A Multiagent Simulation of Collaborative Air Traffic Flow Management

Shawn R. Wolfe, Peter A. Jarvis, Francis Y. Enomoto, Maarten Sierhuis, Bart-Jan van Putten
NASA Ames Research Center

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 alleviate 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, 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.

1. Introduction

Air traffic in the United States of America (U.S.A.) is forecasted to double or triple by the year 2025 (Federal Aviation Administration, 2000). Recent simulations (Grabbe, Sridhar, & Mukherjee, 2007) 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 alleviate air traffic problems. One such area is Collaborative Air Traffic Flow Management (CATFM). Today in the U.S.A., the Federal Aviation Administration (FAA) makes the bulk of traffic flow management decisions without consulting the airlines. In CATFM, the airspace users are given more opportunities to express their preferences, choose among options, and take proactive reaction. In theory, this will result in decreased workload for the FAA, increased airline satisfaction, and more efficient traffic flow management.

But will the CATFM concept really work in practice? Will the airlines take advantage of their 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, or just change it? Will the options available to the airlines enable them to substantially increase their benefit, in particular when many factors still remain out of their control? Will the uncoordinated actions of the individual airlines increase the efficiency of the system as a whole, even though each airline is only concerned with their own gain? Might potential efficiency gains be offset by the actions of rogue operators, who purposely seek to interfere with the operations of a competitor?

Given the current stage of development of the CATFM concept, and the fact that it 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 a future concept of CATFM operations so that strengths and weaknesses can be evaluated long before more costly human in the loop simulations or limited field deployments are attempted. Our simulation is still in an early stage, with several more years needed for a complete, detailed model, but we have already found several interesting and important properties of the CATFM concept.

Though our study is certainly most relevant to air traffic flow management (ATFM), aspects may relate to other forms of traffic. Our methodology (see Section 4.2) can be applied to other concepts of operation. Many of the basic concepts (e.g., choosing routes, traffic congestion, independent and uncoordinated agents) are the same and the overall structure is similar (see Section 2.1). However, there are important differences. Perhaps most the important is that aircraft must remain within a narrow range of speeds during flight – stopping is out of the question, only a slight speed reduction is possible without stalling, and significant speed increases can lead to instability. 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). CATFM distributes some elements of decision making, but ATFM generally has more centralized control than other forms of traffic management, Finally, a significant portion of air traffic is comprised of fleets (i.e., airlines) – essentially cabals of drivers who are interested in cooperating for the common good of the company.

When viewed abstractly, CATFM is an excellent domain to develop and evaluate system designs that could be applied in other agent-based systems, particularly those that model people. The constraints of air traffic management are essentially a competition for limited and shared resources. The agents in the system are neither wholly cooperative (which is rarely realistic when agents are self-interested), nor entirely competitive (which can lead less efficient overall performance). Rather, we have two types of entities: 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 beyond our control, may include antagonistic elements, and are generally unknowable, complicating matters. Yet, this situation occurs often in practice, not only in government-controlled situations, but also in any system with central authority, such as companies, organizational bodies, and games of all 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 was meant to address. Our approach to developing a simulation of this concept of operations is presented, and we describe our simulation thus far. We discuss the comparative results of different CATFM approaches and different airline strategies. We conclude with an analysis of these experiments, and present our goals for future development of the simulation.

2. Background

2.1. Introduction to ATFM

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

Air Traffic Flow Management (ATFM) is a system level function within ATC to manage the flow based on capacity and demand. ATFM is the responsibility of a Traffic Management Unit (TMU) within each ARTCC and at the ATCSCC. The ATCSCC develops strategic plans to ensure balanced flow throughout the NAS over a planning horizon of two to eight hours. The center TMUs develops 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 initiatives. 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. For example, a TMU may identify a Flow Constrained Area where a capacity-demand imbalance may occur due to severe convective weather. The TMU would then analyze which type of traffic management initiative should be invoked to alleviate the traffic imbalance. Since TMIs may affect adjacent centers, either directly or through ripple effects, ATCSCC approval is needed before invoking a TMI. TFM issues are discussed in a bi-hourly planning teleconference, involving representatives from the ATCSCC, each ARTCC, and airlines.

