

A Multi-Agent Simulation of Collaborative Air Traffic Flow Management

<i>Shawn R. Wolfe</i> NASA	<i>Peter A. Jarvis</i> Perot Services Government Systems	<i>Francis Y. Enomoto</i> NASA NASA Ames Research Center United States of America	<i>Maarten Sierhuis</i> USRA-RIACS/ Delft University of Technology	<i>Bart-Jan van Putten</i> USRA-RIACS/ Utrecht University	<i>Kapil S. Sheth</i> NASA
-------------------------------	--	--	--	--	-------------------------------

Today's air traffic management system is not expected to scale to the projected increase in traffic over the next two decades. Enhancing collaboration between the controllers and the users of the airspace could lessen the impact of the resulting air traffic flow problems. We summarize a new concept that has been proposed for collaborative air traffic flow management, the problems it is meant to address, and our approach to evaluating the concept. We present our initial simulation design and experimental results, using several simple route selection strategies and traffic flow management approaches. Though our model is still in an early stage of development, these results have revealed interesting properties of the proposed concept that will guide our continued development, refinement of the model, and possibly influence other studies of traffic management elsewhere. Finally, we conclude with the challenges of validating the proposed concept through simulation and future work.

Keywords: agent-based modeling and simulation, air traffic flow management, collaboration, competition

Introduction

Air traffic in the United States of America (U.S.A.) is forecasted to double or triple by the year 2025 (Pearce, 2006). Recent simulations (Mukherjee, Grabbe, & Sridhar, 2008) of this increase in demand using current air traffic management techniques yielded an increase in average delay per flight from four minutes to over five hours – a clearly unacceptable situation. Accordingly, the National Aeronautics and Space Administration (NASA) is currently exploring several new concepts that may reduce or alleviate air traffic problems. One such concept is Collaborative Air Traffic Flow Management (CATFM), which seeks to lessen the impact on airspace user operations rather than eliminate the problem. Today in the U.S.A., the Federal Aviation Administration (FAA) makes the bulk of Air Traffic Flow Management (ATFM) decisions with only limited consultation with the airlines. In CATFM, the airspace users are given more opportunities to express their preferences, choose among options, and take proactive actions. It is presumed that this will result in decreased workload for the FAA, increased airline satisfaction, and more efficient traffic flow management (Odoni, 1987).

Several questions arise when evaluating if the CATFM concept will work in the future environment. Will the airlines take advantage of new opportunities for action, or will they be passive and let the FAA continue to solve traffic problems independently? Will increasing

airline involvement decrease the FAA's workload? Will the options available to the airlines enable them to substantially increase the efficiency of their operations, in particular when many factors still remain out of their control? Will the uncoordinated actions of individual airlines increase the efficiency of the system as a whole, even though each airline is only concerned with their own operations (Waslander, Raffard, & Tomlin, 2008)? Might potential efficiency gains be offset by the actions of rogue operators, who purposely seek to interfere with the operations of a competitor (Hardin, 1968)?

Given that the CATFM concept involves many independent entities with their own beliefs and desires, we feel that the first step to answering some of these questions is through agent-based modeling and simulation. Our goal is to build a simulation of CATFM so that its strengths and weaknesses can be evaluated long before more costly human-in-the-loop simulations or limited field deployments are attempted (Wambsganss, 1996). Our simulation is in an early stage of development, but we have already found several interesting and important properties of CATFM (presented in our conclusions).

Though our study is certainly most relevant to air traffic, certain aspects are relevant to other forms of traffic as well. Our methodology can be applied to any concept of operations in these domains. Many of the basic concepts (e.g., choosing routes, traffic congestion, independent and uncoordinated agent actions) are the same and the overall structure is similar. Nonetheless, there are important differences. An aircraft's airborne speed must remain in a narrow range: significant speed increases are usually unachievable; slower speeds can produce stalls; and halting is impossible. This greatly constrains the actions that are available, and is further limited by the amount of fuel onboard (which is minimized to reduce operating costs).

ATFM generally has more centralized control than other forms of traffic management: In contrast, CATFM increases information sharing and distributes some elements of decision making. Finally, a significant portion of air traffic is comprised of fleets (i.e., airlines) – essentially allied pilots who are interested in cooperating for the common good of the company.

When viewed abstractly, systems developed and evaluated for CATFM could be generalized to other agent-based systems, particularly those that model people. Like many other real-world systems, the air traffic system involves a competition for limited and shared resources. The participants of this system are neither wholly cooperative (which is rarely realistic given self-interest), nor entirely competitive (which can lead to less efficient overall performance). Rather, there are two types of participants: a controlling entity, which seeks to maximize some global property such as system performance; and participating operators, which seek to maximize their own utility. The challenge is to design a robust system of constraints so that the actions of the participants work towards maximizing the desired global property. The utility functions of the participants are self-determined, may include antagonistic elements, and are generally unknowable, complicating matters. Yet, this situation occurs often not only in government-controlled systems, but also in any system with central authority, such as companies, organizational bodies, and games of many types.

We begin with a description of ATFM and related work. We describe the main features of the CATFM concept of operations and the observed operational problems it is meant to address. Our approach to developing a simulation of this concept of operations is presented, and we describe our simulation of the flight routing phase. We discuss the comparative results of different CATFM approaches and different airspace user strategies. We conclude

with an analysis of these experiments, and present our goals for future development of the simulation.

Background

Introduction to ATFM

Air traffic control (ATC), a superset of ATFM, provides safe, orderly, and efficient flow of aircraft operating within a given airspace (Nolan, 2003). Generally, an Air Traffic Service Provider (ATSP) is the authority responsible for providing air traffic management; the FAA is the ATSP for the U.S.A.'s National Airspace System (NAS). The FAA has four major types of facilities that participate in ATC. ATC towers manage the aircraft arriving, departing, and taxiing on the ground. Terminal radar approach control facilities control airspace within approximately thirty miles of a major airport. Air Route Traffic Control Centers (ARTCCs) are responsible for the remainder of controlled airspace in the NAS. There are twenty such ARTCCs in the continental United States, and each ARTCC is further subdivided into sectors. Finally, the Air Traffic Control System Command Center (ATCSCC) develops nation-wide strategic plans for traffic flow management throughout the NAS. It has final approval of all national flight restrictions and is responsible for resolving inter-facility issues. Our research has focused on ATFM at the ARTCC level, which consists mostly of "en route" traffic flying on instrument flight rules (meaning they rely on instrumentation and FAA guidance). The FAA usually assigns traffic to predefined air routes (essentially "sky highways") in order to increase the predictability of the traffic flow.

ATFM is a system-level function to manage the traffic flow based on capacity and demand. ATFM is the responsibility of a Traffic Management Unit (TMU) within each ARTCC and the ATCSCC for regional and national problems, respectively. The ATCSCC TMU develops strategic plans to ensure balanced flow throughout the NAS over a planning horizon of two to eight hours. The ARTCC TMUs develop tactical plans to manage air traffic within their local airspace over a planning horizon of up to two hours that are consistent with any relevant ATCSCC restrictions. The TMUs constantly monitor for potential conditions that could reduce airspace capacity such as adverse weather, and for excessive traffic demand that could overload a sector controller's ability to safely handle traffic (Adams, Kowitz, Milner, & Odoni, 1996). For example, a TMU may identify a Flow Constrained Area (an airspace region with a capacity-demand imbalance) due to anticipated severe convective weather. The TMU would then analyze which type of restriction should be invoked to alleviate the traffic imbalance. Since restrictions may affect adjacent centers, either directly or through ripple effects, ATCSCC approval is needed before invoking such a restriction. ATFM issues are reported during a bi-hourly planning teleconference, involving representatives from the ATCSCC, each ARTCC, and airspace users.

A variety of restrictions are available to the FAA, depending on the nature of the traffic flow problem (Sridhar, Chatterji, Grabbe, & Sheth, 2002); we describe some commonly used restrictions. A re-route procedure assigns a new route to an aircraft to avoid a problem area, such as a severe thunderstorm or congested airspace. (This is the only restriction we have implemented in our current simulation.) A Ground Delay Program (GDP) is used to delay aircraft at departure airports in order to manage the demand at an arrival airport. Flights are assigned delayed controlled departure times, thus changing their expected arrival time at the impacted airport. GDPs are implemented when capacity at an arrival airport has been reduced

for a sustained period, due to weather or excessive demand. Miles-in-Trail (MIT) restrictions enforce an increased spatial separation between aircraft transiting through some point in the airspace, but may shift traffic problems upstream. Time-based metering provides dynamic sequence and schedule advisories to controllers to reduce delays for arrival aircraft approaching capacity-constrained airports.

Airlines manage their fleet of aircraft in an Airline Operations Center (AOC). Each AOC has a coordinator that monitors the restrictions and participates in the planning teleconference to make their concerns known to the FAA. A major thrust of the CATFM concept is to increase the role of the AOCs in ATFM.

