Comparing Route Selection Strategies in Collaborative Traffic Flow Management

Shawn R. Wolfe
NASA
Shawn.R.Wolfe@nasa.gov

Peter A. Jarvis
Perot Systems Government Services
pjarvis@email.arc.nasa.gov

Francis Y. Enomoto
NASA
Francis.Y.Enomoto@nasa.gov

Maarten Sierhuis
RIACS
Maarten.Sierhuis-1@nasa.gov

Abstract

Today's air traffic management system is not expected to scale to the projected increase in traffic over the next two decades. Enhanced collaboration between the controllers and the users of the airspace is one of the new traffic flow management concepts being developed, and may substantially improve airspace utilization. We argue that an agent-based simulation is well suited to explore and validate these new concepts, and present our initial simulation design. We evaluate several simple route selection strategies and assignment policies that nonetheless show interesting properties of the proposed concept, and conclude with the challenges of validating the proposed concept through simulation and future work.

1. Introduction

The Federal Aviation Administration (FAA) is responsible for directing aircraft through the United States National Airspace System (NAS). Safety is the highest priority in aircraft operations, and concerns for it necessarily take precedence over other objectives, such as on-schedule performance and airline satisfaction. Though safety can typically be accommodated without significant compromise in nominal conditions, disruptions to the NAS from reduced airspace capacity (e.g., due to poor weather) or from increased demand often necessitates mitigating actions – resulting in flight delays and an increased cost in airline operations.

The Joint Planning and Development Office (JPDO), a public/private partnership between several government agencies and industry leaders, is responsible for planning the future of the air transportation system. Their forecast of air traffic in 2025 shows an increase of two to three times over present day levels [1]. Recent simulations [2] show that the NAS, as it is managed today, cannot reasonably handle this increase in demand, as the average delay per flight would increase from four minutes to over five hours. Accordingly, NASA is investing in several fundamental research projects that explore new aviation technologies and operational concepts that may revolutionize the NAS.

The Next Generation Air Transportation System (NGATS) project is a multi-faceted research effort to address issues with the NAS. One such facet is the area of Collaborative Traffic Flow Management (CTFM). The act of guiding an aircraft through the airspace naturally entails cooperation between two parties, namely the controller of the aircraft (i.e., the airline) and the controller of the airspace (i.e., the FAA). The final responsibility for the aircraft remains with the pilot, but in today's system, the flow of traffic is primarily handled by three entities: the FAA's Air Traffic Control System Command Center (ATCSCC) and Traffic Management Units (TMUs), and the individual airlines' Airline Operation Centers (AOCs). A new CTFM concept of operations [3] has been developed to increase both the efficiency of the NAS, and the satisfaction level of the airlines.

In this paper, we describe some of the main features of this new concept for traffic flow management, and our efforts to simulate it with a multi-agent simulation environment. The work we present is our simulation model developed in the first half-year of a multiyear effort. We discuss the comparative results of different CTFM strategies between the airlines' AOCs and the FAA's TMUs. The simulation is agent-based, meaning that we simulate the AOCs and TMUs as simulated organizations that communicate their traffic flow preferences, based on a continuously changing airspace due to individual airline flight prioritizations and external events, such as weather. The simulation is run with each airline using a different flight prioritization strategy. Running the simulation multiple times, with different strategies and schemes allows us to compare the impact of these choices on the traffic flow as a whole, as well as for the individual airlines.

2. Approach

The development of the concept of operations began with field observations of work practice in several representative TMUs, AOCs and the ATCSCC [4], in support of Distributed Air/Ground Traffic Management (DAG-TM) operational concept, a research effort that predated NGATS. These field observations show how work is actually performed, and continue to act as a guide to our agent development effort. Based on these observations, a list of operational issues was compiled, and several changes were suggested to address these issues in the CTFM concept of operations.

The primary finding of this study was that the current TFM system limited the degree of collaborative decision making that actually occurred in traffic flow management. First, the FAA and the airlines do not have shared impact assessment tools, and therefore make decisions based on divergent predictions of airspace availability. Second, the FAA generally does not know the specific concerns of the airline, and so must make routing decisions without regard to airline preferences. Finally, the bulk of the problem solving responsibility falls upon the FAA, increasing an already high workload. Necessarily, mitigations must be chosen that do not push this workload beyond acceptable levels, but such decisions often come late, are overly conservative, and decrease airline efficiency.

