

Effects of Airline Strategies on Airline Operations and Infrastructure Utilization

Guillermo Calderón-Meza (Ph.D. Candidate)*

Center for Air Transportation Systems Research,
George Mason University, Fairfax, Virginia, 22030, USA

Lance Sherry (Ph.D)[†]

Center for Air Transportation Systems and Research,
George Mason University, Fairfax, Virginia, 22030, USA

Plans for NextGen include the implementation of Net-Centric Operations (NCO) as the underlying mechanism to more effectively manage National Air Transportation System (NAS) resources. NCO mechanisms will provide the means for all agents operating in the NAS to self-synchronize in the presence of changing situations and probabilistic forecasts. Applications of NCO in other domains (e.g. military war fighting and health care) have identified the importance of developing rules and procedures for NCO to ensure that the behavior of all agents yields a net benefit to the system.

This paper describes the results of a multi-agent simulation of the evolution of Airline Operation Center (AOC) decision-making for trans-continental route selection when demand for the routes is in excess of the available slots. The results show the emergence of *hawk* and *dove* alliances in AOC strategies. In some cases, these strategies yield improved performance and utilization of NAS resources. The implications of these results are discussed.

I. Introduction

Plans for NextGen¹ include the implementation of Net-Centric Operations (NCO)²⁻⁴ as the underlying mechanism to more effectively manage National Air Transportation System (NAS) resources. NCO mechanisms will provide the means for all agents operating in the NAS to self-synchronize in the presence of changing situations and probabilistic forecasts. Applications of NCO in the NAS (e.g. Collaborative Decision Making for Ground Delay Programs)⁵ and in other domains (e.g. military war fighting and health care) have identified the importance of developing rules and procedures for NCO to ensure that the behavior of all agents yields a net benefit to the system.² Wojcik,⁶ Campbell et al.,⁷ and Waslander and Tomlin⁸ have explored the adaptive behaviors of agents in the NAS, e.g. airlines, to better understand the effects of competition and cooperation on system-wide performance and resource utilization. This studies the selection of trans-continental routes, like the Playbook Routes^a, by selfish, but informed airlines.

This paper describes the results of a multi-agent simulation of the evolution of Airline Operation Center (AOC) decision-making for trans-continental route selection when demand for the routes is in excess of the available slots. The results show the existence of a tradeoff between *performance* and *utilization* when all the agents exhibit homogeneous strategies. The results also show the emergence of *hawk* and *dove*⁹ alliances in AOC strategies when the strategies are allowed to differ between airlines. In some cases, these strategies yield improved performance and utilization of NAS resources. The implications of these results are discussed.

The paper is structured as follows. The first section of the paper defines the objective. The second section explains the methodology of the experiments and provides definitions of concepts used in the experiments.

*Research Assistant, Email: gcaldero@gmu.edu

[†]Director, Email: lsherry@gmu.edu

^a<http://www.fly.faa.gov/PLAYBOOK/pbindex.html>

This section also describes the tools used and behaviors of the agents. The third section presents the designs of the experiments and their results. The last two sections are about conclusions and future work.

II. Objective

The general objective of the research is to *forecast* the consequences in *utilization* and *performance* of implementing NextGen, i.e. Net-Centric Operations, in the trans-continental air route selection in the National Air Transportation System of the United States of America. Three specific objectives of the research are the following:

- Compute the *tradeoff* between performance, i.e. airline cost, and route utilization in the presence of NextGen
- Measure the increase in the *effective capacity* through *cooperation / gaming* in the selection of trans-continental routes
- Evaluate the role of *availability* and *latency* of information in the trans-continental route selection

This study has the objective of *forecasting* the effects of airline *strategies* on airline *performance* and infrastructure *utilization*.

III. Methodology

III.A. An instance of the trans-continental route selection problem

The scope of the research is limited to the *route selection problem*. In particular, the situations in which flights must choose among trans-continental routes in the United States. Airline Operations Center (AOC) choose the routes for the flights before the take off, and flights are not rerouted after they are airborne.

Currently, TFM faces this choice when capacity reductions keep flights from using their intended routes. To manage these situations TFM can choose among standard alternate routes defined in a document called *Playbook*^b. There are 19 East to West trans-continental route, 3 South to Northeast route, and 21 West to East trans-continental routes in the Playbook. Each route affects a particular set of facilities. Facilities can be navigation aids, route segments, or airports. Any flight whose flight plan included one of these facilities is sent to an alternate route from the Playbook.

Figure 1 shows an example with two alternate West to East trans-continental routes, drawn in blue. The distances of the routes are not equal. Both routes share the same destination airport or metroplex, circled in green. And all the origin airports, red dots, have access to both routes also.

In this situation the controllers direct the traffic based on the instructions from the Traffic Flow Management (TFM). TFM decisions are intended to balance the traffic among routes and reduce congestion. NextGen is intended to allow AOCs to do part of this task (i.e. self-synchronization) by information sharing.

In this research, flights can choose routes by themselves, communication with the AOC is assumed, based on available information and on some individual criteria, i.e. the strategy of the airline.

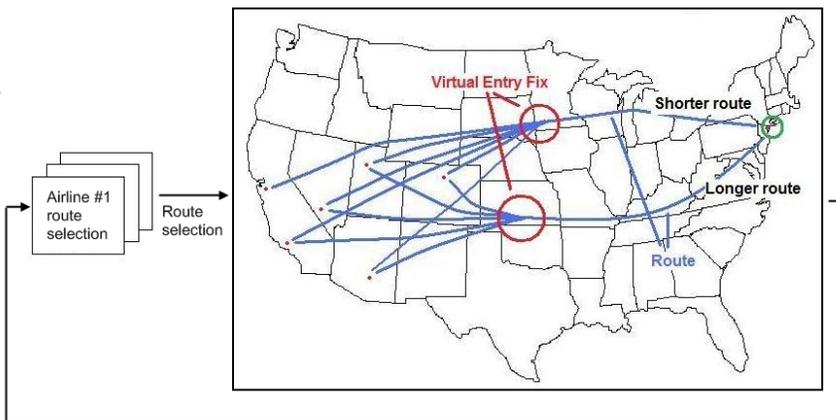


Figure 1. Example of alternate west to east trans-continental (playbook like) routes with alternate airline (AOC) route selection strategy.