Agent-based ATFM Simulations

The Airspace Concept Evaluation System (ACES) (Sweet, Manikonda, Aronson, Roth, & Blake, 2002) is a distributed agent-based simulation of the entire NAS, including but not restricted to ATFM (Couluris, Hunter, Blake, Roth, Sweet, & Stassart, 2003). ACES uses a layered architecture to support several simulations at various levels of fidelity. Airspace participants, ranging from individuals to larger entities, are represented as agents. Given its broad coverage, ACES is able to perform cost-benefit evaluations on new concepts whose effects go beyond that of a particular element.

IMPACT (Intelligent agent-based Model for Policy Analysis of Collaborative Traffic flow management) is a swarm-based agent model of FAA agents and airline agents, used to evaluate three possible responses to capacity reductions: no advanced planning, GDPs without information sharing, and GDPs with shared airline schedules (Campbell, Cooper, Greenbaum, & Wojcik, 2000). In each scenario, the FAA agents decide whether or not to impose GDPs, based on predefined policies. The airline agents choose actions that minimize the estimated cost to their operations. As expected, their simulation measured the best performance when schedule information was shared, but found that GDPs without shared information (as occurs in today's operations) resulted in a *greater* average cost per flight than when no advanced planning occurred.

STEAM (Tambe, 1997) has been used to evaluate a collaborative system for real-time traffic synchronization (Nguyen-Duc, Briot, Drogoul, & Duong, 2003). Real-time traffic synchronization is the work of the individual sector controllers as they manage flights that run through multiple sectors. The airspace user agents do not participate in the collaboration: Only the sector controller agents and a few higher-level coordinating entities coordinate their problem-solving actions.

The Man-Machine Integrated Design and Analysis System (MIDAS) is an agent-based model of human performance when coupled with machine interfaces. MIDAS has been applied to ATFM (Corker, 1999), and emphasizes the capabilities and limitations of human cognitive ability instead of complex decision making.

Issues and Problems

Characterizing Operations and Issues through Field Observations

To characterize current problems in air traffic flow management, field observations were conducted at several operational centers (Idris, Evans, Vivona, Krozel, & Bilimoria, 2006). A diverse set of facilities was included to provide a wide scope of operational characteristics and corresponding issues, including five ARTCCs, five AOCs, and the ATCSCC. The

ARTCCs managed areas of varied geographical size with assorted weather characteristics and differing traffic patterns. The AOCs included both large and small carriers, with different operational models and customers. Finally, the ATCSCC provided a unique perspective of national air traffic flow management.

These field observations supported the development of the CATFM concept of operations in three ways. First, they made it possible to characterize the operational situations that result in air traffic flow constraints. These operational situations typically stem from two immediate causes: either from a decrease in airspace capacity (e.g., due to weather or airspace restrictions); or through an increase in demand (e.g., from pop-up traffic, overscheduling, or from traffic rerouted from another area). Second, once the flow constraint situations and their immediate causes were identified, the underlying operational issues that often lead to inefficient handling of these situations were identified. Finally, these observations provide a valuable record of *work practice*. By analyzing how the work is done, potential solutions were developed, and a corresponding agent-based model of ATFM operations was built.

Identified ATFM Issues

The primary finding from the field observations was that the current ATFM system limited the potential for collaborative problem solving. Primarily two factors cause these issues. First, the sharing of information between the FAA and airlines is limited. Thus, planning must be conducted without accurate information about the other entity's view of the current state, priorities and plans. These three elements correspond to the belief, desire and intention agent framework (Bratman, 1999). Second, the bulk of the problem solving activities falls upon the FAA, but their workload limits the solutions they can realistically pursue. We present a summary of these findings; the complete list can be found in (Idris, Vivona, Penny, Krozel, & Bilimoria, 2005).

Inaccurate Problem Assessment

Efficient management of traffic flow issues begins with an assessment of the problem. Incorrect assessments of either the demand or the capacity can lead to inaccurate problem assessments, including over- or underestimating the problem severity, missing a problem or incorrectly raising a non-existent problem. Factors that lead to inaccurate *demand* assessments include erroneous prediction of pop-up traffic, changes in departure times, flight plans or cancellations, and displacement of traffic from flow constraints elsewhere. Factors that lead to inaccurate *capacity* assessments include incorrect weather and airspace restriction predictions. These inaccuracies may lead to divergent assessments between the FAA and AOCs, resulting in inconsistent plans.

Differing Evaluations of Identified Problem

Once the traffic flow problem is identified, the FAA and the airlines regard the problem differently, for after safety, their concerns diverge. The FAA will seek to minimize the effect of the problem on the NAS and limit controller workload. The airlines are only concerned with the affect on their own flights and not the flights of competing airlines. Each airline seeks solutions that adhere to their business model, often with a goal of minimizing costs while limiting the negative effect on their customers. Moreover, different carriers will have different business models, therefore addressing cost, reliability and on-time service

differently. Thus, even with a consensus on the traffic flow problem, different entities will often prefer different solutions.

Limited Mitigations

The ARTCC and ATCSCC TMUs have a limited set of restrictions available when choosing mitigations to a traffic flow management issue. These restrictions are typically coarse-grained and are applied uniformly to all airspace users. Often, the mitigations are overly restrictive, and because they are not selective, may disproportionately impact some airspace users.

High TMU Workload

Two factors contribute to a high TMU workload when the disruptions to the NAS grow severe. First, the reliance on direct synchronous communications such as teleconferences and phone calls increases the cost of communication, decreasing both the time available for such communications and for other activities. Secondly, actions targeting individual flights (such as rerouting) greatly increase the quantity of tasks that must be performed by the TMU. As a result, TMU workload becomes a limiting factor for the possible solutions.

Limited Coordination between FAA and Airlines

Due to the problems with communication and TMU workload, coordination between the FAA and the airspace users *decreases* as problems become more severe. Unfortunately, this means there is little or no coordination exactly at the times when it is needed most. The FAA and the AOC assess, evaluate and plan independently from one another outside of the planning teleconferences run by the ATCSCC. This is exacerbated by the relative unpredictability of both parties, potentially leading to a double penalty for either: The TMU may choose unnecessary mitigations and be unprepared for the actual problem, while the AOC may independently avoid one restriction only to be impacted by another, unanticipated restriction. Moreover, due to the decrease in communication caused by a high workload, the FAA may be late in notifying all interested parties that a restriction has been removed, resulting in some parties needlessly avoiding a problem that no longer exists.

Solutions and Recommendations

CATFM Concept of Operations

The CATFM concept of operations recommends several changes to address these issues. Most changes fall under the following three categories, listed by order of increasing emphasis. First, automation must be used to reduce the workload of TMU personnel, reducing the need for the TMU planners to perform mundane tasks and lessening the cost of communication. Second, more information should be shared between the FAA and airspace users. By doing so, assessments can be made with more complete information, common assessments are possible, and actions are more predictable. Finally, and most importantly, when possible, the AOCs should be more involved in the traffic flow management process.

We summarize the four phases of the ATFM process in the CATFM Concept of Operations below; a more complete description can be found in (Idris, Vivona, Penny, Krozel, & Bilimoria, 2005).

Common Problem Identification

As described previously, ATFM problems are caused by situations where the demand for an airspace exceeds its capacity. Demand is best predicted by the airspace users who create it, whereas capacity is determined by the FAA, as it is an assessment of the FAA's ability to manage traffic in the affected area. This leads naturally to a collaborative situation where information is shared to produce a more accurate problem assessment, and to minimize the divergence of problem assessments.

Shared Impact Assessment

Various restrictions could address a given ATFM issue, each with a different impact on airline and FAA operations. By establishing a shared impact assessment, options can be evaluated more accurately and better contingency plans can be developed. Moreover, if early indications of probable TMU actions are provided, the AOCs may be able to adjust their plans to coincide with such actions, potentially reducing or eliminating the need for the proposed TMU action.

Traffic Flow Planning with AOC Input

Once a possible set of ATFM actions have been identified, along with their impact, a specific ATFM plan is instantiated to address the traffic flow problem. Instead of a planning decision being made unilaterally by the TMU (as occurs today), the AOCs can provide preferred solutions. These become additional inputs to the TMU's planning process, allowing for the accommodation of airspace user preferences when they do not violate other constraints. In addition, when the TMU workload allows it, the AOCs can suggest alternative plans that may result in an overall better solution.

Joint Plan Implementation

Once an ATFM plan with a set of actions has been chosen, it must be instantiated at the level of individual flights. In some cases, particularly with reroutes, choices must be made, such as which flights should be given the new route. When possible, the airlines should choose which of their flights are impacted by the ATFM action, according to their individual business plan. This reduces the workload of the TMU by shifting the burden of implementation to the AOC, and allows the airline to maximize their own benefit by directly choosing the most acceptable options.