A variety of TMIs are available to the FAA, depending on the nature of the traffic flow problem; we describe some commonly used TMIs. A re-route procedure directs an aircraft onto a new route to avoid a problem area, such as a severe thunderstorm or congested airspace. This is the only TMI we have implemented in our current simulation. Re-routing can impact both ATC and NAS users: workload of the sector controllers receiving the diverted traffic increases, expected aircraft arrival times may change, and more fuel may be needed for the aircraft to follow the new route. A Ground Delay Program (GDP) delays aircraft at the departure airport in order to manage the demand at the arrival airport. Flights are assigned new (delayed) 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 a certain separation between aircraft transiting through some point (e.g., an airport, sector boundary, or route). MITs are used to apportion traffic into a manageable flow, as well as provide space for addition traffic (merging or departing) to enter the flow of traffic. A MIT procedure can cause the traffic flow to back up, potentially resulting in a larger MIT restriction in the upstream center (known as a passback) or delayed departures.

Airlines manage their fleet of aircraft in an Airline Operations Center (AOC). The AOC typically has a position called the ATC coordinator that monitors the TMIs issued by the TMUs and ATCSCC and participates in the planning teleconference to make their concerns visible to the FAA. A major thrust of the CATFM concept is to increase the role of the AOC in ATFM, as detailed in Section 4.1).

2.2. Agent-based ATFM Simulations

The Airspace Concept Evaluation System (ACES) (Sweet, Manikonda, Aronson, Roth, & Blake, 2002) is a distributed agent-based simulation of the NAS that uses a "activity centric paradigm". 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 ATFM (Couluris et al., 2003), but is not its only focus. In ACES, individual reasoning entities are represented as agents, as well as the different simulation layers connected through HLA.

IMPACT (Intelligent agent-based Model for Policy Analysis and of Collaborative Traffic flow management) is a swarm-based agent model of FAA agents and airline agents, simulating several possible responses to capacity reductions: no TMIs, GDPs without information sharing, and GDPs with shared airline schedules. (Keith C, Cooper, Greenbaum, & Wojcik, 2000). In each scenario, the FAA agents decide to impose GDPs or not based on predefined policies, and 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 surprisingly found that GDPs without shared information (as occurs in today's operations) resulted in a *greater* average cost per flight than when no TMIs were instituted.

STEAM (Tambe, 1997) has been used to model 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, and is complementary to our focus on the traffic flow level. Unlike our model, where the collaboration also includes the airspace users, only the sector controllers and a few higher-level coordinating entities coordinate to find solutions to the ATFM problems.

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), but at a different level of granularity than in our work, as it emphasizes the capabilities and limitations of human cognitive ability instead of complex decision making at a more abstract level.

3. Issues, Controversies, Problems

3.1. Characterizing Operations and Issues through Field Observations

In order to more accurately characterize the problems present in air traffic flow management today, and to identify likely challenges in future operations, field observations were conducted at several operational centers (Husni Idris, Evans, Vivona, Krozel, & Bilimoria, 2006). A diverse set of collection of facilities were included in these observations to provide a wide scope of operational characteristics and corresponding issues, including five TMUs, five AOCs, and the ATCSCC. The five TMUs included centers of varied geographical size, assorted weather patterns, and differing dominant traffic patterns (such as terminal or overflight operations). The AOCs included both large and small carriers, with different operational models and customers (such as passenger or cargo). Finally, the ATCSCC provided a unique perspective of air traffic flow management across the country.

These field observations supported the development of the CATFM concept of operations primarily 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 lead to the inefficient handling of these situations (detailed in the following section) could be identified as well. Finally, these observations of how the ATFM participants actually did their work provide a valuable record of *work practice* (see section 4.2). By analyzing how the work is done, potential solutions (see Section 4.1) that could feasibly be adopted were developed, and a realistic agent-based model of ATFM operations can be built.

3.2. Identified ATFM Issues

The primary finding from the field observations was that the current ATFM system limited the amount of possible collaborative problem solving. Two factors dominate the source of these issues. First, the sharing of information between the FAA and airlines is limited. This means that the bulk of the planning must be done 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 framework and validate the choice of an agent-based model for simulation). Second, the bulk of the problem solving activities falls upon the FAA, but the TMUs face workload limitations that in turn limits the solutions they can realistically pursue. We present a summary of these findings: the complete list can be found in (H. Idris, Vivona, Penney, Krozel, & Bilimoria, 2005)

3.2.1. Inaccurate Problem Assessment

Efficient management of traffic flow issues begins with an assessment of the problem. Inaccurate assessments of either the demand or the capacity can lead to problematic problem assessments, including over- or underestimating the problem severity, missing a problem or incorrectly identifying a non-existent problem. Factors that lead to inaccurate *demand* assessments include inaccurate prediction of pop-up traffic, changes in airline flight intent (e.g., change in departure time, flight plan or cancellations), and displacement of traffic from flow constraints elsewhere. Factors that lead to inaccurate *capacity* assessments include incorrect weather and airspace restriction predictions. In addition to the inaccuracy of assessments, the FAA and airlines are likely to have divergent assessments, resulting in inconsistent plans.