^b<http://www.fly.faa.gov/PLAYBOOK/pbindex.html>

It is assumed that flights push back from the gates at the scheduled departure time plus a *push back delay* that is exponentially distributed. The use of an exponential distribution is justified by analysis, made by the Center for Air Transportation Systems and Research^c, of historic data published by the *Bureau of Transportation Statistics* (BTS).

The routes are divided into slots due to safety rules, e.g. *separation distances*. Controllers assign routes to flights before allowing them to enter the route. Usually, flights wait on the ground until the controllers allow them to fly just-in-time to enter an available slot in a route. Once a flight enters a route, it flies at a speed that allows it to maintain a safe separation with the flights before and after it is in the route.

Each route has an *entry queue* associated, which is also divided into slots. For modeling purposes, the entry queue is located before the *virtual entry fix*, marked by the red circles in Figure 1. The virtual entry fixes are merging points for flights coming from the airports of departure. The number of slots from the virtual entry fix to the destination is called the *distance* of the route. The ratio between the two distances is a parameter for the simulations called *distance ratio* and defined by Equation 1.

$$distance_ratio = \frac{distance(Route_0)}{distance(Route_1)} \quad (1)$$

A simplification used in this study is that all the airports of departure are merged into a single airport. So there is only one schedule for each airline.

Another simplification is that airlines use only one type of aircraft. So the speed of the aircraft are equal; it is normalized to 1.0^d. Also, the slots of the routes, i.e. *distance slots*, are of fixed distance. Consequently, the time flights wait on the ground is measured also in *time slots* of fixed duration: the time it takes a flight to fly a distance slot at constant speed.

It is also assumed that airports can accommodate all the scheduled flights. It is not necessary to change the speed of the flights while they are flying or to reroute them either while they are flying or after they have selected a route.

The airborne costs are equal between airlines, as are the ground costs. The cost of flying a distance slot, i.e. airborne cost, is normalized to 1.0. And the cost of waiting on the ground, i.e. ground cost, is expressed as a fraction of the normalized airborne cost. The ratio of the ground cost to the airborne cost is a parameter of the simulation and it is defined by Equation 2. Airborne and ground costs are functions of the number of slots. Utilization is reported also as number of slots.

$$cost_ratio = \frac{ground_cost}{airborne_cost} \quad (2)$$

Three airlines are considered in the experiments. Each airline has 30 flights scheduled to depart. The schedules are flat, i.e. there is one flight per time slot for each airline. The schedules are uniform, i.e. there are no empty time slots between flights.

III.B. Architecture

The goals of the study are achieved by integrating several components (see Figure 2): a NAS simulator, a decision making module, a multi-objective Evolutionary Computation platform, a data repository, and a visualization module. The following subsections explain each of these architectural components.

III.B.1. NAS simulator

The NAS simulation provides a neutral (i.e. unbiased towards any particular strategy or regulation) representation of flights, airports, routes, weather, and other elements of the NAS that are relevant to the experiments.

There are many general purpose, free, open source, and portable *multi-agent simulation* (MAS) tools that can serve the purpose of simple NAS simulator. In the experiments of this study MASON^{e10} was used because it is easy to use and because there was direct support from the developers. MASON was developed at George Mason University. Another reason to use MASON was that the experiments did not require a detailed simulation of the NAS, but a very abstract situation.

^c<http://catsr.ite.gmu.edu/>

^dThe actual value is not relevant as long it remains constant.

^e<http://cs.gmu.edu/~eclab/projects/mason/>

III.B.2. Multi-Objective Evolutionary Computation platform

The multi-objective Evolutionary Computation platform provides a way to achieve adaptation of the strategies. The platform uses the simulation (see the stereotype `«uses»` in Figure 2) to evaluate combinations of strategies. The platform supports the computation of Pareto surfaces based on the results of the executions of the Evolutionary Algorithms. Hence, it must be multi-objective.^{11,12}

The platform *EC3* was used in the experiments of this study.¹³ *EC3* is an academic Evolutionary Computation tool intended to demonstrate Evolutionary Algorithms. *EC3* includes several standard Evolutionary Algorithms, e.g. Evolution Strategies,¹⁴ Genetic Algorithms,¹⁵ Evolutionary Strategies.¹⁶ It is written in Java and provides a basic GUI to control the algorithms and display results. The source code was provided by the author. The tool required some modifications to make it suitable for a multi-objective optimization problem. Some of these modifications are still in being developed.

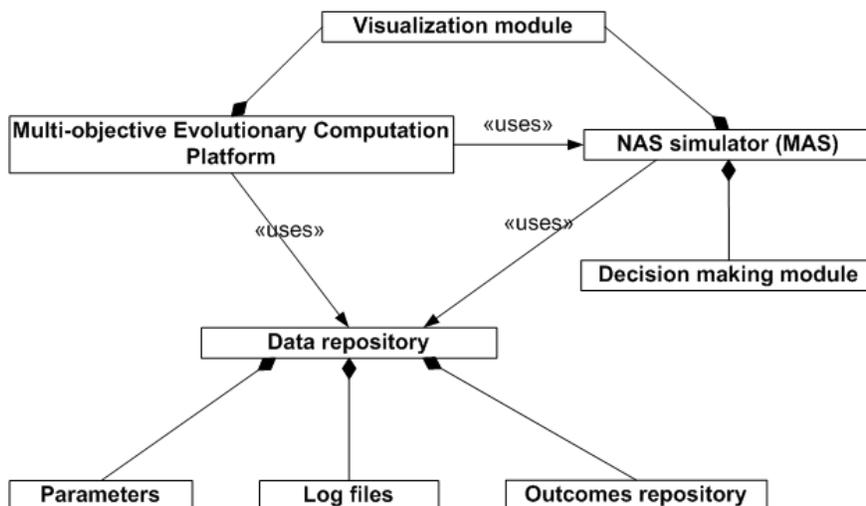


Figure 2. UML class diagram of the architecture of the simulation.