Approach

We have built an initial agent-based simulation of CTFM with Brahms (Clancey, Sierhuis, Kaskiris, & Hoof, 2003). Brahms is a modeling and simulation environment for developing intelligent software agents, particularly to analyze 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*, illuminating how formal process flow descriptions relate to people's actual situated activities in the workplace (Clancey, Sachs, Sierhuis, & Hoof, 1998). 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 models not only individual agent and group behavior, but also systems and artifact behavior, as well as the interactions of people, systems, objects, and the environment. Most other multi-agent languages leave out artifacts and the interaction with the environment, making it difficult to develop a holistic model of real-world situations (Wooldridge & Jennings, 1995). 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 (Sierhuis, 2001).

A methodology for designing and simulating future work systems has been developed and used with Brahms (Clancey, Sierhuis, Seah, Buckley, Reynolds, Hall, & Scott, 2007). The process begins with detailed observations of work practice, which is used to build a model of current operations. After model validation, a new concept of operations is developed, and a simulation of the future work system is created using validated components of the model of current operations whenever possible. After testing the concept in implementation, the process repeats. We have adapted this methodology to our circumstances, taking advantage of the pre-existing CATFM concept of operations and work practice observations. We are developing the model iteratively, building successively more accurate models from increasingly detailed sources of information. At every stage, we evaluate the concept of operations based on the findings of our simulation, modify the concept accordingly, and then increase model fidelity in the next stage.

So far we have built a rudimentary model of ATFM using second-hand sources of information such as the work practice observations described previously, other ATFM literature, and the concept of operations itself. In the next stage, we will interview subject matter experts and incorporate their conception of work practice into the model. This will allow us to fill in details not discernable from the recorded observations of work practice. To validate the model at this stage, historical situations will be simulated and the results will be compared with the historical outcomes. Likewise, historical data may also be used to infer behavior, either by intuition or through data mining techniques. In the third stage, we will perform new observations of work practice, enabling us to build a detailed model at the level of individual (rather than organizational) participants in the ATFM process. The model at this stage can also be validated by comparing the simulated behavior to the behavior observed in the actual system. Subsequent evaluation of the concept will require human subjects to participate in the CATFM process, with humans and agent proxies participating in a human-in-the-loop simulation.

Initial ATFM Simulation Design

We have created a simplified model of a subset of ATFM, including only the Joint Plan Implementation phase (see earlier section of same name) when flights are assigned routes. In order to simplify this selection process, we have redefined capacity to be a property of a route, rather than a sector, and assumed that the routes are independent. Route capacity, flight schedules and agent strategies are static throughout the simulation. In contrast, route demand changes dynamically throughout the simulation as the agents choose routes. We do not model runway constraints or temporal ordering, treating all flights as if they have the same departure time. Our simulation only deals with pre-flight planning and does not simulate the flights themselves.

Figure 1 provides an overview of our current agent architecture. We have built our initial model at the organizational level, with each organization (i.e., TMUs and AOCs) modeled as single agents. Each agent (TMU or AOC) has different responsibilities, with route selection performed by either the TMU agent or the AOC agents (see below). The AOC agents provide the TMU with their flight schedules and the value of each flight. The TMU agent informs the AOC agents of the current status of the airspace by aggregating the current demand on a given route, comparing this with the capacity, and broadcasting the route status (under capacity, at capacity, or oversubscribed) to the AOC agents. In the initial simulation, the TMU does not reroute flights or choose among AOC requests: It approves them all when the route is at or below capacity, and denies all requests when demand exceeds capacity (thus leaving the route unused). To be consistent with U.S.A. law against anti-competitive practices, no communication occurs between AOC agents in order to prevent coalitions or other AOC-AOC negotiations. We do not model communication issues, treating them as reliable, instantaneous, and clear.

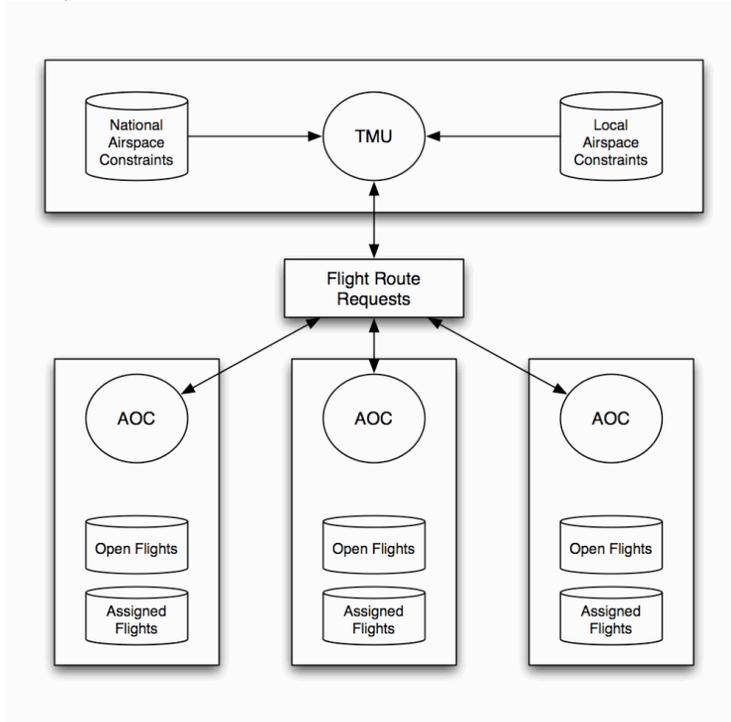


Figure 1. Agent architecture.

For each origin-destination airport pair, we created three routes arbitrarily: a direct route and two alternate routes, 1.25 and 1.5 times the length of the direct route. The capacities of these routes vary, with typically the direct route having insufficient capacity for all scheduled traffic. Our fundamental question is: how will the CATFM concept perform in this simplified model? In order to answer this question, we created four ATFM approaches:

- **Blue Sky:** All capacities are infinite, so every flight takes the direct route. This is not a realistic approach but provides an upper bound on performance that we use as a baseline.
- **Current Operations:** The TMU agent makes the route selection, putting flights on the best available routes (i.e., under capacity routes) in a random order without inspecting

the flight value. This approach is closest to the current operations where the FAA makes route assignments with little input from the airlines.

- **Global Optimum:** The TMU agent makes the route selection as in the Current Operations approach, but does so in order of greatest flight value. This greedy algorithm produces the best overall system performance, according to our metrics, but may give preferential route assignments to one airline over another due to differences in flight value distribution.
- **Airline Planning:** The AOC agents make the route selections, with each agent initially requesting the best route for every flight regardless of the strategy used. After the TMU agent broadcasts the status of all routes, the AOC agent may independently choose a new route for each flight. The process repeats iteratively (six iterations unless where noted otherwise) until the time for planning is exhausted. Within a simulation run, a given AOC agent will use the same strategy on each iteration (i.e., no changes in strategy during a run). We used the following simplified strategies:
 - **Aggressive:** An AOC agent with the Aggressive strategy will always request the best route for every flight at each iteration, regardless of the situation.
 - **Moderate:** An AOC agent with the Moderate strategy will request the next best route for some of its flights when faced with an overcapacity situation, repeating the prior request for the other flights.
 - **Conservative:** An AOC agent with the Conservative strategy will request the *worst* route for some of its flights when faced with an overcapacity situation, repeating the prior request for the other flights. The assumption is that the worst route is the least likely to fill up, so the conservative AOC agent attempts to forgo a chance at a better route assignment in exchange for a greater likelihood of finding an available route.

All approaches except Current Operations are deterministic.

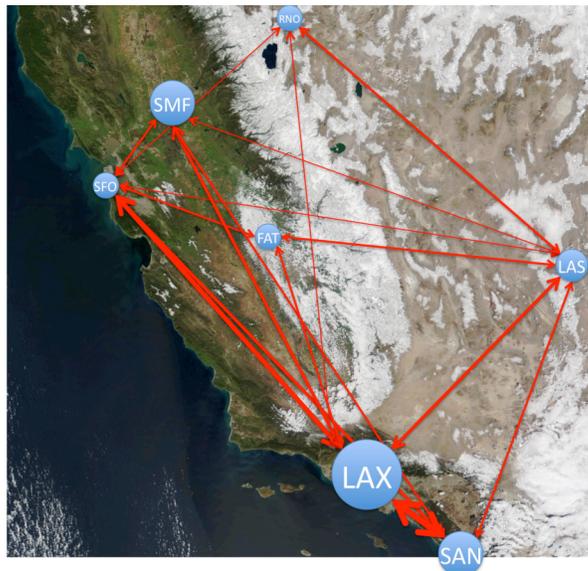


Figure 2. *Local traffic scenario involving seven airports.*

Experiment on a Local Traffic Scenario

We created a local traffic scenario (see Figure 2) that corresponds to traffic generated by three major carriers among several airports in the southwest of the U.S.A. The schedules and aircraft types were chosen based on our observations of the flight schedules of these carriers. Information on connecting crew, passengers, and route capacities were not available, however, so we used our best judgment based on nominal conditions, expected passenger behavior and operational patterns. In all cases, sufficient aggregate capacity was available among the three routes such that every flight could have *some* route assignment.