3.2.2. Differing Evaluations of Identified Problem

Once the traffic flow problem is identified, the FAA and the airlines are likely to have differing evaluations of the problem. After safety, they have divergent concerns. The FAA will seek to minimize the affect of the problem on the NAS and limit the amount of workload for the controllers. The airlines, on the other hand, are not concerned with the NAS except as it affects their own flights (e.g., they are not seeking to minimize the negative effect on competing airlines). Instead, the airline seeks solutions to problems that adhere to their business model, often with a goal of minimizing costs (such as fuel usage) while limiting the negative effect on the customers. Moreover, different carriers will have different business models, therefore regarding cost, reliability and on-time service differently. Thus, even with a consensus on the traffic flow problem, different entities will often prefer different solutions.

3.2.3. Limited Mitigations

The TMUs and the ATCSCC have a limited set of TMIs to choose from when seeking mitigations to a traffic flow management issue. These TMIs are typically coarse-grained and are applied uniformly to all users. As such, the initiatives are often overly restrictive, and because they are not selective, may impact user operations unevenly.

3.2.4. High TMU Workload

Two factors typically contribute to a high TMU workload when the disruptions to the NAS grow more severe. First, the reliance on direct synchronous communications such as telecons and phone calls increases the cost of communication, decreasing both the time available for such communications and for other activities. Secondly, actions at the level of an individual flight (such as rerouting) greatly increase the quantity of tasks that must be performed by the TMU. The end result is that TMU workload becomes a limiting factor, and restricts the possible solutions the FAA can pursue.

3.2.5. 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. As such, the FAA and the AOC assess, evaluate and plan largely independently from one another. This is exacerbated by the relative unpredictability of both parties, potentially leading to a double penalty for either party; the TMU may choose unnecessary mitigations and be unprepared for the actual resulting problem, and the AOC may independently avoid one restriction only to be impacted by an unanticipated mitigation. In addition, 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 unnecessarily avoiding a problem that no longer exists.

4. Solutions and Recommendations

4.1. CATFM Concept of Operations

The CATFM concept of operations recommends several changes to the ATFM process to address the issues presented in Section 3.2. Most of these changes fall under the following three categories, listed by order of increasing emphasis. First, more automation must be used to reduce the workload of TMUs. This automation may reduce the need for the TMU planners to perform mundane tasks and lessen the cost of communication. Second, more information should be shared between the ATSP and the users. By doing so, assessments can be made with more complete information, common assessments are possible, and since the information used to make decisions is shared, actions are more predictable. Finally and most importantly, when possible, the AOCs should be allowed to participate directly in the traffic flow management process, shifting the burden from the ATSP to the users.

We summarize the four phases of the ATFM process in the CATFM Concept of Operations below; a more complete description can be found in (H. Idris et al., 2005)

4.1.1. Common Problem Identification

As described previously, ATFM problems are caused by situations where the demand for the airspace exceeds its capacity. Each element is best predicted by a different party: demand, by the airspace users who create it; and capacity, by the ATSP, as it is an assessment of the ATSP's ability to manage traffic in the effected area. This leads naturally to a collaborative situation where both parties share information about their respective components to produce a more accurate problem assessment, and to minimize the divergence of problem assessments.

4.1.2. Shared Impact Assessment

The ATFM issues identified in the previous phase are constraint violations that may be rectified through a variety of actions, each of which will have different impacts on airline and ATSP operations. Automatically translating the problem assessments into potential impacts will greatly aid the planning process of the AOCs and TMUs. Without it, as in today's system, uncertainty increases, leading to overly conservative planning. 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.

4.1.3. 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 that adequately addresses 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 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.

4.1.4. Joint Plan Implementation

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

4.2. Approach

We have built an initial agent-based simulation with Brahms (Clancey, Sierhuis, Kaskiris, & Hoof, 2003). 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 (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 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. (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).