The CTFM concept of operations addresses these issues in several ways. First, impact assessment information is to be shared amongst the FAA and the airlines. The transmission of this information must be fast and reliable in order to be effective, through electronic transmissions, rather than by phone, as is often the case today. Second, preference information must available in the planning process, either by communicating these to the FAA or directly by the airlines. Support tools (such as planning tools) will be needed to allow the FAA to consider these preferences without increasing their workload. Finally, when appropriate, the airlines should be allowed to choose their own mitigations. This enables the airlines to participate directly in the problem solving process, allowing them to choose actions that suit their business models, rather than forcing the FAA to dictate less preferable solutions.

We have built an initial simulation using Brahms [5], a multi-agent simulation environment. Traditional aviation simulation methods that focus on physics are not well suited to recreate the subtleties of human decision-making; by contrast, an agent-based paradigm allows modeling people and their behavior more naturally. Agents offer a convenient method of encapsulating the beliefs, desires and intentions (BDI) of people (both at the organizational level and at the individual level) that form the basis of our simulation, which focuses on human decision-making rather than physical aspects of air flight. Brahms is particularly well suited to this simulation as it is a BDI architecture, and was developed to simulate work practice, the groundwork for the operational concept and our ultimate simulation goal. As collaboration is our key concept, we must correctly simulate communications and maintain distinct internal states for each entity, which is easily modeled in the agent paradigm.

In our simulation, agents act as proxies for their human counterparts. This allows us to simulate the people and their interactions initially, and later replace proxies with human operators in a humans-in-the-loop simulation. Proxy agents that perform well in simulation may also lead to improvements in automation, as their internal logic may be transferred to assistive tools or automate certain tasks [6] [7] [8].

3. Brahms

Brahms is a modeling and simulation environment for analyzing human work practice and for developing intelligent software agents to support work practice in organizations. Brahms can run in different simulation and runtime modes on distributed platforms, enabling flexible integration of people, hardware-software systems, and other simulations. Brahms was originally conceived as a business process modeling and simulation tool that incorporates the *social systems of work*, by illuminating how formal process flow descriptions relate to people's actual located activities in the workplace [9]. To simulate human behavior at the work practice level, one must model how people work together as individuals in organizations, performing both individual and teamwork activities. The Brahms language is unique in that it not only models both individual agent and group behavior, but also systems and artifact behavior, interpersonal interaction, as well as interaction of people, systems and objects with the environment. Most other multiagent languages leave out artifacts and the interaction with the environment, making it difficult to develop a holistic model of real-world situations (c.f. [10]). Brahms is an agent language that operationalizes a theory for modeling work practice, allowing a researcher to develop models of human activity behavior that corresponds with how people actually behave in the real world [11].

4. Related Work

The Future ATM Concepts Evaluation Tool (FACET) [12] is a NASA-developed tool for simulating air traffic flow. FACET contains modules that concentrate on trajectory modeling, weather modeling, and also contains a model of the airspace structure, including the ARTCC regions, sectors, and air routes. FACET can act either as a simulator or as a playback mechanism, using either from historical data or from a live data feed from the FAA. FACET has been integrated into a commercial product, Flight Explorer [13], which is used by the majority of major U.S. airlines. FACET is not an agent-based simulation, concentrating primarily on the physical aspects of air traffic flow, but does include other concepts such as controller workload and traffic management initiatives.

The Airspace Concept Evaluation System (ACES) [14], also developed by NASA, is a distributed agent-based simulation of the NAS. ACES supports the Department of Defense's High Level Architecture (HLA), which has enabled the integration of several simulations into the overall system. As ACES is focused on the entire NAS, the simulation includes traffic flow management [15], but is not specifically focused on TFM. The agents of ACES follow an "activity centric paradigm," which is compatible with the Brahms framework but does not model as much of the internal state of the agent. Though FACET and ACES seem superficially similar, they were designed to address different simulation needs and have different strengths.