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

- p_c = passengers with connecting flights
- p_u = passengers without connecting flights
- c_c = onboard crew members with a connecting flight
- t_a = the actual flight time of F , in minutes
- t_o = 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 of the flight to the airline. We define v_F , the flight value of F , as

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

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

$$d_F = t_a - t_o \quad (2)$$

When F is not assigned a route, we assume a standard sixty minutes of delay in a later stage that we do not simulate. Traffic demand naturally rises and falls throughout the day, so we assume that the level of demand falls significantly after our simulation ends. Other factors may also cause delays in practice but are not part of our model.

Finally, we seek to measure in our experiments the total passenger delay incurred by flight F , either through an immediate delay or through missed connections. We assume that when a passenger with a connecting flight is delayed, on average, that passenger will experience an additional two-hour delay. When connecting crew members are delayed, their personal delay is not counted (since they are not considered passengers in our simulation), but they are likely to delay the departure of 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 hours of passenger delay. Combining this with the above formulae, we calculate the total incurred passenger delay incurred by flight F , d_T , in minutes, as

$$\begin{aligned} d_T &= (p_u * d_F) + (p_c * d_F) + 120p_c + 300c_c \text{ when } d_F > 0 \\ d_T &= 0 \text{ when } d_F = 0 \end{aligned} \quad (3)$$

We ran the experiments once for each deterministic approach and fifty times for the randomized Current Operations approach, yielding some surprising results (Wolfe, Jarvis, Enomoto, & Sierhuis, 2007). The Airline Planning approach is highly sensitive to the

strategies employed by the AOC agents and often performs poorly. Figure 3 shows an example with several strategies, where the light shaded bars indicate delay incurred by selecting a longer route, and the dark shaded bars indicate delay from failing to get an approved route assignment. Further examination of specific trials showed that the Aggressive strategy is disruptive to the system as a whole by pushing demand beyond capacity on the best routes. However, the best performing combination of airline strategies outperformed the Current Operations approach (see Figure 4), indicating the potential for improvement under the CATFM concept. The number of planning cycles can also affect solution quality in the Airline Planning approach, as shown in Figure 5.

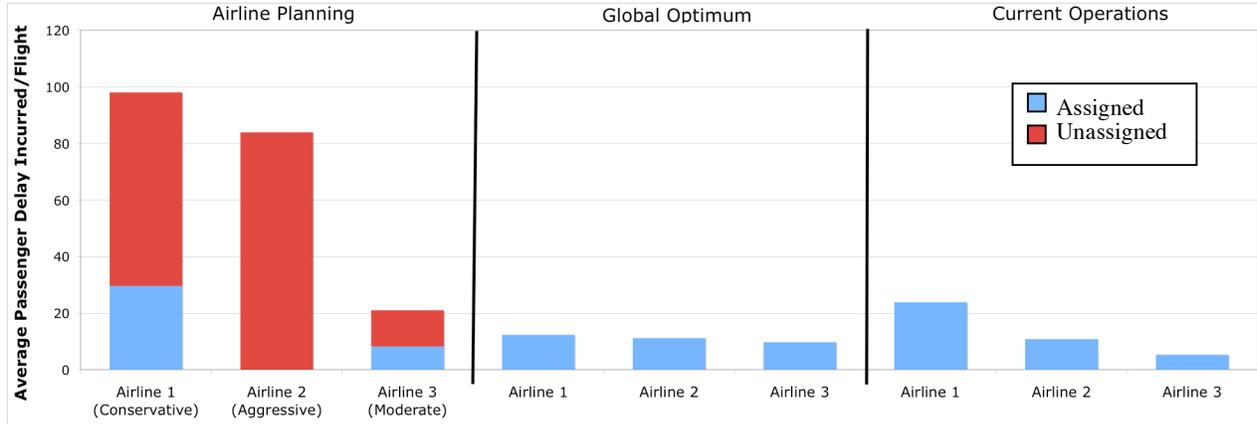


Figure 3. Comparing ATFM approaches on the local scenario.

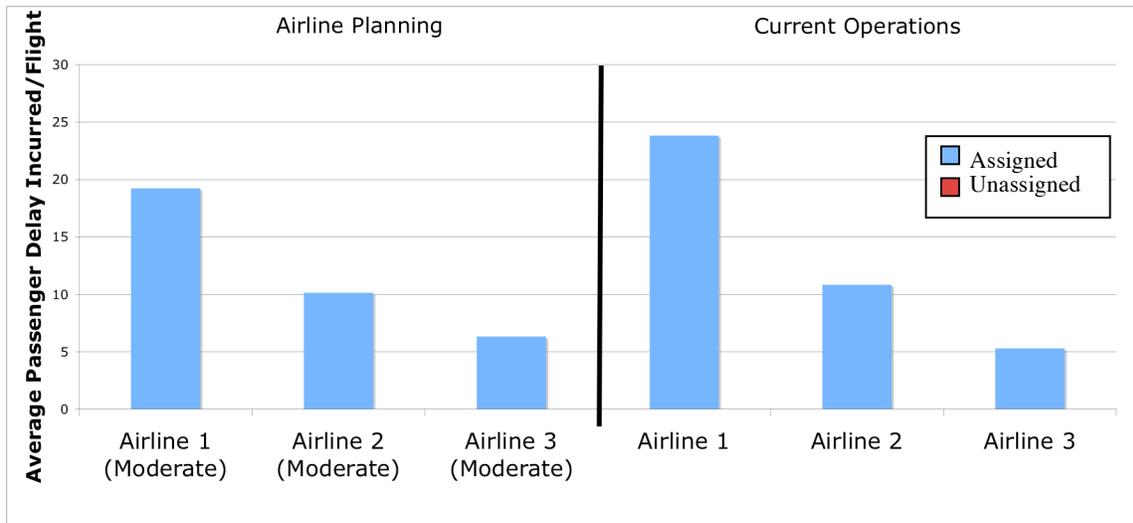


Figure 4. Best airline planning combination compared with Current Operations approach.

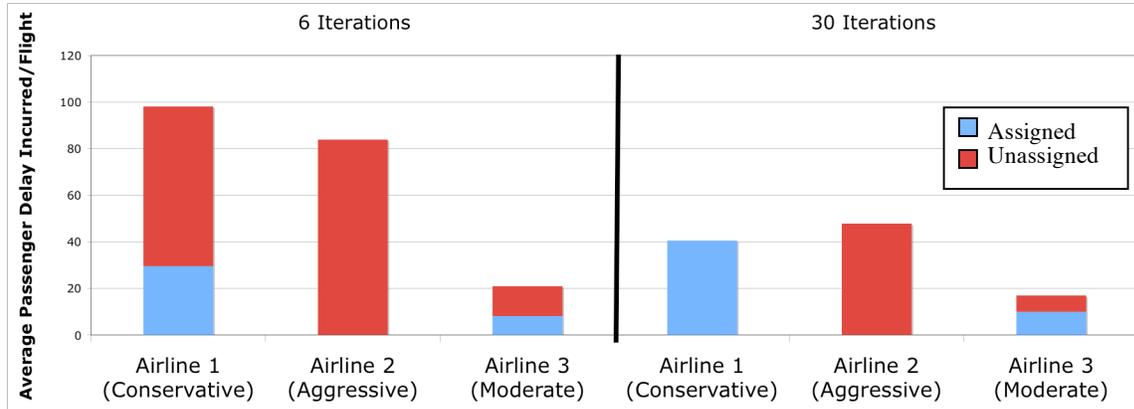


Figure 5. Effect of additional planning cycles with Airline Planning approach.

Single Origin-Destination Experiment

In our previous experiment, a given AOC agent would use the same strategy on all origin-destination pairs, regardless of the situation. In reality, an airline is likely to use several strategies, matching them to the situation at hand. Since we aggregated the results over the origin-destination pairs, we could see how a strategy performed overall but could not isolate the specific situations where it performed well or poorly. We also wanted to evaluate new approaches that could address concerns that arose from our previous set of experiments, leading to the following additions:

- **Mixed:** This combines the Airline Planning and Optimal approaches. The airlines schedule their flights as before in the Airline Planning approach. Once the planning phase is over, however, the TMU agent will assign any unassigned flights using the Optimal approach. This ensures that any unused capacity will be utilized by flights for which the AOC agents failed to choose an acceptable route.
- **Equitable:** This is a variant of the Optimal approach. Each AOC agent gives a ranking of their flights but does not supply flight values. The TMU agent gives top priority to first-ranked flights, followed by second-ranked flights, and so on. This gives each airline an equal share of each route’s capacity, regardless of the value of their flights.

