An Arrival Scheduling Model for Incorporating Collaborative Decision-Making Concepts into Time-Based Flow Management

Yeming Hao[1], David J. Lovell[1,3], Michael O. Ball[2,3]
[1] Department of Civil and Environmental Engineering
[2] R.H. Smith School of Business
[3] Institute for Systems Research
University of Maryland
College Park, MD, USA
{yhao, lovell, mball}@umd.edu

Sergio Torres, Gaurav M. Nagle
Leidos, Inc.
Gaithersburg, MD, 20878
{sergio.torres, gaurav.nagle}@leidos.com

Abstract—This paper proposes a flight scheduling scheme – 2-opt-swap, which assigns controlled times of arrival (CTAs) for flights reaching the Freeze Horizon and allows certain slot swapping between different flights with the goal of reducing total controlled arrival delay cost over all carriers. The allowable swaps are predicated on models of carrier preferences following a Collaborative Decision-Making paradigm. Monte Carlo simulations were designed to prove the benefits of this new CTA scheduling scheme, compared to a baseline model of first-come-first-served discipline, which is currently used in Time-Based Flow Management.

Keywords- collaborative decision-making; flow management; Time-Based Flow Management

I. INTRODUCTION

The purpose of the research described in this paper was to explore the potential savings that can be realized in managing flight arrivals at busy airports. Currently, the FAA uses Time-Based Flow Management (TBFM) to arrange arrivals to an airport so flights arrive in an orderly fashion consistent with airport capacity and with adjusted inter-arrival spacing so landing services can be offered safely and efficiently. The TBFM system uses a first-come-first-served (FCFS) protocol to determine the Controlled Times of Arrival (CTAs) for flights at various points in the national airspace system. This method of scheduling is solely based on the Estimated Times of Arrival (ETAs) and does not take into consideration user preferences. The focus of the research is to investigate opportunities for borrowing collaborative decision-making (CDM) concepts from other air traffic management regimes and adapting them as a means of incorporating more carrier preferences into the TBFM process.

For example, if an airline gets assigned a pair of slots by TBFM at a metering location, the airline does not have the ability to indicate which flights should go into each slot based on factors such as cost and number of connecting passengers. This project focused on schemes that could be applied after the initial FCFS scheduling to allocate CTA slots to airlines based on the cost of assigning each flight a particular slot. Cost considerations include factors such as the size of the aircraft, a dollar value/minute/seat, and the difference between the flight's CTA and ETA. In addition to exploring different scheduling approaches, we also analyzed the operational envelope where queue optimization approaches present the highest benefit. The latter is a key contribution to the advancement of knowledge in the CDM field.

In real-world conditions, there are various perturbations that a flight might encounter from the time a landing slot has been assigned (per the CTA) to the time it enters the final approach. Because of this 'meet time error', the flight may incur additional delay, which we refer to as the service delay. Since the flight deviated from its CTA, the final approach and tower controllers may need to send the flight into a delay absorption deviation or holding to accommodate the flight among the other flights competing for the same resources. Therefore, it is important to understand the extent to which potential optimization gains are subsequently erased due to these CTA meet time errors.

The analysis described in this paper was conducted in a Monte Carlo simulation environment coded in MATLAB. It was based on a simulation package originally designed by Leidos, Inc. and described in [1]. The simulation was improved to accommodate the CDM schemes described above, as well as to incorporate more realistic models of en route perturbations, winds, and other factors.

II. RELATED WORK

Collaborative Decision Making (CDM) was developed in the mid-1990s for the planning and control of ground delay programs (GDPs). See [2] for early references. Fundamental principles of CDM include rich information exchange among all relevant parties and the use of resource allocation mechanisms that encourage truthful information sharing. Further, the decision and control architecture should enable air carriers to make those decisions that involve only their own resources, should seek to treat all operators equitably and, to the extent possible, should infuse carrier preference information within FAA allocation processes. Vlachou and Lovell [5] modeled a compact vocabulary for carrier priorities and acceptable flight delays that could be incorporated in a fair air resource allocation procedure such as that described by Vakili [6]. Raj [7] modeled the use of a similar vocabulary in the Collective Trajectory Options Program, with additional controls on the vocabulary to prevent strategic gaming.

