2.4 Control Activities

The last set of models an ATCAgent maintains for each aircraft represent the state of control and associated clearance activities an aircraft is eligible for. Figure 4 shows control activity models for the lateral, vertical, and speed dimensions. The actions that effect each transition correspond to a clearance of the given type (e.g., a clearance direct to a waypoint establishes the lateral control model in the 'direct-to' state). The dashed transition in Figure 4 indicates that an aircraft flying direct to a waypoint on the charted route automatically transitions to the 'charted route' state upon crossing the waypoint. The 'charted' states are outlined in bold in Figure 4 to indicate that they are both the initial states of aircraft entering the case study airspace, and the goal states for the Final controller (i.e., the Final controller should establish each aircraft on the charted approach before relinquishing control). Specific clearance values are stored within a clearance state node in the control model (e.g., the 'heading vector' node holds the vector issued, say, 060). Control models are transferred to a downstream controller when aircraft ownership changes. This simulates coordination activities between controllers observed whenever aircraft are in off-nominal clearance states when control is transferred. For example, an upstream controller might tell
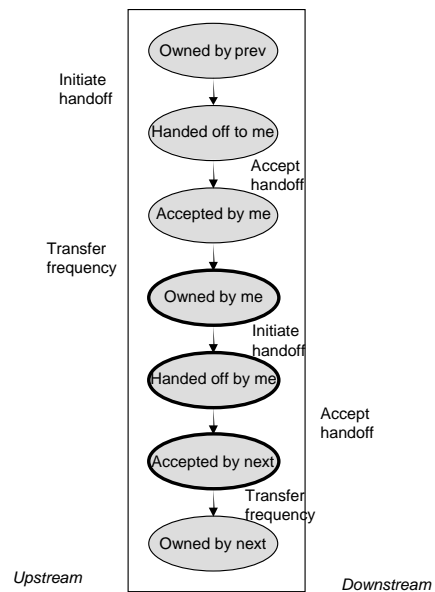

Figure 2. Flight progress model.
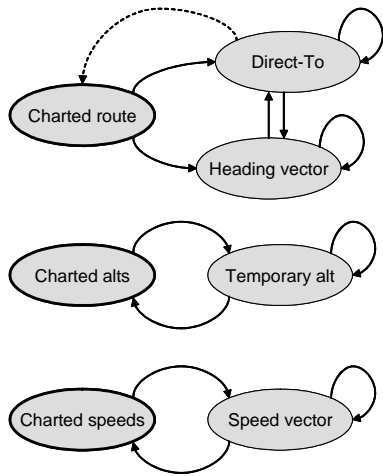

Figure 3. Ownership model.

Figure 4. Control activity models for each control dimension

a downstream controller, "aircraft 123 is on a 060 heading."

ATCAgents select control activities by evaluating alternatives. For a set of aircraft in conflict, an ATCAgent consults its picture and extracts the control status of each aircraft in the set. Links leaving the current control states in each dimension identify the set of feasible clearances for each aircraft. For example, if two aircraft are both in the 'charted' states in all dimensions, an ATCAgent could issue a direct-to clearance, heading vector, speed vector, or temporary altitude (or some combination) to either (or both) aircraft.

For agenda-management purposes, each aircraft in an ATCAgent's picture has a 'control' placeholder. If an aircraft's control placeholder appears in the current agenda, the ATCAgent assesses potential clearances according to specified control strategies (described below). When the agent finalizes its agenda, control activities are interleaved with any required ownership activities. For example, in this case study, the Feeder controller ATCAgent can determine that an aircraft needs a clearance, but that it must first accept the handoff to attain ownership of the aircraft. The Feeder agent then performs these activities in the required order.