The methodology used in OCAMS (Clancey et al., 2007) describes how to design and simulate a future work system. 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 aspects of the concept of operations, modify the concept according to the findings from our simulation, and then increase model fidelity in the next stage.

In our current, initial stage, we have built a rudimentary model of ATFM (see Section 4.3) using second-hand sources of information such as the original work practice observations, other ATFM literature, and the concept of operations itself. Subject matter experts have validated the design and behavior of our current model. 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 models 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.

4.3. Initial CATFM Simulation Design

We have created a simplified model of ATFM in our initial simulation, concentrating on the Joint Plan Implementation phase (see Section 4.1.4) where routes are assigned to flights. 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. This also makes the problem more like other forms of traffic that tend to have route-based capacities. Airspace conditions are held static throughout the simulation and the flight schedules are fixed, though these parameters are known only to the FAA, and demand on a given route will fluctuate according to previous selections. We have abstracted the details of staggering flights and airport limitations, dealing with all flights simultaneously within a given time window.

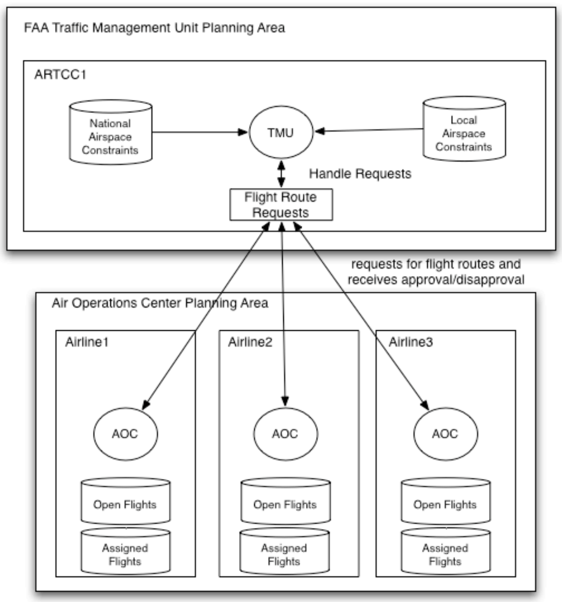

Figure 1. ***Agent architecture.***

Figure 1 provides an overview of our agent current 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 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 agents 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 to the (under capacity, at capacity, or oversubscribed) to the AOC agents. The TMU also has the responsibility of ensuring that no route goes above capacity; it does this by preventing all flights from accessing a route if demand exceeds capacity when the planning phase is over.

For each origin-destination pair, we created three routes: 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 the traffic. Our fundamental question is: how will the CATFM concept perform in this simplified model? In order to answer this question, we made several variations of the model, one that uses a CATFM approach and several other approaches for comparison:

- **Blue Sky (BS)**: In this variation, all capacities were set to infinite so every flight could take the direct route. This is not a viable approach but an upper bound on performance that we use as a baseline.
- **Current Operations (CO)**: Here the TMU agent is responsible for the route selection. It puts flights on the best available routes (i.e., under capacity routes), but does so in an arbitrary order without inspecting the flight value. This approach is closest to the current operations where the FAA must make route assignments without input from the airlines.
- **Global Optimum (G)**: Again, 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.

- **Airline Planning (AP)**: Here the responsibility for route selection falls upon the AOC agents. After the TMU agent gives feedback on the route status, the AOC agent may choose a new route or continue to request the shortest route. This continues for several cycles until the time for planning is exhausted. We developed several simple strategies for the AOC agents:
    - **Aggressive (A):** An AOC agent with the Aggressive strategy will always request the best route for every flight, regardless of the situation.
    - **Moderate (M):** An AOC agent with the Moderate strategy will request the next best route for some of its flights when faced with an overcapacity situation.
    - **Conservative (C):** An AOC agent with the Conservative strategy will request the *worst* route for some of its flights when faced with an overcapacity situation. 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.

4.4. 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.