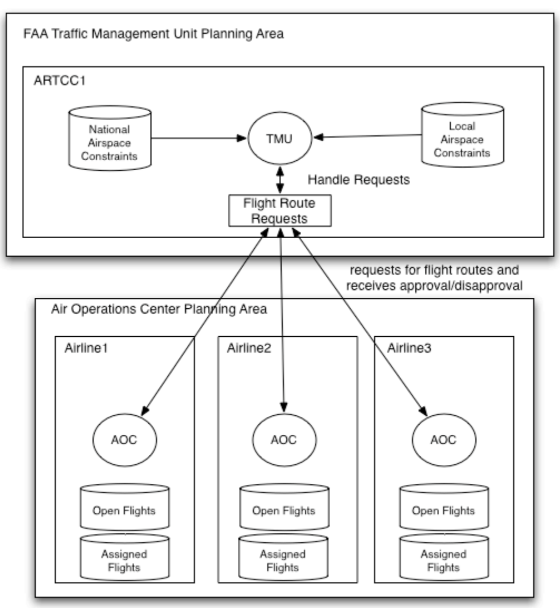

Figure 1. Agent architecture.

5. Simulating CTFM

Commercial aircraft generally follow structured air traffic routes, a sort of "highway in the sky". Though typically not observable to the naked eye, aircraft are often queued up in these flight routes, much like cars on a busy freeway. These routes typically pass through many sectors, each with their own sector controllers. In our simplified model of air traffic flow management, we have combined the sectors and their controllers into a modified concept of air traffic routes, which we define as having a capacity in lieu of the sectors. Air traffic follows the routes from beginning to end, and any disruption to the route (due to weather or other factors) affects the capacity on the route as a whole. Unanticipated events (such as weather) may change the capacity during the planning phase. We have limited the scope of our simulation to include only a single time window for the flights, so delayed flights are not explicitly scheduled in our simulation.

Figure 1 provides an overview of our agent architecture. We have created agents corresponding to the decision-making entities of the real world, namely a TMU agent for each ARTCC and an AOC agent for each airline. Real world operations involve national constraints as well as local constraints, and necessitate communication between TMUs (as routes will pass through multiple ARTCCs), but we have restricted our simulation to traffic within a single ARTCC. As a result, a TMU agent need only communicate with the AOC agents, and not other TMU agents. The TMU agent is responsible for assessing the impact of constraints on the ARTCC airspace, calculates the route demand based on requests from the AOC agents, and broadcasts the route status (under capacity, at capacity, or oversubscribed) to the AOC agents.

The AOC agents consult their flight schedules and communicate the priority of each flight to the TMU agent. Depending on the scheme used (see section 4), either the TMU agent assigns routes to the flights and communicates this to the AOC agents, or the AOC agents request route assignments that the TMU agent must either accept or reject.

6. Scenario Generation

In order to test our simulation, we have created several scenarios that correspond to low, medium, and high traffic conditions. Though artificial, these scenarios are built from observations of the actual flight schedules and are reasonably accurate, given the overall fidelity of our model. Scenarios are created from a list of airports, air traffic centers, airlines, and flight schedules. All routes are defined as the great circle between city pairs, with costlier alternate routes defined. We assume there are no spatial conflicts or shared airway segments amongst routes, though these do exist in today's NAS. Though our model can handle an arbitrary number of air traffic centers, airports, and airlines, we have restricted ourselves to one ARTCC, up to seven airports, and three airlines.

Airspace demand is generated from flight schedules that approximate the distribution of flights flown throughout the day for the airlines used in the scenarios. A flight is defined by the airline, city pair, requested departure time, and airspeed. Each flight also describes the number of passengers, connecting passengers, and connecting flight crew that are used to define that flight's "value" to the airline.

Our scenario generator also creates KML files in order to display graphically the many of the important features of the scenario. This allows display of the scenario in any KML-compatible browser, such as Google Earth (see Figure 2).