III.B.3. Data repository

The data repository provides a way to analyze and reproduce experiments. It is also the source of parameter values for the simulation. It can be implemented as a database, a set of log files, or a combination of both.

The data repository includes three separated elements: the parameters, the log files, and the outcomes repository. The following paragraphs explain these components.

Parameters

In the current implementation of the simulation, the parameters are given in a combination of *command line parameters* and configuration files. The MASON multi-agent simulation receives command line parameters. The *EC3* application uses configuration files in text format, but the GUI allows the modification of those files.

Log files

The current implementation of the multi-agent simulation can write log files in text format, and store them in a database. Figure 3 illustrates the current design of the database that stores execution data and outcomes.

EC3 can write outputs and execution data in text files.

Outcomes repository

In the current implementation of the simulation outcomes are part of the data recorded in the log files and in the database (See Figure 3).

These data are used to generate charts and test the execution of the simulation.

III.B.4. Decision making module

This module implements the behavior of the agents being simulated, i.e. airlines and flights. It can be changed from experiment to experiment.

From the implementation perspective, the module is distributed among the agents and embedded in their source code.

III.B.5. Visualization module

The visualization allows the presentation of results to the user. It eases the analysis by displaying charts, tables, messages, or any other useful visualization devices.

This module allows the control of the Evolutionary Algorithm, and the multi-agent simulation.

From the implementation perspective, it is distributed among other components of the architecture (see the composition relation in Figure 2).

III.C. Definitions

III.C.1. The definition of cost

The performance of the system being simulated is measured using two *criteria* or *objective functions*.

The first criterion is the *cost*. For the interest of this research, cost is operative: it is a function of the fuel consumption of flights, i.e. *airborne cost*, and of the time the flight waits on the ground, i.e. *ground cost*.

It is assumed that when the flight is waiting on the ground it has either the engines off or idling to use only little fuel.

Equation 3 shows how the airborne cost of a flight is computed during the simulation. In the equation, d_r is the distance of the route to be flown by the flight. The factor f_f is a conversion factor from distance slots (used by d_r) to cost units.

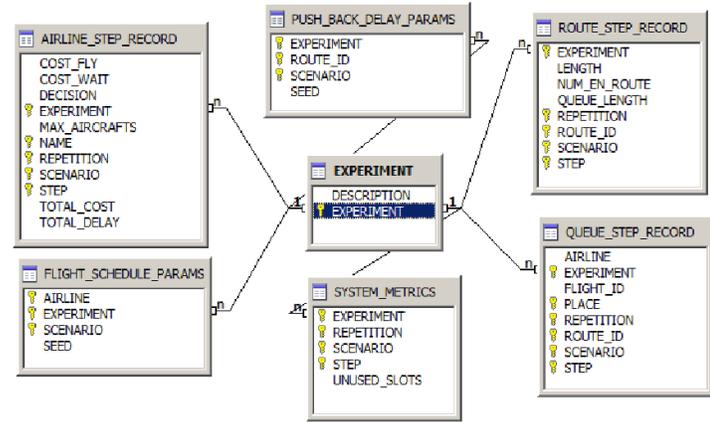


Figure 3. Data log and outcomes database diagram.

$$cost_{airborne} = f_f \cdot d_r \quad (3)$$

Equation 4 shows how the ground cost of a flight is computed during the simulation. In the equation, t_r is the time a flight waits on the ground after pushing back from the gate and before starting to fly. The factor w_f is a conversion factor from time slots (used by t_r) to cost units.

$$cost_{ground} = w_f \cdot t_r \quad (4)$$

The cost of the i -th flight is computed with Equation 5. Where the sub-index “ i ” is used to identify each flight.

$$cost_{flight_i} = cost_{airborne_i} + cost_{ground_i} \quad (5)$$

The cost for the j -th airline, called *airline cost*, during the execution of a simulation is computed with the Equation 6. Where the sub-index j is used to identify an airline.

$$cost_{airline_j} = \sum cost_{flight_i} \quad (6)$$

The cost across all airlines in a simulation, called *aggregated airline cost*, is computed using Equation 7. This is the performance metric of the system used in the simulations.

$$cost_{aggregated} = \sum cost_{airline_j} \quad (7)$$

III.C.2. The definition of utilization

Intuitively, utilization is defined as number of unused route slots in the execution of a simulation.

At each time step of the simulation, all the flights for which the departure time is reached must choose a route and enter the entry queue of the route. If a route could have been choose, but it is not chosen adds one unused slot to the *total unused slots*. This is the utilization metric of the system used in the simulations.

The implication of this algorithm is that unused slots are only counted once, even if they remain empty in subsequent time steps.

III.D. The airline, flights, and route behaviors

In these experiments the airlines do not show active behavior. They only have two real numbers, $0 \leq \beta \leq \alpha \leq 1$, as parameters. These two parameters define the valid strategies of the flights associated to the airline.

Flights have a single parameter, φ . The value of this parameter is determined by the condition $\beta \leq \varphi \leq \alpha$. Intuitively, this parameter determines the importance the flight gives to the airborne cost. Each flight decides which route to choose when its schedule departure time plus some push back delay is reached. All flights use Equation 8 as decision function to choose routes. It is Equation 5 expanded and with the parameter φ introduced.

In Equation 8, d_r is the distance of the route “r” in distance slots; t_r if the length of the entry queue of route “r” in time slots in the moment of the choice; f_f is a conversion factor from distance slots to cost units; w_f is another conversion factor from time slots to cost units.

Equation 8 computes the route “r” with the minimal predicted cost for the flight.

$$r = \underset{r}{\operatorname{argmin}} [(\varphi \cdot f_f \cdot d_r) + ((1 - \varphi) \cdot w_f \cdot t_r)]. \quad (8)$$

Neither the distance nor the length of the queue are known to the pilots in the current NAS. An assumption in these experiments is that the information is perfect, i.e. real-time and exact, and available to all the flights.