We created three scenarios with the same origin-destination, with one primary route and two alternates as defined previously. In all three scenarios we had three AOC agents, each with four flights to schedule. The scenarios varied in the amount of capacity available:

- **Demand<Capacity:** each route can accommodate five flights.
- **Demand=Capacity:** each route can accommodate four flights.
- **Demand>Capacity:** each route can accommodate only three flights.

Therefore, all flights could be assigned a route on the Demand<Capacity and the Demand=Capacity scenarios, but this was not possible in the Demand>Capacity scenario.

We ran each scenario with all combinations of the three strategies for the three AOC agents, using both the Airline Planning and Mixed approaches, resulting in twenty-seven runs for each. Figure 6 and Figure 7 show the average performance (across all agents and competitor strategy combinations) for each strategy. Table 1 and Table 2 compare strategy alternatives by measuring whether an agent would do as well or better with a different strategy in the given situation, while keeping the competitor strategies constant. For instance, in the Demand<Capacity scenario under the Mixed approach (Table 2), an agent using the

Aggressive (A) strategy would have performed as well or better with the Moderate (M) strategy in only 4% of the simulated situations— indicating that the Aggressive strategy was the better choice.

Several patterns emerge from this analysis. The Aggressive strategy is a poor choice when using the Airline Planning approach, consistent with earlier findings, because its insistence on the best route makes that route unusable, potentially leaving its flights unassigned. In contrast, the Aggressive strategy is a good choice when using the Mixed approach with adequate overall capacity. In such cases, the Aggressive strategy will either succeed in putting all of its flights onto the best route, or it will prevent all other airlines from using the best route. In the latter case, none of the Aggressive airline’s flights will be scheduled, and the best route will be completely available when the TMU assigns the remaining flights, leading to a greater share of the best route. However, when there is not sufficient capacity, this strategy performs poorly because not all of its flights will be assigned.

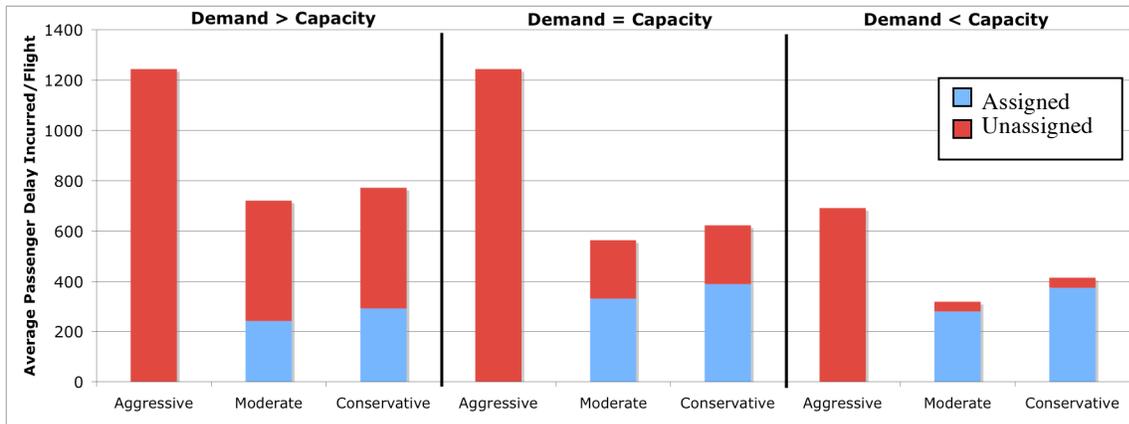


Figure 6. Strategy performance averaged over all agents with Airline Planning approach.

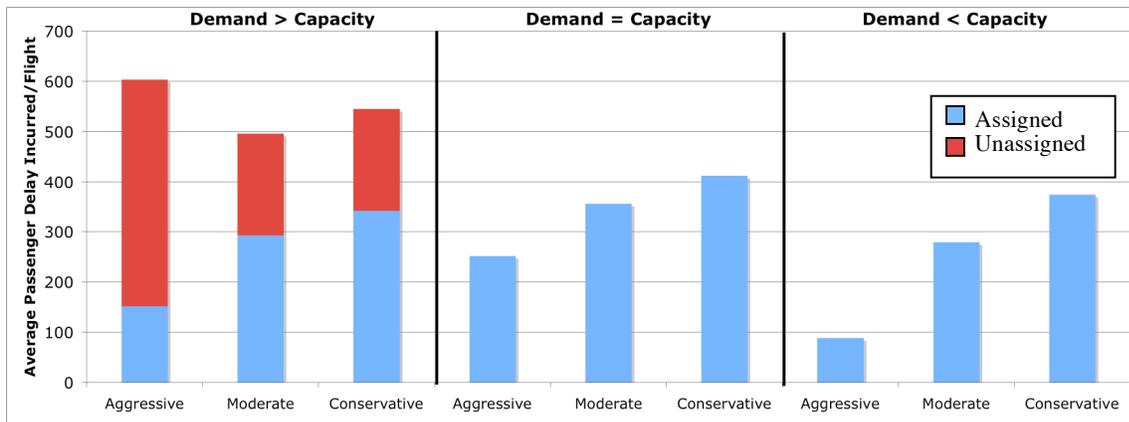


Figure 7. Strategy performance averaged over all agents with Mixed approach.

Demand > Capacity	Chosen Strategy			Demand = Capacity	Chosen Strategy			Demand < Capacity	Chosen Strategy					
	A	M	C		A	M	C		A	M	C			
Alternate Strategy	A	-	0%	0%	Alternate Strategy	A	-	0%	0%	Alternate Strategy	A	-	44%	44%
	M	100%	-	78%		M	100%	-	89%		M	56%	-	100%
	C	100%	22%	-		C	100%	11%	-		C	56%	0%	-

Table 1. Airline Planning approach: Cases equal or improved with alternate strategy.

Demand > Capacity	Chosen Strategy			Demand = Capacity	Chosen Strategy			Demand < Capacity	Chosen Strategy					
	A	M	C		A	M	C		A	M	C			
Alternate Strategy	A	-	19%	33%	Alternate Strategy	A	-	85%	100%	Alternate Strategy	A	-	100%	100%
	M	81%	-	78%		M	26%	-	93%		M	4%	-	100%
	C	67%	22%	-		C	11%	7%	-		C	0%	0%	-

Table 2. Mixed approach: Cases equal or improved with alternate strategy.

	Airline 1	Airline 2	Airline 3	Total
Current Operations	3552	4332	2939	10823
Optimal	3314	2806	3300	9420
Equitable	2969	3407	3073	9449

Table 3. Total Incurred Passenger Delay for three airlines with forty flights each.

As shown by the rows of Table 1 and Table 2, the best strategy cannot be determined only by the scenario and approach (with the sole exception of the Aggressive strategy in the Demand<Capacity scenario with the Mixed Approach). This is because the performance of a strategy is affected by the competitors' strategies: in particular, each strategy performed worse when a competitor used the same strategy. Therefore it was often preferable to use a unique but generally less attractive strategy than one used by a competitor.

Finally, we created a larger scenario with primary and secondary routes defined as before, but each with a capacity of forty flights, and three airlines with forty flights each. Table 3 shows the results of experiments on this scenario in terms of the total incurred passenger delay metric. In this case, the Equitable approach performed nearly as well as the Optimal approach; it is worth noting that the distributions of flight values were comparable among the three airlines.

Conclusion

We have described the design and methodology of a multi-agent simulation of ATFM, as well as experimental findings. At this time, our simulation is a coarse-grained model of operations, with agents corresponding to participating entities (i.e., TMUs and AOCs) rather than persons. Since we simplified other components that were not essential to the problem, actual performance in implementation may differ, but should produce similar conclusions under identical conditions, strategies and policy.

We evaluated several approaches to ATFM, and for the Airline Planning and Mixed approaches, also evaluated several simple route selection strategies. Of these, the Moderate strategy is intuitively the most appealing, and had the best overall performance in our experiments. In contrast, the Conservative strategy did not perform as well, but was usually preferable when it was different than all competitors' strategies. This theme was repeated throughout our experimental results; in nearly every case, the best strategy could not be chosen independently, as it was dependent on the strategies used by the other AOC agents. Finally, the Aggressive strategy worked very well with the Mixed approach when there was adequate capacity, casting doubt on the suitability of the Mixed approach. The Aggressive strategy also did well when the other AOC agents removed their flights from the best route, thus accommodating the aggressive AOC.