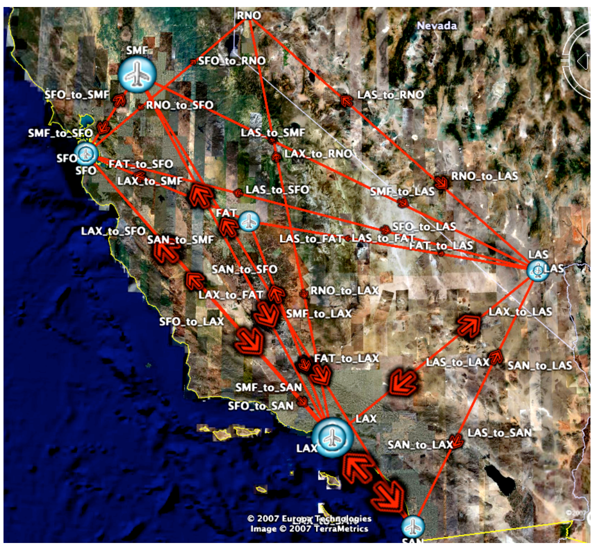

Figure 2. Medium traffic scenario.

7. Airline Strategies

We have developed conservative, moderate, and aggressive route allocation strategies for the airlines. All strategies will initially choose the best route for every flight, and only differ in how they react to an oversubscribed route. The aggressive strategy (A) is the simplest, as the aggressive airline always requests the best route for all flights even when that route is known to be oversubscribed. The hope is that the other airlines using that route will choose alternate routes and solve the problem. The risk of the aggressive strategy is that the route remains oversubscribed, and some flights will not be assigned a route (even though an alternate route may exist).

The conservative strategy (C) is designed to find a route assigned to every flight. When a flight is assigned to an oversubscribed route, the conservative airline will not only reassign a percentage of flights to an alternate route, but it will also choose the *least* preferred route for the flight. This is done because the second most preferred route may also become oversubscribed, and since planning time is limited, always choosing the next best available route may result in no route assignment. Therefore, the conservative strategy is designed to increase the chance that every flight is assigned a route, but may leave unused capacity on more desirable routes.

The moderate strategy (M) is a blend between the aggressive and conservative strategies. When using the moderate strategy, an airline will reassign a percentage of their flights to the next best available route. Unlike the aggressive strategy, the moderate strategy allows an airline to take corrective action when a route is oversubscribed. Unlike the conservative strategy, an airline using the moderate strategy will try to minimize costs by choosing the next best alternative.

Though the moderate strategy may seem to be the only reasonable strategy, all three strategies (or variants thereof) are actually used in airline operations. This is because different airlines have different overall business models, and may also use different strategies depending on the particular situation. Such complexity currently falls outside of our simulation, however.

8. Route Selection Schemes

8.1. Blue Sky

The blue sky scheme gives the ideal performance for each flight by removing all constraints on the airspace. No weather or other disruptions affect airspace capacity in this simulation variant, and all route capacities are infinite. Hence, all flights are given the optimal route. Of course, the blue sky scheme is unrealistic since limited capacity and airspace disruptions are real constraints on route planning, but it does represent a theoretical upper bound on performance.

8.2. Current Operations

In the current operational model, the FAA is unaware of the specific concerns of the airlines on a flight-by-flight basis. Therefore, the FAA is unable to consider the preferences of the airlines when making route allocation decisions. In this scheme, the TMU agent will honor the initial route selection request from any airline if there is enough capacity on the requested route. When the route is at capacity, the TMU agent will reassign the flight to the next best available route without any consultation from the airline. This strategy ensures that the best routes are used, but a less important flight may be assigned a better route than a higher value flight.

8.3. FAA Global Maxima Planning

The FAA Global Maxima Planning scheme takes a step towards the CTFM concept of operations by inserting the specific prioritizations of each flight into the flight planning process. In this scheme, the TMU agent knows the flight value of each flight. Since this value is objective, there is not a possibility for "gaming" the system by choosing artificially high flight values. The TMU agent uses a greedy algorithm to obtain a globally optimal route assignment to the flights. However, the flight assignments for a particular airline are not likely to be optimal. Moreover, airlines with a greater percentage of high valued flights will receive preferential assignments, and so inequities may be present in the solution.

8.4. Direct Airline Planning

The Direct Airline Planning scheme gives even greater freedom to the airlines and largely removes the FAA from the planning process. In this scheme, the TMU agent will continue to accept route requests, calculate and broadcast current route capacities and demand, and approve route assignments only when no conflicts exist. The TMU agent will not reroute flights, as that responsibility has been passed on to the airlines, so the AOC agents must reduce demand on oversubscribed routes independently.