III.E. Considerations for the experimental design

III.E.1. Stochastic behavior and multiple executions

The NAS is a stochastic system. There are several processes in the NAS that cannot be simulated in a deterministic way, e.g. weather, arrival and departure delays. Also the behavior of humans involved in the system is not deterministic.

As a reflex of this, the simulation has some sources of stochastic behavior. One source of stochastic behavior is the introduction of the *pushback delays*. A pushback delay occurs when a flight scheduled to push back from its gate at predetermined time does not push back at that time. Most frequently, flights show small positive delays, they push back late, but only a few minutes. Very rarely, flights push back before the scheduled time, i.e. negative delay. There are cases in which the delay is more significant, but these cases are not frequent.

Studies of the pushback delay show that it is exponentially distributed, with a λ of about 25 minutes, but the value changes from airport to airport.

Another source of stochastic behavior results from the decision making of the agents. In the experiments explained in a section below, flights chose routes based on *predicted costs*. If two or more routes have equal predicted costs, the flight breaks the tie by randomly selecting any of the tied routes.

Later, with the introduction of the Evolutionary Algorithm, another source of stochastic behavior arises. The solutions found by the Evolutionary Algorithm are rarely equal from execution to execution.

A more implicit source of stochastic behavior comes from the multi-agent simulation platform used during the experiments, i.e. MASON. In each iteration of the simulation, MASON processes the agents in random order. So, a flights could see different current situations on each execution depending on whether it has real-time information or not.

In future experiments, any simulation or Evolutionary Computation Platform used in the experiments would bring similar sources of stochastic behavior.

The experimental design must consider these sources of stochastic behavior. Each scenario must be executed several times and only mean values and standard deviations should be reported as results.

The exact number of times to execute a scenario should be computed considering execution time constraints, as well as statistical criteria, e.g. required confidence intervals.

III.E.2. Validation of results

Castle and Crooks¹⁷ recommend to perform the so-called *internal validation* when the system involves a multi-agent simulation. The internal validation is a form of *unit testing* of the simulation. The assumption behind this validation is that if the parts of the system work well separated, there is a higher probability for them to work well together.

This approach has been used informally during the experiments, and it will be used more formally in future experiments.

The research is evaluating plausible futures after the implementation of NextGen. The implementation of NextGen is still mostly in the future and the results of the experiments are *forecasts*. So, there are no data to compare to the results of the simulation.

The approach in this research is to design experiments that are statistically sound. This will build more confidence in the results of the experiments. This is another form of internal validation.

As external validation researchers usually build their simulations and train them on historic data. If the simulation is able to “predict” the past then, the researchers assume, forecasts of the future will also be correct.¹⁸ In a system like the NAS there is no guarantee of that premise to be true.

Another approach for validation purposes is to present the results to an expert in the field and ask for advice and opinions on the plausibility of the results. This approach has been used in the experiments and will be used in future experiments too.

IV. Experiments and results

Figure 4 is the Universal Modeling Language (UML) class diagram of the implementation, in MASON, of NAS simulator for the experiments. The classes in red are “active agents”, i.e. they are part of the situation being simulated. The airlines of the simulation are actually *PolicyAirline* objects. The flights are *FlightAgent* objects. A single *Universe* object represents the environment and the shared pool of information. The blue classes are support agents or classes that only provide services to the simulation. For instance *SystemInspector* computes statistics of the simulation, and *DatabaseHandler* writes the data to the database if it is used.

IV.A. Experiment #1: Fixed homogeneous airline strategies

The experiment #1 includes no Evolutionary Computation, i.e. there is no adaptation. The results of this experiment are presented in more detail in Calderón-Meza et al.¹⁹ In the present study the charts of cost and the charts of under-utilization in the former study are merged into one Pareto surface chart. Also, not all the Pareto surfaces are shown in the present study, but only one for the cost ratio of 0.3 and the distance ratio of 0.68.

In this experiment, all the airlines use the same pair of parameters, and the two values of the parameters are equal, i.e. $\beta = \alpha = \varphi$. Effectively, all flights in the experiment behave the same way in each execution. The behavior of the flights is fixed during each execution of the experiment.

The experiment is divided into scenarios. Each scenario is defined by a combination of three parameters: the *Airline/Flight behavior* defined by $\alpha = \beta = \varphi$, the *distance of the Route₀*, and the *ground to airborne cost ratio* (See Table 1). Since there are sources of stochastic behavior, e.g. the pushback delays, and random resolution of cost ties during the decision making of the agents, each scenario is repeated 30 times. The results of each repetition is recorded, but only the mean and standard deviation are reported for a scenario.

The parameters vary as follows. The behavior parameter, i.e. $\alpha = \beta = \varphi$, simply identified by φ can take the values 0.0, 0.10, 0.20, 0.33, 0.50, 0.67, 0.75, 1.00. The value 0.00 represents the strategy in which all the flights minimize their ground cost. The value 1.00 represents the strategy in which all the flights minimize the airborne cost.

The Route₀ distance varies from 39 to 75 slots in increments of 2 slots for a total of 19 different values. The distance of the Route₁ remains constant at 57 slots.

The ground to airborne cost ratio can take three different values: 0.3, 0.5, and 0.7. The meaning of this value is that the cost of waiting on the ground one time slot has a cost that is 0.3, 0.5, or 0.7 times smaller than the cost of flying one distance slot. The ratio might change to simulate conditions in which, for instance, the cost of fuel is higher.

The experiment has a total of $8 * 19 * 3 = 456$ scenarios. Each scenario is repeated 30 times for a total of 13680 executions.

Table 1. Experimental design for a single scenario of Experiment #1

Scenario					
Airline / Flight behavior ($\alpha = \beta = \varphi$)					
Route ₀ distance (slots)					
Ground to airborne cost ratio					
Repetition	Airline cost (cost units)				Total unused slots (slots)
	Airline ₁	Airline ₂	Airline ₃	Aggregated	
1					
2					
...					
29					
30					
Mean					
Standard deviation					

*Route₁ distance constant at 57 slots.