Given that CDM problems involve complicated exchanges between carriers, various methodological approaches have been adopted to solve these problems. In particular, integer programming models are useful approaches to finding optimal solutions. Relaxation or heuristic approaches are usually adopted for real time solutions [8][9][10][11][12].

A recent paper by Idris et al. [13] has analyzed modifying the TBFM allocation algorithm based on an accrued delay prioritization scheme. While this approach does not explicitly employ user preference information, it does lead to more equitable solutions, a key CDM principle.

Our goal in this paper is to demonstrate the potential benefit of the development of a CDM mechanism that infuses carrier cost/preference data into TBFM slot assignment. Specifically, we modify the TBFM allocation algorithm to take into account carrier cost information. We develop a simulation model that quantifies the improvements achieved by our modifications over using the standard TBFM FCFS mechanism. Follow on research is required to develop a specific CDM mechanism that can provide more carrier information to achieve the benefits demonstrated by this research. Such a mechanism should be fair and should encourage airlines to provide truthful information.

III. Evaluation Approach

In this section, we describe the data sources used to feed the simulation model, and the architecture of the simulation itself.

A. Data Preparation

The data inputs are broken into a table of flights, a table of aircraft seats by aircraft type, a table of aircraft speeds by aircraft type, a table of wind speeds, a table of Airport Acceptance Rates (AARs), and other data. These are described in more detail in the following subsections.

1) Flights Table

A single run of the simulation model considers an entire day's worth of flights arriving to a given airport. As our study airport, we chose Atlanta Hartsfield-Jackson International Airport (ATL) due to its high traffic volumes and extensive use of TBFM for approach traffic smoothing. The traffic data were taken from Monday, October 7, 2019, which was a typical traffic day into the airport with no severe weather conditions. There were 1293 incoming flights used for the analysis (after removing a few for which only incomplete data were available).

The data fields in the flight data set include carrier identification, flight ID, aircraft type, and scheduled and actual landing times. All timing data used in the project were converted to the number of seconds past midnight, Eastern Daylight Time.

It is important to note that because these are historical data, the actual landing times were accomplished after a TBFM process was applied to smooth the incoming flows. Thus, their inter-flight headways and spacings are more regular than would be expected of flights arriving to a metering fix. On the other hand, the conditions encountered by the metering system at the time the CTAs are allocated are different because it deals with flights before pre-conditioning (that is, before they reach the freeze horizon, FH). In this simulation, the freeze horizon was assumed to be 250 nm upstream of the metering fix.

Before arrival queue preconditioning, aircraft do not cross the FH at regular intervals; instead, the time at the FH depends on the past history of the flight (departure delays, pilot actions, wind, etc.). At the FH crossing time, therefore, the ETAs for these flights are random (not metered). This randomness in ETA is the condition we need to simulate. The inputs to the FCFS algorithm are precisely the ETAs known at the FH crossing time. This was remedied by adding a random time shift to the actual landing times (ATAs) to approximate the landing ETA at the FH. It is assumed that this delay is due primarily to a departure delay distribution. This presumes an offset for the time between metering fix and the runway, which should be very uniform. Analysis of the departure delay distribution at the ATL airport on 10/07/2019 shows an exponential distribution with a mean of 11 min. Historical data has shown that on average, arrival delay is less than departure delays [14]. We thus add an exponential distribution with a mean of 7 mins to the ATAs to simulate the arrival delays. There are other factors such as wind, pilot behavior, ATC actions affecting flights' ETAs, among which wind is the major factor. We thus refined this model by adding a Gaussian component to account for wind deviations. The previously described exponential distribution adds only a positive time shift, while this Gaussian component can add positive or negative noise to simulate the ETAs. $\sigma_{\text{wind}} = T/64$ was used as the standard deviation for this noise for each flight, where T is the time from departure to the FH, and is estimated as (total travel time – 38 min), with 38 min being the nominal FH to MFX flying time [15].