2.5 Sector-Specific Knowledge

A particular airspace sector has several characteristics that distinguish it from other sectors. The ATCAgents therefore include this sector-specific knowledge, specified in text files that are read upon initialization. The first element is the area of regard that roughly corresponds to the airspace viewable on an air traffic controller's scope. Figure 5 shows the case study airspace, along with the areas of regard for the Feeder and Final controllers.
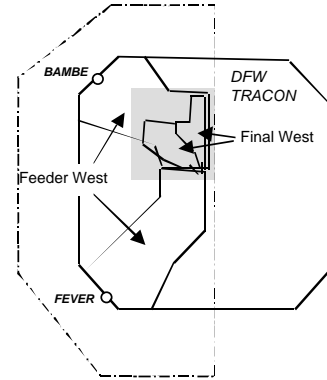

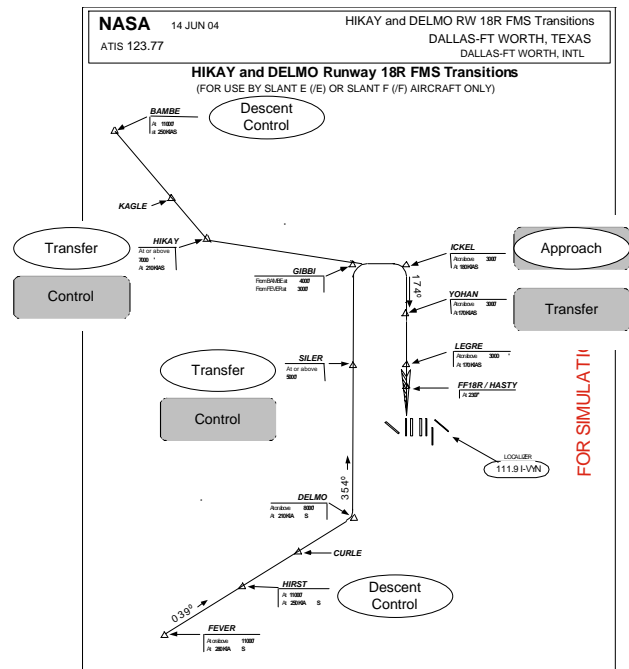Figure 5. Case study airspace showing Feeder (dashed) and Final (gray) areas of regard.

Figure 6. Charted routes that define standard flows, annotated with urgency reference points.

Sectors are also characterized by standard traffic 'flows' [9]. ATCAgents identify aircraft that belong to specific traffic flows and use sector-specific flow knowledge to assign the 'urgencies' of various activities, in order to guide agenda formulation. Figure 6 shows two charted FMS flows through the case study airspace, annotated with the urgencies various locations are associated with. Ovals designate Feeder ATCAgent urgency reference points; rounded rectangles designate Final ATCAgent urgency reference points. For example, for aircraft in the northwest flow to runway 18R, the descent transition clearance and control urgency reference point is BAMBE. SILER is the reference point the Feeder ATCAgent uses to determine the urgency of transferring control of southwest flow aircraft to the Final ATCAgent. SILER is also the point the Final ATCAgent uses to

establish the urgency of controlling southwest flow aircraft. In cases where ATCAgents use flow points alone to establish the urgency of activities (i.e., before it has performed some control activity on an aircraft), the sector-specific reference points in essence establish a first-come-first-served priority for addressing arriving aircraft.

2.6 Control Strategy Models

Control strategy models dictate what clearance alternatives to evaluate, and how they should be evaluated. ATCAgents implement a ClearanceEvaluator class that represents a particular control strategy. Together with file-specfied sector knowledge and aircraft-specific activity models, this represents a significant architectural improvement from air traffic controller agents developed previously because it affords considerable flexibility in testing the effectiveness of various control strategies. Various schemes, such as rules [5] or cost functions [1], can be implemented for comparison. For the present case study, Feeder and Final control strategies are represented in a single model. A control strategy might specify that the Feeder ATCAgent should first try to issue speed clearances, while the Final ATCAgent should first try to issue heading vectors.

A control strategy also guides planning by assigning urgencies to nodes in the control models described above. While the en route agents encoded plans explicitly, the ATCAgents need only to assign an urgency to a control model node that indicates it is part of a larger plan. As an example, suppose that an ATCAgent identifies a heading vector as the current clearance to issue. By assigning an appropriate urgency and value to the 'direct to' node, the ATCAgent establishes a plan to revisit the aircraft and clear it direct to the specified waypoint. The assigned urgency is propagated up to the 'control' placeholder node, so that the plan appears on the agenda for execution at the appropriate time. This scheme also supports plan adaptation by enabling an ATCAgent to evaluate other values (e.g., a different waypoint than planned) and links out of the current clearance state are possible (e.g., heading vector instead?). Plan adaptation may also simply entail reassigning the urgency of the planned clearance to issue it later than initially planned.