Figure 5 shows a subset of the results. The charts describes the results when the distance ratio is 0.68 (i.e. Route₀ = 39) and the ground to airborne cost ratio is 0.3. The horizontal axis chart is the aggregated airline cost in cost units. The vertical axis is the total unused slots. The curve shows whiskers, or error bars, for the two criteria (horizontal and vertical axes). The curve is a *Tradeoff surface* or *Pareto surface* for this particular set of parameters. The shape of the curve is typical of problem in which the two criteria are being minimized.²⁰ Curves for other combinations of parameters show similar shapes.

When the control parameter, i.e. the behavior parameter φ , is 1.00 all the flights are minimizing their airborne cost. Then they all choose to fly the shortest route. The longest route remains unused. Since there are a total of 90 scheduled flights, there are also 90 unused slots corresponding to the slots of the longest route. But the aggregated airline cost is minimal.

When φ takes the values 0.50, 0.67, and 0.75, the results are equal the case when $\varphi = 1.00$. Because of the value of the ground to airborne cost ratio, when the flights start giving more importance to the airborne cost, this cost dominates the whole decision process, and the ground cost becomes irrelevant. So the results are equal in all those cases.

When $\varphi = 0.00$ all the flights are minimizing their ground cost. Flights chose whichever route has the shortest wait queue. The result is that traffic is balanced between the routes: utilization is minimal in this case. But cost is higher because half of the flights are flying the longest path.

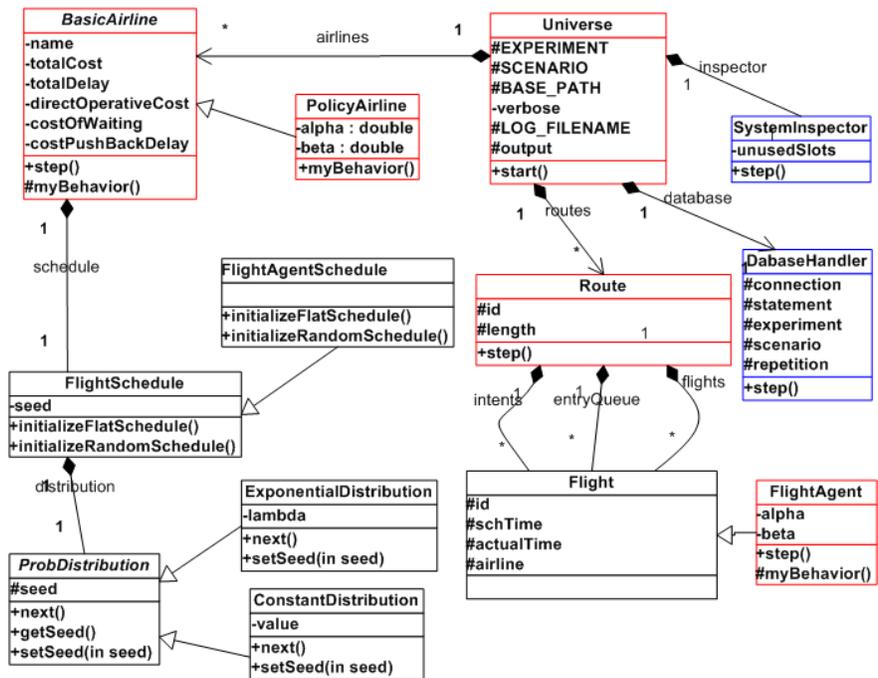


Figure 4. UML class diagram of the NAS simulator, in MASON, for the experiments.

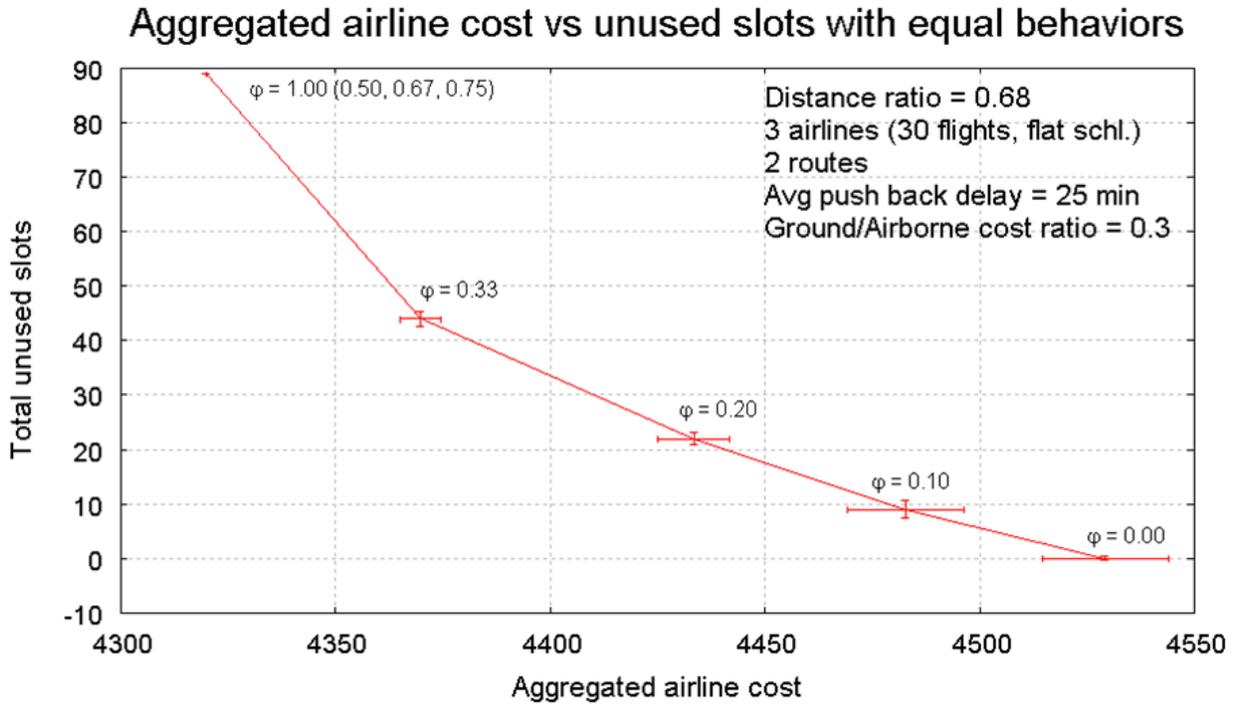


Figure 5. Tradeoff aggregated airline cost vs. unused slots as the result of homogeneous airline strategies.