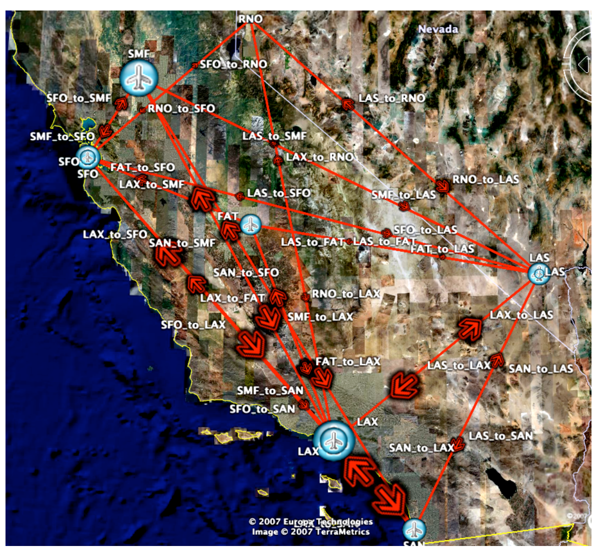

Figure 2. ***Scenario involving seven airports (in Google Earth)****. Only the best route is displayed; larger arrows and airport symbols indicate higher levels of traffic.*

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 of the flight to the airline. We define v_F , the flight value of *F*, as

$$v_F = p_u + 3p_c + 5c_c \tag{1}$$

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

$$d_F = t_F - t_B \tag{2}$$

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.

Finally, we seek to measure the total incurred 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

$$d_T = (p_u * d_F) + (p_c + d_F) + 60p_c + 300c_c \tag{3}$$

The local traffic scenario, when simulated with several different TFM approaches and AOC strategies, yielded some surprising results. When given a mix of strategies, the Airline Planning approach performed poorly. This is shown in

Figure 3, where the light blue bars indicate delay incurred by selecting an alternate route, and the dark red bars indicates delay from failing to make a route assignment. Figure 4 shows that the Airline Planning approach is highly sensitive to the strategies employed by the AOC agents. The best mix of airline strategies outperforms the Current Operations approach (see Figure 5), indicating a potential for improvement under the CATFM concept. Further observation reveals that the Aggressive strategy in particular is disruptive to the system as a whole, and typically a very poor strategy in terms of our passenger delay metric. Figure 6 shows that in some cases, the quality of the solutions found in the Airline Planning approach is affected by the number of planning cycles.

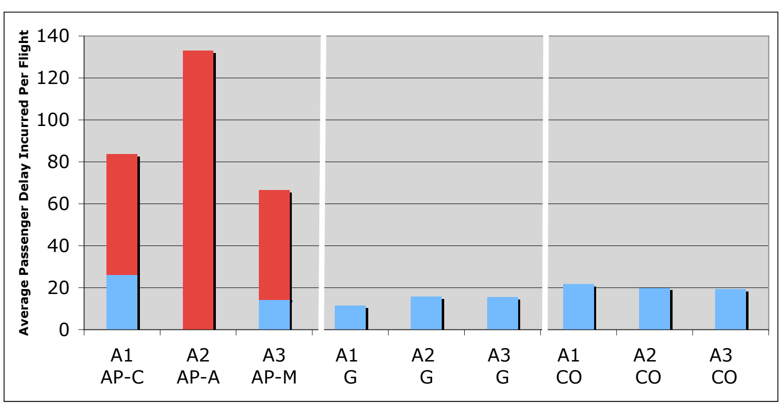

Figure 3. ***Comparing TFM Approaches on the Local Scenario***.