3 Performance Evaluation

For the present case study, four control strategy models were developed, as follows:

- Model I: A baseline 'hands-off' model in which agents addressed flight progress requirements by issuing descent transition and approach clearances, but otherwise kept aircraft on their charted routes. This model was used as a baseline to assess how much control is actually required for the traffic flows in the DAG-TM traffic scenarios.

- Model II: A model that attempted to mimic control strategies observed during the DAG-TM simulation trials in which the Feeder controller issued 'short-cut' clearances to aircraft. The model otherwise used speed clearances to adjust spacing.

- Model III: A model in which both controllers used speed clearances alone, enabling aircraft to remain on their charted FMS lateral routes and vertical profiles.

- Model IV: A model in which both ATCAgents attempted to evaluate and select speed, direct-to, vectoring, and altitude clearances to achieve the required spacing at the final approach fix.

The ATCAgents controlled traffic using each of these control strategy models on each of DAG-TM terminal-area traffic scenarios. Spacing accuracy at the final approach fix was used as the basis for comparison to observed professional air traffic controller performance averaged over sixteen trials (two for each traffic scenario) in the DAG-TM baseline condition.

Figure 7 depicts the results in spacing-accuracy histograms (plotted as lines). Model I indicates many aircraft in the test scenarios are set up reasonably well for spacing. Model II shows that aggressive short-cuts by the Feeder ATCAgent are not effectively managed by the Final ATCAgent using speed clearances alone. A broad range of spacing accuracies, including unsafe spacing, are observed. Model III, the speed control model, compares most favorably to the DAG-TM baseline data; except for a few aircraft in the uncoordinated flows, most aircraft are set up to be safely spaced using speed clearances. Model IV is an attempt to enable ATCAgents to use all sorts of clearances, just as the human controllers did. However, Model IV yields relatively large numbers of aircraft pairs with far too little or too much spacing. In general, the results indicate that performance deteriorates as ATCAgents are less
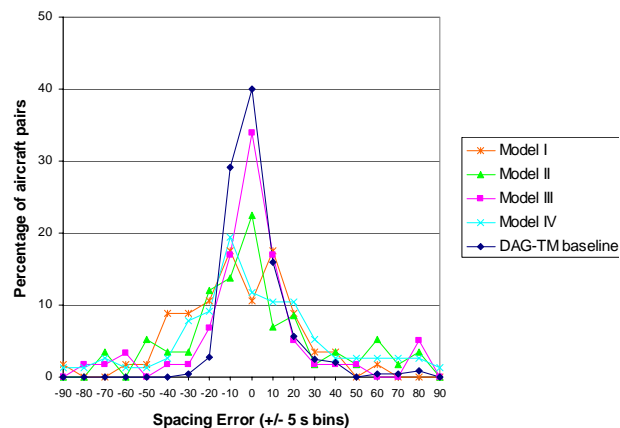

Figure 7. Spacing accuracy results for the case study.

constrained in the clearances they issue, and when heading vector clearances result in aircraft being off, and therefore no longer constrained by, their charted FMS routes. This makes predicting future aircraft locations difficult for the ATCAgents. Additional research is required to develop effective broad-spectrum control strategy models that address these issues.

4   Conclusion

This paper has presented a case study in which computational models of terminal-area air traffic controllers were evaluated against human performance on the same simulated traffic scenarios. The results suggest that the agents perform best when the 'entropy' of the traffic situation remains low; they perform worst using a control strategy with many clearance options, leading to many plans and removal of constraints on the aircraft trajectories. The reasons are thought to stem from deterioration in the quality of predictions about future aircraft locations, and from deficiencies in the preliminary implementation of 'urgencies' that guide when plans should be adapted. Future research will address these issues and provide more detailed validation of agent performance.