The lessons learned from this experiment are that the variances are small enough to say that parameter that control the behavior of the flights is effective to control the aggregated airline cost and the utilization. Also, there is an observable tradeoff between the utilization and the aggregated airline cost.

IV.B. Experiment #2: Heterogeneous behaviors and adaptation

The experiment #2 includes adaptation of the flight behaviors. The adaptation is achieved by an Evolutionary Algorithm that uses the multi-agent simulation to evaluate possible configurations of behaviors.

In this experiment, each airline uses the same values for its two parameters, i.e. $\alpha = \beta$. But the value is different for each airline, i.e. $\alpha_1 \neq \alpha_2 \neq \alpha_3$. As before, all the flights scheduled by an airline behave the same way, e.g. $\varphi_1 = \alpha_1 = \beta_1$ for airline 1. So the behavior of an airline is represented by its *varphi* value.

The goal of the experiment is to compute the Pareto surface for a situation in which the behavior of the flights (airlines) is not homogeneous. Then the utilization and aggregated cost pairs must be ranked in some way to be compared to other values. As explained in the literature review, there are several ways to rank vectors of results. In this experiment the weighted sum of the results is used. An explanation of this technique is given by Fonseca and Fleming.¹¹

The objective function for the ranking of solutions is shown in Equation 9. The this equation *cost* is the aggregated cost, *unused_slots* in the total unused slots, and ω is the weight. Intuitively, ω is the importance or *social value* given to the aggregated cost. It is a real number limited to the range $0 \leq \omega \leq 1$.

$$f = -(\omega \cdot cost + (1 - \omega) \cdot unused_slots). \quad (9)$$

This ω is one more, continuous, parameter of the experiment.

The flights make decisions based on the function shown in Equation 8. But, the values of φ are modified during the experiments in search for “the best” combination of behaviors determined by the objective function.

All the other parameters, ranges, and conditions remain the same of the simulation.

The Evolutionary Algorithm brings with it more parameters. Table 2 shows the names and values of these parameters. It is not the goal of these experiment to determine the best combination of parameters for the Evolutionary Algorithm. So, its parameters remain constant during all the experiment.

Table 2. Parameters of the Evolutionary Algorithm for experiment #2

Parameter	Value
Population size	15 individuals (15 offspring)
Population dynamics	Overlapping populations
Parent selection	Fitness proportional
Survival selection	Fitness proportional
Operators	Mutation and recombination
Iterations	300 (birth counts)

The experiment is divided into scenarios. Each scenario is defined by a combination of six parameters: the *Airline/Flight behavior* for airline₁ defined by $\alpha_1 = \beta_1 = \varphi_1$, the *Airline/Flight behavior* for airline₂ defined by $\alpha_2 = \beta_2 = \varphi_2$, the *Airline/Flight behavior* for airline₃ defined by $\alpha_3 = \beta_3 = \varphi_3$, the *distance of the Route₀*, the *ground to airborne cost ratio*, and the weight defined by ω . To simplify the analysis, the distance of Route₀ remains constant at 39 slots, and the ground to airborne cost ratio also remains constant at 0.3. (See Table 3). Since there are sources of stochastic behavior, e.g. the pushback delays, and random resolution of cost ties during the decision making of the agents, each scenario is repeated 30 times. The results of each repetition is recorded, but only the mean and standard deviation are reported and used as inputs to the objective function for a scenario.

Table 3. Experimental design for a single scenario of Experiment #2

Scenario								
ω								
Route ₀ distance (slots)								
Ground to airborne cost ratio								
Repetition	Best individual			Airline cost (cost units)				Total unused slots (slots)
	φ_1	φ_2	φ_3	Airline ₁	Airline ₂	Airline ₃	Aggregated	
1								
2								
...								
29								
30								
Mean								
Standard deviation								

*Route₁ distance constant at 57 slots.

The parameters vary as follows. The behavior parameters are varied automatically by the Evolutionary Algorithm. The algorithm executes the simulation 30 times with each individual and it reported the mean and standard deviation as outcomes.

The ω varies from 0 to 1 in increments of 0.1.

The experiment has a total of 15 individuals, i.e. scenarios, per generation of the Evolutionary Algorithm. Each scenario is repeated 30 times. There is a limit of 300 births for a total of 20 generations. Since ω can take 11 different values, the total number of executions when the distance and cost ratios are constant is $15 * 30 * 20 * 11 = 99000$ executions of the simulation. If the distance and cost ratios also vary, the total number of executions is $99000 * 19 * 3 = 5643000$.

Table 4 shows the results for experiment #2 when the distance and cost ratios remain constant at 0.68 and 0.3. In nine of the eleven results one of the airlines uses a value of φ that is much higher, in italics, than the values of the other two airlines, i.e. this airline is more aggressive minimizing the its airborne cost. The other two airlines prefer to minimize the ground cost. In the last two rows, when the importance of aggregated cost is higher, ω approaches 1.0, the results are different. In these rows two airlines, also in italics, more aggressively minimize the airborne cost, and only one minimizes the ground cost.

Table 4. With heterogeneous strategies, more aggressive airlines in terms of airborne costs obtain lower individual cost. Further reductions in the aggregated cost require two aggressive airlines and result in under utilization.