9. Metrics

We have instrumented our simulation to provide statistics on an airline's performance. Currently, we use the same metrics to evaluate each airline even though they may have differing business models.

For a specific flight F of airline A_F, we define the following quantities:

p_c = passengers with connecting flights;
p_u = passengers without connecting flights;
c_c = onboard crew members a connecting flight;
t_F, = the actual flight time of F, in minutes;
t_B, = the optimal flight time of F (from the Blue Sky simulation), in minutes.

Each flight is assigned a flight value, which is a heuristic measure of the importance the flight to the airline. We define v_F , the flight value of F, as

$$v_F = p_u + 3p_c + 5c_c \qquad (1)$$

When F is assigned a route, we calculate d_F, the delay for flight F, as follows:

$$d_F = t_F - t_B \qquad (2)$$

When F is not assigned a route, we assume a standard 60 minutes of delay.

We needed a measure the total passenger delay incurred by flight F, either through an immediate delay or through problems with later connecting flights. We assume that when a passenger with a connecting flight is delayed, on average, they will experience an additional two-hour delay. When connecting crew members are delayed, their personal delay does not count (since they are not considered passengers in our simulation), but they are likely to delay the departure for their connecting flight, which in turn impacts many passengers. Therefore, we assume on average, any delay of a connecting crew member results in a total of five additional hours of passenger delay. Combining this with the above formulae, we calculate the total passenger delay incurred by flight F, d_T, in minutes, as

$$d_T = (p_u * d_F) + (p_c + d_F) + 60p_c + 300c_c \qquad (3)$$

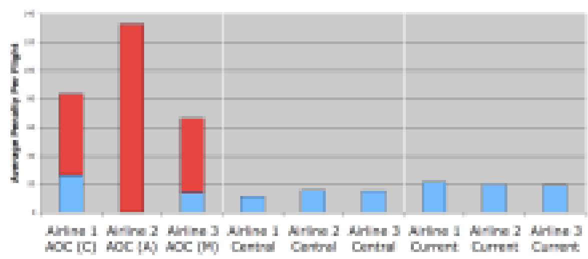

Figure 3. Comparing solution strategies.

10. Simulation Results

In this section, we present the results of several different simulations of the high traffic scenario. All graphs display the average passenger delay per airline, in minutes. The light blue portion shows the delay contributed by flights given route assignments, and the dark red portion shows the delay contributed by the flights without route assignments.

Figure 3 compares the Direct Airline planning, FAA Global Maxima planning, and Current Operations planning schemes. As expected, the FAA Global Maxima scheme produces a better solution than the Current Operations scheme as the flight values are taken into consideration. Surprisingly, the Direct Airline Planning scheme performed poorly, because the airlines were unable to find routes for all their flights, and the resulting delay (shown in dark red) is large.

Figure 4 shows a very different picture when different blends of AOC strategies are used. When every airline uses either the moderate or conservative strategies, all flights were successfully assigned routes and reasonable results were achieved. By contrast, the uncompromising aggressive strategy typically has worse performance overall, but also negatively impacts other airlines. However, the aggressive strategy in the right situation may still outperform other airlines, as can be seen in the last simulation run.

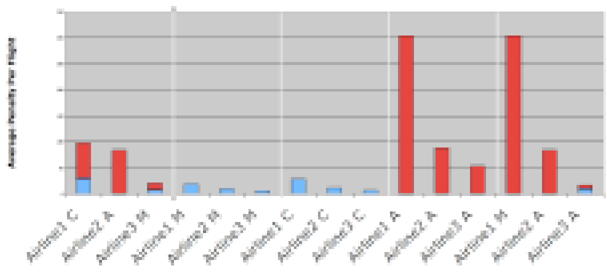

Figure 4. Interactions between AOC strategies.

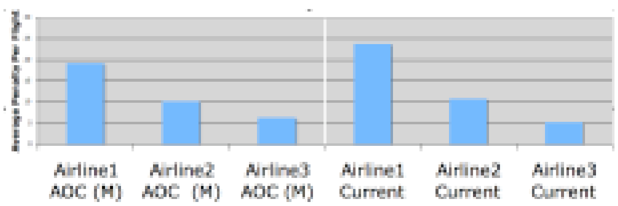

Figure 5. AOC planning compared to present day.