References

[1]   A. Bayen, P. Grieder, and C. Tomlin, "A control theoretic predictive model for sector-based air traffic flow," AIAA 2002-5011, American Institute of Aeronautics and Astronautics, Reston, VA, 2002.

[2]   T. Callantine, "Agents for analysis and design of complex systems," Proc. 2001 IEEE Intl. Conf. Systems, Man, and Cybernetics, Tucson, pp. 567-573, 2001.

[3]   T. Callantine, "Air traffic management domain and control strategy analysis," Proc. 2004 IEEE Intl. Conf. Systems, Man, and Cybernetics, The Hague, pp. 6268-6273.

[4]   T. Callantine, "Modeling and simulation of trajectory-oriented air traffic management concepts," AIAA-2004-5260, American Institute of Aeronautics and Astronautics, Reston, VA, 2004.

[5]   T. Callantine, "Performance evaluation of a computational model of en route air traffic control," Proc. 13th Intl. Symp. Aviation Psychology, Oklahoma City, OK, pp. 86-91, 2005.

[6]   T. Callantine, P. Lee, J. Mercer, T. Prevôt, and E. Palmer, "Air and ground simulation of terminal-area FMS arrivals with airborne spacing and merging, Proc. 6th USA/Europe ATM R&D Seminar, Baltimore, MD, August, 2005.

[7]   T. Callantine, and E. Palmer, "Fast-time simulation studies of terminal-area spacing and merging concepts," Proc. 22nd Digital Avionics Systems Conf., Indianapolis, IEEE: 03CH37449 (CD-ROM), October, 2003.

[8]   K. Corker, B. Gore, K. Fleming, and J. Lane, "Free flight and context of control: Experiments and modeling to determine the impact of distributed air-ground air traffic management on safety and procedures," Proc. 3rd USA/Europe ATM R&D Seminar, Naples, 2000.

[9]   H. Davison and J. Hansman, "Using structure as the basis for abstraction in air traffic control," Proc. 12th Intl. Symp. Aviation Psychology, The Ohio State University, Columbus, OH, 2003.

[10] K. Harper, S. Guarino, A. White, M. Hanson, K. Bilimoria, and D. Mulfinger, "An agent-based approach to aircraft conflict resolution with constraints," AIAA 2002-4552, American Institute of Aeronautics and Astronautics, Reston, VA, 2002.

[11] H. Hexmoor, and T. Heng, T. (2000). Air traffic control and alert agent. Proc. 4th Intl Conf on Autonomous Agents, Berlin, pp. 237-238, 2000.

[12] J. Krozel, M. Peters, K. Bilimoria, C. Lee, and J. Mitchell, "System performance characteristics of centralized and decentralized air traffic separation strategies," Proc. 4th USA/Europe ATM R&D Seminar, Santa Fe, NM, 2001.

[13] K. Leiden, "Human performance modeling of en route controllers (RTO-55 Final Report)," Micro Analysis and Design, Inc., Boulder, CO, 2000.

[14] S. Mondoloni, "Development of an enroute conflict resolution rulebase for the reorganized air traffic control mathematical simulator," NARIM-A09008-01, CSSI, Inc., Washington, D.C., 1998.

[15] C. Niessen, C. and K. Eyferth, "A model of the air traffic controller´s situation awareness," *Safety Science*, Vol. 37, pp. 187-202, 2001.

[16] A. Nunes, and R. Mogford, "Identifying control strategies that support the 'picture'", Proc. 47th Annual Meeting of the Human Factors and Ergonomics Society, Denver, pp. 71-75, 2003.

[17] A. Pritchett, S. Lee, M. Abkin, A. Gilgur, R. Bea, K. Corker, S. Verma, A. Jadhav, T. Reynolds, L. Vigeant-Langlois, G. Gosling, "Examining air transportation safety issues through agent-based simulation incorporating human performance models," Proc. 21st Digital Avionics Systems Conf., Irvine, 7.A.5-1-7.A.5-13 (CD-ROM), 2002.