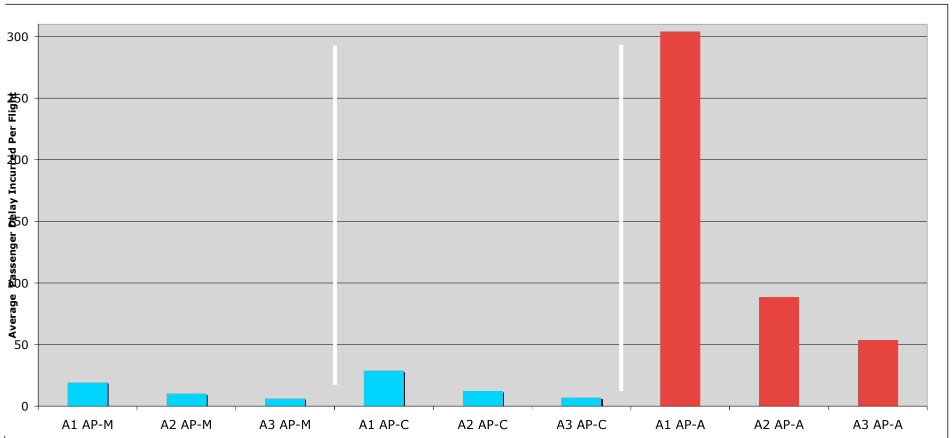

Figure 4. ***Interactions between AOC Strategies.***

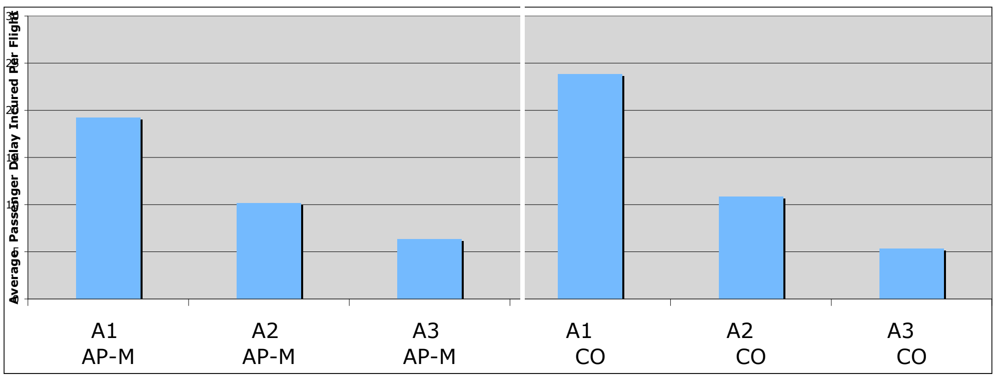

Figure 5. ***Best AOC Strategy Mix Compared with Current Operations Approach.***

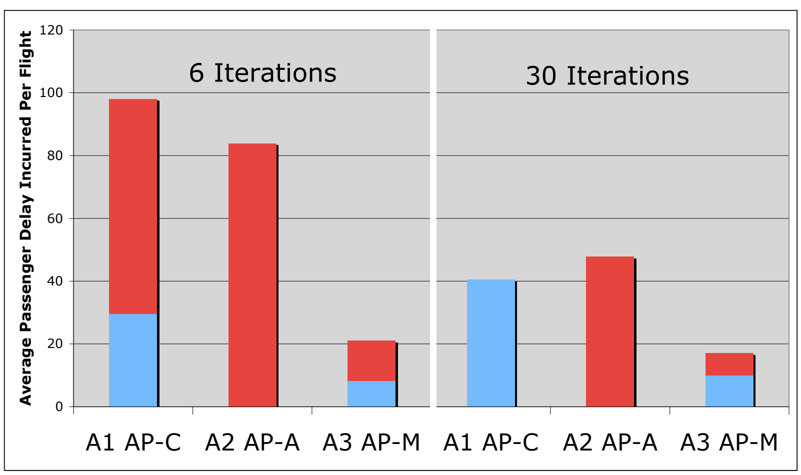

Figure 6. ***Effect of Additional Planning Cycles with Airline Planning Approach.***

4.5. 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 (M)**: 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 that the AOC agents failed to choose routes for.
- **Equitable (E)**: 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: in the Excess Capacity scenario, each route could accommodate five flights; in the At Capacity scenario, each route could accommodate four flights; in the Under Capacity scenario, each route could accommodate three flights. Therefore, all flights could be assigned a route on the At Capacity and Excess Capacity scenarios, but this was not possible in the Under Capacity scenario as demand exceeded capacity.

We ran each scenario with all combinations of strategies for the three AOC agents using both the Airline Planning and Mixed approaches, resulting in twenty-seven runs for each. From these runs, we aggregated the situations where one strategy was used against another. For each pair of strategies, we calculated the average incurred passenger delay per flight as well as percentage of time that one strategy outperformed the other. Table 1 and Table 2 show the delay average metric and winning percentage, respectively, when using the Airline Planning approach; Table 3 and Table 4 show the delay average metric and winning percentage, respectively, when using the Mixed approach.

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, leaving its flights unassigned. However, head-to-head performance improves for the aggressive AOC agent when there is excessive capacity, as AOC agents with more flexible strategies (i.e., Moderate or Conservative) make room for the aggressive airline. When the Mixed approach is used, the Aggressive strategy does even better. In situations with adequate capacity, the Aggressive strategy performs best overall. The Aggressive strategy will either succeed in putting all of its flights onto the best route, or is will succeed in preventing any other airline from using the best route (though stranding its own flights in the process), thus making the best route available for its flights when the TMU assigns the remaining flights. However, when there is not sufficient capacity, this strategy performs poorly because not all of its flights will be assigned.