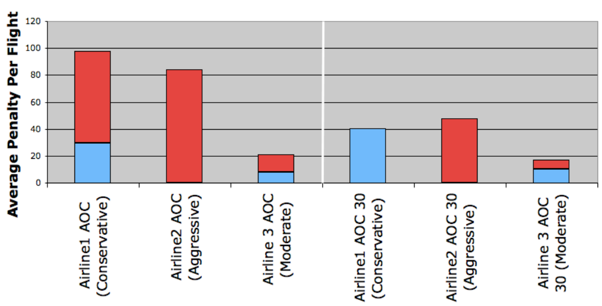

Figure 6. Effect of additional time.

Figure 5 shows that the more successful mixes of airline strategies compare well with the Current Operations scheme. Indeed, using all moderate strategies produced lower overall delay (averaged across all airlines), but also produced average delays that are slightly more equivalent across the airlines. This is important, because it establishes more delay equity across the airlines. The FAA Global Maxima scheme produces a better solution as well as better equity, but the latter may be due, in part, to a similarity in flight priority in our chosen airlines.

Finally, Figure 6 shows that increasing the time available for planning greatly improves the simulation results, even when using a less beneficial blend of AOC strategies. In our simulation, the AOC agents create their plans from scratch in the time available. In reality, nominal route assignments exist, and in most situations can be reused in the planning process.

11. Observations

The limitations (in both scope and fidelity) of our route simulator prevent us from making strong conclusions based on our overall results and delay calculations. Nonetheless, we observed interesting phenomena that may correspond to real world behavior and could influence the development of CTFM.

In our simulation, delays on the ground (when no route was chosen) were far more costly than delays in the air (when a longer route was chosen). Though this corresponds with operations today, the reasons are different: airlines will delay flights on the ground when no feasible route is available, but in our simulation, a reasonable alternative always existed. In particular, this shows the importance of using available routes, something that should be stressed in the AOC planning strategies.

In particular, the aggressive strategy failed primarily because of its unwillingness to compromise and use available (but suboptimal) routes. The aggressive strategy may be effective in certain limited situations, but is clearly ineffective overall. Surprisingly, the aggressive strategy was damaging not only to the AOC agent implementing it, but also to the overall system. Through its unreasonable demands, the aggressive AOC agent created demand-capacity imbalances that were difficult, or even impossible, for the more cooperative airlines to resolve.

The problems caused by aggressive behavior indicate limitations of a laissez-faire approach by the FAA. In our simulation, the TMU agent will not resolve any overcapacity problems- the airlines must reduce the demand until it is equal or less than the capacity, and if they fail, the route remains completely unused. This leaves the overall system vulnerable to a rogue (i.e., uncooperative) AOC. No airline is likely to pursue this strategy to the extreme that we have simulated it, but a crafty airline may be able to exploit specific situations to detriment of their competitors. In order to insure the integrity of the overall system, the FAA needs additional means – either through incentives or direct action – to stop the airlines from inappropriately creating problems for their competitors.

Finally, when reasonable strategies were used, the Direct AOC Planning scheme outperformed the Current Operations scheme in our simulation. This suggests the potential of CTFM, but significant work remains to validate the proposed operational concept.

12. Future Work

Our current model focuses only on the route selection aspect of CTFM and makes many additional simplifications. Actual traffic flow management is much more complicated. Traffic flow managers have a variety of actions and restrictions they may use, including miles in trail (dictating the distance between aircraft passing over a fixed point), ground delay programs (delaying flights destined for a particular airport), and coded departure routes (large scale reroutes of traffic to avoid impacted areas). Also, the FAA policies and airline strategies implemented were overly simplistic and performed poorly in our experiments. Further development will include refinement and enhancement of these strategies and approaches in both the model and the concept of operations. This may include alternatives such as market-based approaches and increased distributed decision-making.