Distance ratio								0.68
Ground to airborne cost ratio								0.30
ω	Best individual			Airline cost (cost units)				Total unused slots (slots)
	φ_1	φ_2	φ_3	Airline ₁	Airline ₂	Airline ₃	Aggregated	
0.0	<i>0.756</i>	0.009	0.025	<i>1248.6</i>	1757.7	1511.7	4518.0	1
0.1	<i>0.756</i>	0.009	0.025	<i>1249.5</i>	1757.7	1510.8	4518.0	1
0.2	<i>0.888</i>	0.010	0.027	<i>1248.9</i>	1758.3	1510.8	4517.4	1
0.3	<i>0.565</i>	0.020	0.031	<i>1252.8</i>	1632.3	1632.9	4516.0	1
0.4	<i>0.565</i>	0.020	0.031	<i>1254.0</i>	1632.0	1632.0	4516.8	1
0.5	0.010	<i>0.115</i>	0.011	1632.6	<i>1249.2</i>	1636.2	4516.3	1
0.6	0.002	0.016	<i>0.436</i>	1634.1	1634.1	<i>1249.8</i>	4515.7	1
0.7	0.014	<i>0.884</i>	0.010	1636.5	<i>1248.0</i>	1633.5	4516.8	1
0.8	<i>0.508</i>	0.028	0.031	<i>1252.2</i>	1630.2	1636.6	4512.9	1
0.9	<i>0.244</i>	<i>0.520</i>	0.021	<i>1378.2</i>	<i>1303.5</i>	1701.3	4383.1	19
1.0	<i>0.716</i>	<i>0.651</i>	0.034	<i>1322.7</i>	<i>1322.1</i>	1675.2	4320.4	33

*Route₁ distance constant at 57 slots.

The individual airline cost is always less for the airlines that are aggressive in flying the shortest route (in italics) than for airlines that minimize the ground cost. Using an aggressive strategy brings less cost as a benefit for the airline.

The aggregated cost decreases as ω increases. Since ω is the importance of the aggregated cost in the objective function (see Equation 9), this is the expected behavior.

In the search for lower aggregated cost, there is a point from which the cost is reduced by sacrificing utilization. In the first nine rows of the table, there is only one unused slot, but in the last two (in bold) there are 19 and 33 unused slots. These results intuitively indicate that there is a tradeoff, but finding it depends on the value of ω .

The values in the table are examples of the best individuals in each scenario. There are other different individuals, not reported in the table, with similar outcomes. The goal of the experiment is to approximate the Pareto surface. So, the Evolutionary Algorithm must be extended so that it can monitor the population and record all the best individuals found during the execution. This set of individuals have a higher probability of being on the Pareto surface. This extension was not implemented in this experiment.

The lessons learned from this experiment are that normalization of the outcomes or criteria is needed. The value of utilization is much smaller than the value of aggregated airline cost. Since the criteria are linearly combined, the aggregated airline cost tends to dominate the objective function. The computation of scaling values is no easy task and was not done in this experiment. Instead, fixed values of 5000 and 3510 were used as bounds for the aggregated airline cost. For the utilization the bounds are 0 and 90. Some work has been done after the experiment to create a formula to compute the lower limit of the cost. Theoretically, there is no upper limit since the pushback delays are exponentially distributed and this distribution has no upper bound.

Another lesson learned in this experiment is that there is the need to implement the Hall of Fame as a means to record all the best individuals. The changes to the EC3 are proposed and being implemented.

Finally, the use of the weighted sum as an objective function is subjective. The importance of cost is actually a social value, it can change from community to community, from agent to agent. In fact, almost any result can be achieved by properly changing the value of ω . Because of this, another experiment was attempted.

V. Conclusions

This study focuses on a simplified version of the problem of trans-continental route selection by the Airline Operations Centers in the presence of NextGen, i.e. Net-Centric Operations. Two multi-agent simulation experiments were carried out to forecast consequences of the implementation of NextGen in the route selection problem.

The agents select the route with minimal *predicted cost* of the route. The cost of a route is the combination of airborne and ground cost. The selection strategy of an agent is controlled by one parameter. The agent can go from selecting only the shortest route, to flying only the route with the least ground (wait) time.

The simulation has two outcomes: the performance, and the utilization. The performance is the aggregated airline cost: total of the actual costs of the flights across all airlines. The utilization is measured by the under-utilization: the total number of route slots not used by any flight when slots could have been used.

Experiment #1 uses fixed *homogeneous strategies*: all agent strategies have the same value for their parameter. The results of the experiment show that there is a performance vs. utilization *tradeoff*. When the flights choose to fly the shortest route, the aggregated airline cost is low, but the under-utilization is high. When the flights choose not to wait on the ground, the cost becomes higher, but the under-utilization is minimal, i.e. the traffic is balanced between the routes. In the *absence of rules* about utilization performance considerations by the airlines lead to *under-utilization* of the system: case when all the agents choose only the shortest route. If the utilization rules are too restrictive, e.g. forcing perfect traffic balance, the cost of the system rises.

Experiment #2 included an Evolutionary Algorithm to adapt the airline strategies independently, i.e. *heterogeneous strategies*. The results exhibit *emergent behavior*. Airlines adapt to use strategies of different type. Some of the airlines use a “hawk”, i.e. dominant, strategy in terms of reduction of cost by flying the shortest route. While other airlines use a “dove” strategy, reducing cost by choosing the route with the shortest ground wait time. The airlines with the “hawk” strategy achieve the lowest airline cost among the airlines, and the doves receive higher costs. This shows the *lack of equity* when there are no rules to limit the behavior of the airlines.

VI. Future work

At this stage, extensions and refinements are necessary to compute the Pareto surface, i.e. find the tradeoff, in the case of heterogeneous strategies. The Evolutionary Algorithm must be extended to record all the best individuals found in an execution and report them as results.

The method ranks the combinations of strategies must also be improved, the literature offers several possibilities that must be explored.