The Conservative strategy performs poorly compared with the other strategies. This is not surprising, as it prefers the worst route. However, it is generally not favorable to use the same strategy as another AOC agent because both AOCs will create demand on the same routes. In our experiments, the Conservative strategy fared better in terms of our delay average metric when matched against a moderate opponent, and vice versa. In particular, in the Under Capacity experiment, the Conservative strategy is the best when faced with a moderate opponent.

Finally, we created a larger scenario with a primary and secondary routes defined as before, but each with a capacity of forty flights, and three airlines with forty flights each. Table 5 shows the results of experiments on this scenario in terms of the total 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.

Airline Strategy	Excess Capacity: Competitor's Strategy A	M	C	At Capacity: Competitor's Strategy A	M	C	Under Capacity: Competitor's Strategy A	M	C
A	1244	415	415	1244	1244	1244	1244	1244	1244
M	454	252	252	706	559	426	770	742	652
C	567	337	337	792	498	577	843	724	748

Table 1. ***Average Incurred Passenger Delay (minutes) per Flight with Airline Planning Approach***

Airline Strategy	Excess Capacity: Competitor's Strategy A	M	C	At Capacity: Competitor's Strategy A	M	C	Under Capacity: Competitor's Strategy A	M	C
A	50%	67%	67%	50%	0%	0%	50%	0%	0%
M	33%	50%	100%	100%	50%	61%	100%	50%	61%
C	33%	0%	50%	100%	39%	50%	100%	39%	50%

Table 2. ***Head-to-Head Winning Percentage with Airline Planning Approach***

Airline Strategy	Excess Capacity: Competitor's Strategy A	M	C	At Capacity: Competitor's Strategy A	M	C	Under Capacity: Competitor's Strategy A	M	C
A	167	50	50	301	245	210	566	632	614
M	335	252	252	381	374	313	444	561	483
C	448	337	337	460	385	389	512	556	567

Table 3. ***Average Incurred Passenger Delay (minutes) per Flight with Mixed Approach***

Airline Strategy	Excess Capacity: Competitor's Strategy A	M	C	At Capacity: Competitor's Strategy A	M	C	Under Capacity: Competitor's Strategy A	M	C
A	50%	100%	100%	50%	89%	100%	50%	28%	39%
M	0%	50%	100%	11%	50%	67%	72%	50%	61%
C	0%	0%	50%	0%	33%	50%	61%	39%	50%

Table 4. ***Head-to-Head Winning Percentage with Mixed Approach***

	A1	**A2**	**A3**	**Total**
Current Operations	3552	4332	2939	10823
Optimal	3314	2806	3300	9420
Equitable	2969	3407	3073	9449

Table 5. ***Comparing Current Operations, Optimal, and Equitable Approaches***

5. Conclusion

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; nonetheless, it was the best choice when competing against moderate AOCs in situations without excess capacity. 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 in the when other AOC agents moved their flights off the best route, thus accommodating the aggressive AOC.

In our evaluation of the CATFM concept, we observed that nearly all the approaches that incorporated the flight value yielded better results than the Current Operations approach. This supports the claim that utilizing user preferences in ATFM should lead to better solutions. However, this was not the case in all of our experimental results; certain mixes of strategies with the Airline Planning approach produced unacceptably poor results. Moreover, 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, not collaboration, is most likely the key to reducing FAA workload. This is consistent with the CATFM concept of operations. 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 is.

In the end, the challenge of refining the CATFM concept will not be designing effective AOC agent strategies, as in practice, they will be determined by the airlines and our of our control. 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 incentivizes 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 too vulnerable to aggressive AOC agents; likewise, the Mixed approach often incentivized 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 experiment)s 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.

6. Future Research Directions

We have completed the initial stage of development and will continue to expand the CATFM model according to the methodology presented in Section 4.2. We have begun work on the next stage, expanding our model to capture the breadth of the CATFM concept of operations, covering all the phases presented in Section 4.1. 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 are the ones with the greatest impact on operations.

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 show 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. 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 to 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.

7. References

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). *Advantages of Brahms for Specifying and Implementing a Multiagent Human-Robotic Exploration System.* Paper presented at the The 16th International FLAIRS Conference, St. Augustine, FL.