Units and individuals in each organization need to be modeled in detail, in contrast to our current simulation which models organizations as a single entities. Agents must learn about the effectiveness of choosing a strategy in particular situations, allowing them to choose a strategy based on previous experience. Modeling at the level of individuals will also make the communication clearer, both within an organization and between different organizations.

As with any simulation effort, creating an accurate model is challenging. Thus far, we have primarily relied on the guidance of experts (through interviews and literature review) to build our model and will continue to use such methods. However, expertise is generally in short supply and there is a risk that the resulting model captures the expert's conception more than it does reality. One possible alternative is to use past records of traffic management decisions to induce the model (e.g., using machine learning methods) and to validate the accuracy of the model (by testing to see if the simulation generates similar results given identical conditions). Both approaches are challenging, however, and validation is particularly so, because even different traffic flow managers do not necessarily act predictably in the same set of circumstances. For validation, what is needed is not a test of predictability, but rather a Turing test to evaluate if the agents act in a manner that seems reasonable.

Finally, we will make use of FACET to simulate the physical elements of our overall simulation, such as the properties of the airspace, routes, weather conditions, and the movement of the aircraft. This integration will require the reconciliation of the time-based model of FACET and the discrete event-based model of the Brahms simulator.

13. Acknowledgements

The authors would like to thank the NGATS ATM-Airspace project for their support of this work, as well as Dr. Karl Bilimoria for his guidance and expertise.

14. References

[1] Federal Aviation Administration, "FAA Long-Range Aerospace Forecasts Fiscal Years 2015, 2020 and 2025".

[2] Grabbe, S., B. Sridhar, and A. Mukherjee, "Regional Management of Chicago Arrival Flows", *presentation to Code AFC*, NASA Ames Research Center, 2007.

[3] Idris, H., et al., "Operational Concept for Collaborative Traffic Flow Management based on Field Observations", *AIAA 5th Aviation, Technology, Integration, and Operations Conference (ATIO)*, 2005.

[4] Idris, H., et al., "Field Observations of Interactions Between Traffic Flow Management and Airline Operations", *AIAA 6th Aviation, Technology, Integration, and Operations Conference (ATIO)*. 2006.

[5] Clancey, W.J., et al., "Advantages of Brahms for Specifying and Implementing a Multiagent Human-Robotic Exploration System", *16th International FLAIRS Conference*, St. Augustine, FL, 2003.

[6] Wolfe, S.R., "Supporting Air Traffic Flow Management with Agents", *AAAI Spring Symposium: Interaction Challenges for Intelligent Assistants*, 2007

[7] Clancey, W.J., et al., "Automating CapCom Using Mobile Agents and Robotic Assistants", *American Institute of Aeronautics and Astronautics 1st Space Exploration Conference*, Orlando, FL, 2005.

[8] Sierhuis, M., et al, "NASA's Mobile Agents Architecture: A Multi-Agent Workflow and Communication System for Planetary Exploration", *i-SAIRAS*, München, Germany, 2005.

[9] Clancey, W.J., et al., "Brahms: Simulating practice for work systems design", *International Journal on Human-Computer Studies,* **49**, 1998, pp. 831-865.

[10] Wooldridge, M. and N.R. Jennings, "Intelligent Agents: Theory and Practice", *Knowledge Engineering Review,* **10**(2), 1995, pp. 115-152.

[11] Sierhuis, M., "Modeling and Simulating Work Practice; Brahms: A multiagent modeling and simulation language for work system analysis and design", *Social Science Informatics (SWI)*, Ph.D. dissertation, 2001.

[12] Bilimoria, K., et al., "FACET: Future ATM Concepts Evaluation Tool", *3rd USA/Europe ATM 2001 R&D Seminar*, Napoli, Italy, 2000.

[13] *Flight Explorer*, Flight Explorer Inc, http://www.flightexplorer.com.

[14] Sweet, D.N., et al., "Fast-Time Simulation System for Analysis of Advanced Air Transportation Concepts", *AIAA Modeling and Simulation Technologies Conference and Exhibit*, Monterey, CA, 2002.

[15] Couluris, G.J., et al., "National Airspace System Simulation Capturing the Interactions of Air Traffic Management and Flight Trajectories", *AIAA Guidance, Navigation, and Control (GNC) Conference*. 2003.