In our evaluation of the CATFM concept, we observed that nearly all the approaches that utilized our flight value metric (Equation 3) yielded better results than the Current Operations approach. This supports the claim that utilizing airspace user preferences in ATFM should lead to better solutions. However, this was not the case in all of our experimental results; certain combinations of strategies with the Airline Planning approach produced unacceptably poor results. Moreover, based on current experiments, we did not observe any indication that increasing AOC involvement would reduce FAA workload. In the Optimal and Equitable approaches, the TMU agent continued to perform route selection, and with additional criteria, so this represents an increase in workload. In the Airline Planning approach, the TMU did not perform route selection but the results were often unacceptable; in the Mixed approach, the results were good, but often the TMU would still make many route selections and inadvertently rewarded aggressive behavior. Therefore, automation is most likely the key to reducing FAA workload. Finally, the AOC agents usually found better solutions when more planning cycles were available. This puts an emphasis on the earlier stages of the CATFM process, which we did not simulate – the earlier situational information is available, the better the likely solution.

In the end, the challenge of refining the CATFM concept will *not* be designing effective AOC agent strategies, as they will be determined by the airlines rather than the system designers. Each airline is likely to have a somewhat different strategy, geared towards their private business model and influenced by the people executing it. Nor is it reasonable to assume that these strategies would necessarily be optimal in all cases. Rather, the challenge is to design a *system* that rewards behavior yielding desirable system performance. In game-theoretical terms, this amounts to redesigning the game itself, rather than the player strategies. In our experiments, the Airline Planning approach was vulnerable to aggressive AOC agents; likewise, the Mixed approach often rewarded the Aggressive strategy. The Optimal approach is unlikely to be deployable in practice, as it would be difficult to create a single objective utility function (flight value in our experiments) over all airlines. Based on our experiments, the Equitable approach is the most promising, as it produced results on par with the Optimal approach (when airlines had comparable flights), but did so without relying on a universal flight evaluation.

Future Research Directions

We have completed the initial stage of development and will continue to expand the CATFM model. We have begun work on the next stage, expanding our model to capture the breadth of the CATFM concept of operations, covering all phases. Our current study

simulated the instantiation of the ATFM plan (namely the selection of routes), which was necessary to evaluate the result of the process; however, as earlier phases produce inputs to later phases, it may be that the earlier phases have the greatest operational impact.

In addition to broader scope, a higher degree of fidelity would support stronger claims about the CATFM concept of operations. A more sophisticated flight model would eliminate many simplifying assumptions, such as simplified schedules, and route capacities in lieu of sector capacities. Modeling organizational roles and concentrating on interactions at the level of individual people would reveal the complexity of the proposed work practice and lead to more accurate characterizations of workload. Interviews with subject matter experts, case studies, and additional observations of work practice will yield insight as to how these processes work today.

The results from our initial experiments can be used to guide refinements to the concept of operations and develop policies that are more likely to be successful. Further experimentation with the Equitable approach in a wider array of situations is needed to evaluate its suitability. Additionally, more complex ATFM approaches and airline strategies may yield better overall solutions. Identifying likely airline strategies is of great importance, but difficult, due to their proprietary nature. Since the situations we are simulating are characteristic of *future* operations, rather than today's operations, airlines may not have developed appropriate strategies, and if they have, they may not be willing to share them.

Building a model of *future* operations is difficult at any stage of development. Our approach has been to build and validate a model of current operations, and then to modify that model to fit the future concept. Even validating the current model is a challenge, given the complexity of operations. Modifying a model of current operations to yield a model of future operations introduces uncertainty. We have dealt with this by simulating a variety of possible actions, essentially modeling several possibilities. Game theory can be utilized to develop likely strategies and to analyze properties of the system as a whole. Approaches to traffic management problems in other domains may translate to ATFM, and vice versa.

References

- Adams, M., Kolitz, S., Milner, J., & Odoni, A. (1996). Evolutionary Concepts for Decentralized Air Traffic Flow Management. *Air Traffic Control Quarterly*, 4(4), 281-306.
- Bratman, M. E. (1999). *Intention, Plans, and Practical Reason*. Chicago, Illinois: University of Chicago Press.
- Clancey, W. J., Sachs, P., Sierhuis, M., & Hoof, R. v. (1998). Brahms: Simulating practice for work systems design. *International Journal on Human-Computer Studies*, 49, 831-865.
- Clancey, W. J., Sierhuis, M., Kaskiris, C., & Hoof, R. v. (2003, May). *Advantages of Brahms for Specifying and Implementing a Multiagent Human-Robotic Exploration System*. Paper presented at the 16th International Florida Artificial Intelligence Research Society (FLAIRS) Conference, St. Augustine, Florida.
- Clancey, W. J., Sierhuis, M., Seah, C., Buckley, C., Reynolds, F., Hall, T., & Scott, M. (2007, October). *Multi-Agent Simulation to Implementation: A Practical Engineering Methodology for Designing Space Flight Operations*. Paper presented at the Eighth Annual

A Multi-Agent Simulation of Collaborative Air Traffic Flow Management

International Workshop "Engineering Societies in the Agents World" (ESAW 07), Athens, Greece.

Corker, K. M. (1999, December). *Human Performance Simulation in the Analysis of Advanced Air Traffic Management*. Paper presented at the 1999 Winter Simulation Conference, Phoenix, Arizona.

Couluris, G. J., Hunter, C. G., Blake, M., Roth, K., Sweet, D. N., Stassart, P. A., (2003, August). *National Airspace System Simulation Capturing the Interactions of Air Traffic Management and Flight Trajectories*. Paper presented at the American Institute of Aeronautics and Astronautics (AIAA) Guidance, Navigation, and Control (GNC) Conference, Austin, Texas.

Hardin, G. (1968). The Tragedy of the Commons. *Science*, 162(3859), 1243-1248.

Idris, H., Evans, A., Vivona, R., Krozel, J., & Bilimoria, K. (2006, September). *Field Observations of Interactions Between Traffic Flow Management and Airline Operations*. Paper presented at the American Institute of Aeronautics and Astronautics (AIAA) 6th Aviation, Technology, Integration, and Operations Conference (ATIO), Wichita, Kansas.

Idris, H., Vivona, R., Penny, S., Krozel, J., & Bilimoria, K. (2005, September). *Operational Concept for Collaborative Traffic Flow Management based on Field Observations*. Paper presented at the American Institute of Aeronautics and Astronautics (AIAA) 5th Aviation, Technology, Integration, and Operations Conference (ATIO), Arlington, Virginia.

Campbell, K. C., Cooper, W. W., Greenbaum, D. P., & Wojcik, L. A. (2000, June). *Modeling Distributed Human Decision-Making in Traffic Flow Management Operations*. Paper presented at the 3rd USA/Europe Air Traffic Management R&D Seminar, Napoli, Italy.

Mukherjee, A., Grabbe, S., & Sridhar, B. (2008). *Alleviating Airspace Restriction through Strategic Control*. Paper presented at the American Institute of Aeronautics and Astronautics (AIAA) Guidance, Navigation, and Control Conference, Honolulu, Hawaii.

Nguyen-Duc, M., Briot, J.-P., Drogoul, A., & Duong, V. (2003, October). *An Application of Multi-Agent Coordination Techniques in Air Traffic Management*. Paper presented at the 2003 Institute of Electrical & Electronic Engineers/ Web Intelligence Consortium (IEEE/WIC) International Conference on Intelligent Agent Technology (IAT 2003), Halifax, Canada.

Nolan, M. S. (2003). *Fundamentals of Air Traffic Control* (4th ed.) Pacific Grove, California: Brookes/Cole Publishing.

Odoni, A. (1987). The Flow Management Problem in Air Traffic Control In *Flow Control of Congested Networks* (pp. 269-288). Berlin, Germany: Springer-Verlag.

A Multi-Agent Simulation of Collaborative Air Traffic Flow Management

Pearce, R. A. (2006). The Next Generation Air Transportation System: Transformation Starts Now. *Journal of Air Traffic Control*, 7-10.

Sierhuis, M. (2001). *Modeling and Simulating Work Practice, BRAHMS: A multiagent modeling and simulation language for work system analysis and design* (Vol. 2001-10). Amsterdam, the Netherlands: University of Amsterdam.

Sridhar, B., Chatterji, G., Grabbe, S., & Sheth, K. (2002, August). *Integration of Traffic Flow Management Decisions*. Paper presented at the American Institute of Aeronautics and Astronautics (AIAA) Guidance, Navigation, and Control Conference, Monterey, California.

Sweet, D. N., Manikonda, V., Aronson, J. S., Roth, K., & Blake, M. (2002, August). *Fast-Time Simulation System for Analysis of Advanced Air Transportation Concepts*. Paper presented at the American Institute of Aeronautics and Astronautics (AIAA) Modeling and Simulation Technologies Conference and Exhibit, Monterey, California.