Clancey, W. J., Sierhuis, M., Seah, C., Buckley, C., Reynolds, F., Hall, T., et al. (2007). *Multi-Agent Simulation to Implementation: A Practical Engineering Methodology for Designing Space Flight Operations*. Paper presented at the The Eighth Annual International Workshop "Engineering Societies in the Agents World" (ESAW 07).

Corker, K. M. (1999). *Human Performance Simulation in the Analysis of Advanced Air Traffic Management*. Paper presented at the I999 Winter Simulation Conference.

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

Federal Aviation Administration. (2000). *FAA Long-Range Aerospace Forecasts Fiscal Years 2015, 2020 And 2025* (No. FAA-AO-00-5)o. Document Number)

Grabbe, S., Sridhar, B., & Mukherjee, A. (2007). Regional Management of Chicago Arrival Flows, *Presentation to Code AFC*. NASA Ames Research Center, Moffett Field, CA.

Idris, H., Evans, A., Vivona, R., Krozel, J., & Bilimoria, K. (2006, September 25-27, 2006). *Field Observations of Interactions Between Traffic Flow Management and Airline Operations*. Paper presented at the AIAA 6th Aviation, Technology, Integration, and Operations Conference (ATIO), Wichita, Kansas.

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

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

Nguyen-Duc, M., Briot, J.-P., Drogoul, A., & Duong, V. (2003, October 13-17). *An application of Multi-Agent Coordination Techniques in Air Traffic Management.* Paper presented at the 2003 IEEE/WIC International Conference on Intelligent Agent Technology (IAT 2003), Halifax, Canada.

Sierhuis, M. (2001). *Modeling and Simulating Work Practice; Brahms: A multiagent modeling and simulation language for work system analysis and design.* Unpublished Ph.D. Dissertation, University of Amsterdam, Amsterdam, the Netherlands.

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 AIAA Modeling and Simulation Technologies Conference and Exhibit, Monterey, California.

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

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

8. 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. Pujet et al. have developed a simulation of the United Airlines AOC, (Pujet, Feron, & Rakhit, 1998), modeling each AOC employee 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 at 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 using different methods problem-solving methods compete to find the best solution; in simulation, this approach produced better solutions that 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 multiagent 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 NASA-developed tool for simulating air traffic flow that has been integrated into a commercial product used by nearly all major U.S. airlines ("Flight Explorer,"). 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. CTAS (Erzberger, 1994) is another NASA developed simulation system, with a greater emphasis on human in the loop simulations. The Traffic Management Advisor of CTAS 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 similarities to 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, the Post Operations Evaluation Tool (Metron Aviation, 2006b) for analysis support of NAS operations.

The Future Concepts Team is another sub-group of the CDM initiative. Over the past few years, the FCT group has focused their effort on future of collaboration between the service provider and the 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 user's 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 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, TMI 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.

9. Additional Reading References

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

Air Transport Association of America. (2007). *The Airline Handbook*.

Ball, M., Hoffman, R., Chen, C.-Y., & Vossen, T. (2000). Collaborative Decision Making in Air Traffic Management: Current and Future Research Directions [Electronic Version],

Bilimoria, K., Sridhar, B., Chatterji, G., Sheth, K., & Grabbe, S. (2000, June 13-16). *FACET: Future ATM Concepts Evaluation Tool.* Paper presented at the 3rd USA/Europe ATM 2001 R&D Seminar, Napoli, Italy.

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

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). *Center-TRACON Automation System (CTAS)*. Paper presented at the Capacity Technology Subcommittee, FAA Research and Development Advisory Committee.

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

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

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

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

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

Flight Explorer. Flight Explorer Inc.

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

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

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

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

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

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

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

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). *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]. *http://www.metronaviation.com/content/text/collaborativeDecisionMakingFSM_2.pdf*. Retrieved December 15, 2007,

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

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

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

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

Sierhuis, M. (2007). *It's not just goals all the way down - It's activities all the way down*. Paper presented at the 7th International Workshop on Engineering Societies in the Agents World (ESAW 2006).

Sierhuis, M., & Clancey, W. J. (2002). Modeling and Simulating Work Practice: A human-centered method for work systems design. *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*. Paper presented at the International Workshop on Behavioural Game Theory and Experiments.

Tumer, K., & Agogino, A. (2007). *Distributed Agent-Based Air Traffic Flow Management*. Paper presented at the AAMAS 2007.

Usmani, A. (2005). *Upcoming TFM Enhancements: The Role of Simulation*. Paper presented at the TFM in Fast-Time Simulation.