It is also future work to analyze and compare more advanced tools to simulate the NAS and to support evolutionary multi-objective optimizations. With these tools more realistic situations could be analyzed.

Several assumptions and premises used in the experiments could be relaxed or modified to make the experiments more realistic. For instance, the schedules of the airlines could be made more complex and real. More airports and flights could also be included in the simulation. More complex routes, i.e. with more merge and branching points could also be used. The last modification would require a redefinition of the concept of utilization.

Acknowledgments

The authors would like to acknowledge the contributions and help from the following persons and institutions.

The research this study is part of is funded by NASA (NRA NNA07CN32A). Furthermore, María Consiglio, Brian Baxley, and Kurt Neitzke, all NASA employees, have provided suggestions and comments throughout the whole research process.

George Hunter and Huina Gao from Sensis Corporation have also provided suggestions and comments through the whole research process.

From George Mason University the authors would like to acknowledge the help from Dr. Thomas Speller Jr., Dr. Kenneth De Jong, Dr. Robert Axtell, Dr. George Donohue, Dr. John Shortle, Dr. Rajesh Ganesan, Maricel Medina-Mora, John Ferguson, and Keith Sullivan. They have all contributed to the improvement of

the research in their particular areas of expertise.

From Metro Aviation, the authors acknowledge the help and contribution of Dr. Terry Thompson and Norm Fujisaki.

From FAA, the authors acknowledge the help and contribution of Joe Post and Tony Diana.

Finally, thanks to the Ministerio de Ciencia y Tecnología (Minister of Science and Technology) of Costa Rica.

References

- ¹JPDO, *Concept of Operations for the Next Generation Air Transportation System, Version 2.0*, Joint Planning and Development Office, Washington DC, USA, June 2007.
- ²Alberts, D. S., *The Unintended Consequences of Information Age Technologies: Avoiding the Pitfalls, Seizing the Initiative.*, National Defense University Press, Washington, DC, 1996.
- ³Alberts, D. S., *Information Age Transformation. Getting to a 21st Century Military*, Information Age Transformation Series, Command & Control Research Program, Washington DC, April 1996.
- ⁴Cebrowski, A. K. V. U. and Garstka, J. J., "Network Centric Warfare: Its Origin and Future," *Proceedings of the Naval Institute*, Vol. 124, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, January 1998, pp. 28–35.
- ⁵Ball, M. O., Hofman, R. L., Knorr, D., Wetherly, J., and Wambsganss, M., "Assessing the Benefits of Collaborative Decision Making in Air Traffic Management," *Air Transportation Systems Engineering*, edited by G. L. Donohue and A. G. Zellweger, Vol. 193, American Institute of Aeronautics and Astronautics, USA, 2001, pp. 239–250, ISBN 1-56347-474-3.
- ⁶Wojcik, L. A., "Airline Personalities and Air Traffic Flow Management: A Simple Agent-Based Model," *AIAA 4th Aviation Technology, Integration and Operations Forum*, Chicago, Illinois, USA, September 2004, pp. 1–16.
- ⁷Campbell, K. C., Jr., W. W. C., Greenbaum, D. P., and Wojcik, L. A., "Modeling Distributed Human Decision Making in Traffic Flow Management Operations," *Air Transportation Systems Engineering*, edited by G. L. Donohue and A. G. Zellweger, Vol. 193, American Institute of Aeronautics and Astronautics, USA, 2001, pp. 227–237, ISBN 1-56347-474-3.
- ⁸Waslander, S. and Tomlin, C., "Efficient Market-Based Air Traffic Control Flow with Competing Airlines," *45th IEEE Conference on Decision and Control*, San Diego, California, USA, Dec. 2006, pp. 2843–2848.
- ⁹Maynard Smith, J. and Price, G., "The Logic of Animal Conflict," *Nature*, Vol. 246, November 1973, pp. 15–18.
- ¹⁰Luke, S., Cioffi-Revilla, C., Panait, L., and Sullivan, K., "MASON: A New Multi-Agent Simulation Toolkit," *Proceedings of the 2004 SwarmFest Workshop*, Michigan, USA, 2004.
- ¹¹Fonseca, C. M. and Fleming, P. J., "An Overview of Evolutionary Algorithms in Multiobjective Optimization," *Evolutionary Computation*, Vol. 3, No. 1, 1995, pp. 1–16.
- ¹²Deb, K., "Multi-objective Genetic Algorithms: Problem Difficulties and Construction of Test Problems," *Evolutionary Computation*, Vol. 7, No. 3, 1999, pp. 205–230.
- ¹³Jong, K. A. D., *Evolutionary Computation. A Unified Approach*, The MIT Press, Massachusetts, USA, 2006, ISBN 0-262-04194-4.
- ¹⁴Fogel, L. J., Owens, A. J., and Walsh, M. J., *Artificial Intelligence through Simulated Evolution*, John Wiley & Sons, New York, USA, 1966.
- ¹⁵Holland, J., *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor, Michigan, USA, 1975.
- ¹⁶Schwefel, H. P., *Evolutionsstrategie und numerische Optimierung*, Ph.D. thesis, Technical University of Berlin, Berlin, Germany, 1975.
- ¹⁷Castle, C. J. E. and Crooks, A. T., "Principles and Concepts of Agent-Based Modelling for Developing Geospatial Simulations," *Working Papers Series*, University College London Centre For Advanced Spatial Analysis, London, UK, September 2006, pp. 1–60.
- ¹⁸Niedringhaus, W. P., "The Jet:Wise Model of National Airspace System Evolution," *Simulation*, Vol. 80, No. 1, January 2004, pp. 45–60.
- ¹⁹Calderon-Meza, G., Sherry, L., and Consiglio, M., "Multi-agent simulation of NAS infrastructure utilization in the presence of airline policies," *Integrated Communications, Navigation and Surveillance Conference, 2009. ICNS '09.*, May 2009, pp. 1–7.
- ²⁰Collette, Y. and Siarry, P., *Multiobjective Optimization. Principles and Case Studies.*, Springer-Verlag, 2003.