Tambe, M. (1997, July). *Agent architectures for flexible, practical teamwork*. Paper presented at the American Association for Artificial Intelligence Conference (AAAI-2007), Providence, Rhode Island.

Wambsganss, M. (1996). Collaborative Decision Making Through Dynamic Information Transfer. *Air Traffic Control Quarterly*, 4, 107-123.

Waslander, S. L., Raffard, R. L., & Tomlin, C. J. (2008). Market-Based Air Traffic Flow Control with Competing Airlines. *Journal of Guidance, Control and Dynamics*, 31(1), 148-161.

Wolfe, S. R., Jarvis, P. A., Enomoto, F. Y., & Sierhuis, M. (2007, November). *Comparing Route Selection Strategies in Collaborative Traffic Flow Management*. Paper presented at the 2007 Institute of Electrical & Electronic Engineers/ Web Intelligence Consortium/ Association for Computing Machinery (IEEE/WIC/ACM) International Conference on Intelligent Agent Technology (IAT 2007), Fremont, California.

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

Additional Reading

The Brahms simulation environment has its own language (Hoof & Sierhuis, 2007), which is similar but distinct from other belief, desire, and intent frameworks (Sierhuis, 2007). This representation has been developed to support the simulation of work practice (Sierhuis & Clancey, 2002), a major application of Brahms technology. The theoretical basis of Brahms is related to that of situated cognition (Clancey, 2002). The Brahms tool set, simulation environment and additional information are publicly available from the Brahms website (Agent iSolutions).

Agent based modeling and simulation and agent-based techniques have been applied to various aspects of AOC operations. A simulation of the United Airlines AOC has been developed (Pujet, Feron, & Rakhit, 1998), where each AOC employee is modeled as a multi-class queueing server. This model was used to track task execution information, namely which entities performed which task at any given point in time, with the goal of supporting timely decision making. Castro and Oliveira have developed a multi-agent system to handle disruptions in operations by reallocating crew (Castro & Oliveira, 2007). Various agents compete using different methods problem-solving methods to find the best solution; in simulation, this approach produced better solutions than current human operators.

Agent-based solutions have been proposed to solve other areas of ATFM. Tumer and Agogino have developed a multi-agent algorithm for ATFM (Tumer & Agogino, 2007). They use a Monte-Carlo simulation to estimate the congestion within the NAS, based on agents' actions to speed up or slow down traffic. These agents use reinforcement learning to set the separation between airplanes in order to manage the congestion. OASIS is an agent-based system developed to maximize airport arrival throughput by managing aircraft arrival and runway utilization (Ljunberg & Lucas, 1992). Various functions of ATC Tower operations are managed by agents in OASIS, and are implemented in the Procedural Reasoning System (Ingrand, Georgeff, & Rao, 1992). Jonker, Meyer, and Dignum have also advocate the use of multi-agent systems in the ATC Tower operations (Jonker, Meyer, & Dignum, 2005). They describe a market-based control mechanism, and analyze its usage from a game-theoretical perspective.

Agent-based modeling and simulation has also been used to study the effect of increased volume and independent choice in other forms of traffic. A simulation of projected traffic in the seaport of Rotterdam estimated the effect of increased traffic in terms of delay (Ruit, Schuylenburg, & Ottjes, 1995). Automobile traffic has been simulated fairly extensively; of particular relevance to this book chapter are those focused on route selection. Klügl and Bazzan examined how individual drivers could learn to prefer certain routes and how forecasts of traffic influenced this ability (Klügl & Bazzan, 2004). Interestingly, their study showed that the best overall system performance was achieved when most, but not all, drivers had access to these traffic forecasts. Stark et al. (Stark, Helbing, Schönhof, & Holyst, 2006) investigated how cooperative strategies could be learned in a route selection context without any communication between drivers.

Several other relevant ATFM simulation environments are not agent-based. The Future ATM Concepts Evaluation Tool (FACET) (Bilimoria, Sridhar, Chatterji, Sheth, & Grabbe, 2000) is a NASA-developed tool for simulating air traffic flow that has been integrated into Flight Explorer, a commercial product used by nearly all major U.S. airlines. 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. The Center-TRACON Automation System (CTAS) (Erzberger, 1994) is another NASA-developed simulation system, with a single ARTCC focus and a greater emphasis on human in the loop simulations. The Traffic Management Advisor, one of the CTAS suite of tools, is particularly relevant from an ATFM perspective, and has been extended to coordinate among multiple ARTCCs in the McTMA system (Hoang, 2004). The Linking Existing On Ground, Arrival and Departure project (LEONARDO) evaluated the feasibility of implementing Collaborative Decision Making (CDM) in airport processes, both through simulation and a limited deployments (European Commission, 2004). LEONARDO integrated decision support tools to promote information sharing among airport stakeholders, providing them with early and reliable planning updates. SKATE (Skills, Knowledge, and Attitudes for Teamwork), is a model for teamwork measurement developed and used in real-time simulations to validate the use of LEONARDO for CDM (EUROCONTROL, 2004).

The CATFM concept of operations has to the potential to enhance the Collaborative Decision-Making (CDM) initiative (Ball, Hoffman, Chen, & Vossen, 2000; Federal Aviation Administration), a joint government and industry effort was established in the mid-1990s to enhance the interaction and collaboration between the ATSP and the users of airspace. CDM deals with improvement of ATFM through better information exchange among the participants of the aviation community. The goal of CDM is to create solutions for better utilization of airspace resources through technological and procedural solutions for traffic management problems that are encountered in the NAS, without compromising safety. The CDM group consists of several sub-groups, e.g., flow evaluation, future concepts, ground delay program enhancements, weather evaluation, etc., which deal with various aspects of the air traffic flow management problem. Several automation decision support tools have emerged as a result of the CDM effort over the years, including the Flight Schedule Monitor (Metron Aviation, 2006a) for managing arrival/departure times, the Collaborative Convective Forecast Product (National Oceanic and Atmospheric Administration, 2007) for a common assessment of convective weather, and the Post Operations Evaluation Tool (Metron Aviation, 2006b) for analysis support of NAS operations. Preliminary evaluation of CDM initiatives on elements such as GDP is promising (Ball, Hoffman, Knorr, Wetherly, & Wambsganss, 2001).

The Future Concepts Team is a sub-group of the CDM initiative. Over the past few years, the FCT group has focused their effort on future collaboration between the service provider and the airspace users to improve efficiency of operations in the NAS. The two main areas of interest are the Integrated Collaborative Routing (ICR) (Usmani, 2005) and the System Enhancements for Versatile Electronic Negotiation (SEVEN) (Gaertner, Klopfenstein, & Wilmouth, 2007). The ICR effort is geared towards better incorporation of airspace users' preferences for rerouting during events that cause congestion and weather related delays. The SEVEN concept is a longer-term initiative which aims to enhance the collaboration among the participants to a much higher level than what exists today through use of electronic data exchange and to explore the roles and responsibilities of participants, along with identification of associated issues and concerns. This enhanced collaboration encompasses all elements of the Flow Constrained Areas (for establishing areas of impacted traffic), the Ground Delay Programs and Airspace Flow Programs (for managing traffic during bad weather conditions) and Playbook routes (for specific rerouting strategies). The premise for

Concept SEVEN is for the airspace users to provide prioritized flight lists and enabling them to update their options as the constraining events unfold.

Other concepts of operations have elements that are similar to the CATFM concept of operations. The Concept of Operations for the Next Generation Air Transportation System (Joint Planning and Development Office, 2007) defines how the air transportation system shall operate in the year 2025, forming a technological baseline to help stimulate the development of policy. The International Civil Aviation Organization has also developed requirements for an operational concept in 2025 (International Civil Aviation Organization, 2003), emphasizing collaborative decision making. It also provides a comprehensive view of operations, including airspace design, airport operations and collision avoidance, and describes potential benefits and a possible adoption strategy.

The FAA has developed useful training materials that explains terms, techniques, and programs associated with traffic flow management in the NAS (Federal Aviation Administration, 2007). Operational details of ATFM, including the ATFM roles and duties at the ATCSCC, ATFM tools, flight restriction guidelines, and overviews of the traffic patterns within each ARTCC are available from the FAA (Federal Aviation Administration, 2006). Finally, the Airline Handbook (Air Transport Association of America, 2007) provides a brief history of aviation and an overview of important aviation topics, including: the principles of flight, deregulation, the structure of the industry, airline economics, airports, air traffic control, safety, security and the environment, and a glossary of commonly used aviation terms.

Additional Reading References

Agent iSolutions. *Brahms* [Electronic Version]. Retrieved December 15, 2007, from <http://www.agentisolutions.com/>

Air Transport Association of America. (2007). *The Airline Handbook*. Washington, D.C.: ATA Publications.

Ball, M., Hoffman, R., Chen, C.-Y., & Vossen, T. (2000). Collaborative Decision Making in Air Traffic Management: Current and Future Research Directions. In L. Bianco, P. Dell'Olmo & A. R. Odoni (Eds.), *New Concepts and Methods in Air Traffic Management* (pp. 17-30). New York, New York, USA: Springer.

Ball, M. O., Hoffman, R. L., Knorr, D., Wetherly, J., & Wambsganss, M. (2001). Assessing the Benefits of Collaborative Decision Making in Air Traffic Management. In G. L. Donohue & A. G. Zellweger (Eds.), *Air Transportation Systems Engineering, Progress in Astronautics and Aeronautics* (Vol. 193, pp. 239). Reston, Virginia: American Institute of Aeronautics and Astronautics.

Bilimoria, K., Sridhar, B., Chatterji, G., Sheth, K., & Grabbe, S. (2001). FACET: Future ATM Concepts Evaluation Tool. *Air Traffic Control Quarterly*, 9(1), 1-20.

Castro, A. J. M., & Oliveira, E. (2007, November). *Using Specialized Agents in a Distributed MAS to Solve Airline Operations Problems: a Case Study*. Paper presented at the Institute of

A Multi-Agent Simulation of Collaborative Air Traffic Flow Management

Electrical & Electronic Engineers/ Web Intelligence Consortium/ Association for Computing Machinery (IEEE/WIC/ACM) International Conference on Intelligent Agent Technology (IAT 2007), Fremont, CA.

Clancey, W. J. (2002). Simulating activities: Relating motives, deliberation, and attentive coordination. *Cognitive Systems Research: Special Issue on Situated and Embodied Cognition*, 3(3), 471-499.

Erzberger, H. (1994, July). *Center-TRACON Automation System (CTAS)*. Paper presented at the Capacity Technology Subcommittee, FAA Research and Development Advisory Committee, Washington, D.C.

EUROCONTROL. (2004). *A Measure to Assess the Impact of Automation on Teamwork* [Electronic Version]. Retrieved December 15, 2007, from <http://www.eurocontrol.int/humanfactors/gallery/content/public/docs/DELIVERABLES/HF48-HRS-HSP-005-REP-07%20Released-withsig.pdf>

European Commission. (2004). *LEONARDO Final Report* [Electronic Version]. Retrieved December 15, 2007, from http://ec.europa.eu/transport/air_portal/research/doc/rtd_5_leonardo.pdf

Federal Aviation Administration. (2007). *Collaborative Decision Making* [Electronic Version]. Retrieved December 15, 2007, from <http://cdm.fly.faa.gov/>

Federal Aviation Administration. (2006). *National System Strategy Team* [Electronic Version]. Retrieved December 15, 2007, from <http://www.fly.faa.gov/Operations/NSST/nsst2006.pdf>

Federal Aviation Administration. (2007). *Traffic Flow Management for Flight Operations Personnel* [Electronic Version]. Retrieved December 15, 2007, from http://www.fly.faa.gov/Products/Training/Traffic_Management_for_Pilots/TFM_NASv2.pdf

Gaertner, N., Klopfenstein, M., & Wilmouth, G. (2007). *Updated Operational Concept for System Enhancements for Versatile Electronic Negotiation (SEVEN)* [Electronic Version]. Retrieved December 15, 2007, from <http://cdm.fly.faa.gov/Workgroups/ICEFM/2007/32F0907-005-R0%20Updated%20SEVEN%20Operational%20Concept.pdf>

Hoang, T. (2004, September). *A Description of the Mutli-Center Traffic Management Advisor (McTMA) Simulation Environment*. Paper presented at the American Institute of Aeronautics and Astronautics (AIAA) Aircraft Technology, Integration and Operations (ATIO), Chicago, Illinois.

Hoof, R. v., & Sierhuis, M. (2007). *Brahms Language Specification, TM99-0008* [Electronic Version]. Retrieved December 15, 2007, from http://www.agentisolutions.com/documentation/language/ls_title.htm

A Multi-Agent Simulation of Collaborative Air Traffic Flow Management

Ingrand, F. F., Georgeff, M. P., & Rao, A. S. (1992). An architecture for real-time reasoning and system control. *Institute of Electrical & Electronic Engineers (IEEE) Expert* 7(6), 34-44.

International Civil Aviation Organization (2003, November). *ATM Operational Concept Document*. Paper presented at the 11th Air Navigation Conference, Montreal, Canada.

Joint Planning and Development Office. (2007). *Concept of Operations for the Next Generation Air Transportation System* [Electronic Version]. Retrieved December 15, 2007, from http://www.jpdo.gov/library/NextGen_v2.0.pdf

Jonker, G., Meyer, J.-J., & Dignum, F. (2005, December). *Towards a Market Mechanism for Airport Traffic Control*. Paper presented at the 12th Portuguese Conference on Artificial Intelligence (EPIA 2005), Covilha, Portugal.

Klügl, F., & Bazzan, A. L. C. (2004). Route Decision Behaviour in a Commuting Scenario: Simple Heuristics Adaptation and Effect of Traffic Forecast. *Journal of Artificial Societies and Social Simulation*, 7(1).

Ljunberg, M., & Lucas, A. (1992, September). *The OASIS Air Traffic Management System*. Paper presented at the Second Pacific Rim International Conference on Artificial Intelligence (PRICAI '92), Seoul, Korea.

Metron Aviation. (2006a). *Flight Schedule Monitor* [Electronic Version]. Retrieved December 15, 2007, from http://www.metronaviation.com/content/text/collaborativeDecisionMakingFSM_2.pdf.

Metron Aviation. (2006b). *Post Operations Evaluation Tool* [Electronic Version]. Retrieved December 15, 2007, from http://www.metronaviation.com/content/text/collaborativeDecisionMakingPOET_2.pdf

National Oceanic and Atmospheric Administration. (2007). *Collaborative Convective Forecast Product: Product Description Document* [Electronic Version]. Retrieved December 15, 2007, from <http://aviationweather.gov/products/ccfp/docs/pdd-ccfp.pdf>

Pujet, N., Feron, E., & Rakhit, A. (1998, August). Modeling and identification of an Airline Operations Center as a multi-agent queueing system. Paper presented at the American Institute of Aeronautics and Astronautics (AIAA) Guidance, Navigation, and Control Conference (GNC 1998), Boston, Massachusetts.

Ruit, G. J. v. d., Schuylenburg, M. v., & Ottjes, J. A. (1995, June). Simulation of shipping traffic flow in the Maasvlakte port area of Rotterdam. Paper presented at the European Simulation Multiconference (ESM 1995), Prague, Czech Republic.

Sierhuis, M. (2007). It's not just goals all the way down - It's activities all the way down. In G. M. P. O'Hare, A. Ricci, M. J. O'Grady & O. Dikenelli (Eds.), *Engineering Societies in the*

Agents World VII (Vol LNAI 4457/2007). 7th International Workshop (ESAW 2006) (pp. 1-24) Dublin, Ireland: Springer.

Sierhuis, M., & Clancey, W. J. (2002). Modeling and Simulating Work Practice: A human-centered method for work systems design. *Institute of Electrical & Electronic Engineers (IEEE) Intelligent Systems*, 17(5).

Stark, H.-U., Helbing, D., Schönhof, M., & Holyst, J. A. (2006). Alternating cooperation strategies in a route choice game: Theory, experiments, and effects of a learning scenario. In A. Innocenti & P. Sbriglia (Eds.), *Games, Rationality, and Behaviour* (pp. 256-273). Houndmills, England and New York, New York: Palgrave Macmillan.

Tumer, K., & Agogino, A. (2007, May). Distributed Agent-Based Air Traffic Flow Management. Paper presented at the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007), Honolulu, Hawaii.

Usmani, A. (2005, November). *Upcoming TFM Enhancements: The Role of Simulation*. Paper presented at the FAA Eurocontrol Action Plan 9 (AP9) Traffic Flow Management (TFM) in Fast-Time Simulation Technical Interchange Meeting, Atlantic City, New Jersey.

Key Terms & Definitions

Air Traffic Control (ATC): A service operated by the appropriate authority to promote the safe, orderly, and expeditious flow of air traffic.

Air Traffic Flow Management (ATFM): The regulation of air traffic in order to avoid exceeding airport or airspace capacity, and to ensure that available capacity is used efficiently.

Airline Operations Center (AOC): An airline unit responsible for dispatching flights and adjusting schedules in response to restrictions in the airspace system.

Brahms: A set of software tools to develop and simulate multi-agent models of human and machine behavior.

Collaborative Decision Making (CDM): Collaboration involving the system stakeholders in determining the best approach to a given situation. In the context of air transportation, it is the cooperative effort between the government and industry to exchange information for better decision-making.

Traffic Management Unit (TMU): A team of air traffic controllers who analyze the demand and external effects, such as weather, on the airspace system and implement initiatives to balance the demand with capacity.