The objective of this operation is to reconstruct the state of the flow prior to applying the FCFS queuing rule to assign CTAs. The advantage of this approach is that the arrival flow is randomized based on an empirically calibrated distributions that reflect real world conditions.

2) Aircraft Number of Seats Table

A list of unique aircraft types was extracted from the flights table described previously. There are 37 unique aircraft types amongst the 1293 flights in our table. For modeling purposes, we assigned a single value for the typical number of seats for each aircraft type. This ignores some of the flexibility that is afforded some aircraft types for their cabin design, but such deviations are small and negligible for the purposes described here. These numbers are sufficient to represent large-scale differences between aircraft types. In a real-world application, airlines could provide their actual number of seats table. This information is important because it can serve as some representation of the "importance" of an aircraft in a modeling or practical scenario where occupancy information is not publicly available, or where other possible expressions of carrier preference are not feasible.

Most of the values assigned come from an aircraft overview booklet by DVB Bank SE [16]. For those aircraft not listed in this reference, additional information was found in FAA Aircraft Registry, FlyRadius, and Controller websites [17].

3) Speed Table

For each flight in the simulation, it is necessary to have three speed values at hand: the "actual" speed that a particular simulated flight will be assumed to follow, and the minimum and maximum speeds that it could be assigned. These speed extrema allow the model to determine the range of possible CTA slot assignments that are feasible for an aircraft, in the event that swapping flights between slots can improve the overall objective function.

a) Speed extrema

The values for the minimum and maximum possible speed for each aircraft type were determined empirically based on a margin of speeds that is operationally reasonable. For aircraft that can fly this high, a default of 35,000 ft. was chosen as the cruising altitude, presuming this is the nominal altitude at the freeze horizon. Historical average altitude was used for other aircraft. Although TBFM applies to a phase of flight where altitudes and speeds are expected to change as part of the descent, a constant speed assumption is sufficient to distinguish the aircraft types. In a real environment, actual speeds would be known, and it is possible that carriers could indicate desired speed extrema based on fuel consumption or other economic factors, rather than more liberal controller limits or aircraft operational envelopes.

For a limited number of aircraft that could not be categorized this way, we used the values recommended for similar aircraft and ICAO's DOC 8643 Designators [20].

b) Current speed

The current speed assigned to each flight is sampled from a distribution of historical flight speeds at the appropriate altitude for the corresponding aircraft type. A large sample of historical flight plans was analyzed to determine average speeds filed for cruise. The flight plans were submitted in August 2017. A sample of the analyzed statistics is shown in TABLE I.

TABLE I. FLIGHT PLAN SPEEDS BY AIRCRAFT TYPE

actype		*hnpt*	*havg (ft)*	*…*	*avgtas (nm/hr)*
A319		45,497	34,042		451
A320		71,861	33,894		456
A321		38,812	32,846		459
…		…	…	…	…

where *actype* = aircraft type (ICAO designator as filed);
hnpt = number of flight plans analyzed;
havg = average cruise altitude filed (in feet);
avgtas = average true airspeed (TAS) at altitude *havg*.

Note that some aircraft types cannot fly at 35,000 feet (B190, BE58, BE9L, C208, and PC12). For those aircraft types we cannot use the average speed at 35,000 feet. Instead, we use the average TAS at their historical average altitude.

The speed table was joined to the flight table to indicate for each flight its nominal, minimum, and maximum speeds. For the purposes of simulation, zero-mean Gaussian noise with standard deviation of 3.33 knots was added to the existing current speed distribution to approximate the flight's true current speed at the freeze horizon. This step allows different flights on the list to have slightly different current speeds even if they belong to the same aircraft type. With this choice of standard deviation, approximately 99.7% of the current speed values will be within 10 knots of the historical flight plan *avgtas* value for each aircraft type. A value check is conducted to make sure each aircraft's new current speed is within the speed limit range. If it violates the maximum or the minimum speed check, it is clamped to the closest boundary value.

c) Wind Speed

To simulate the real-world environment, it is important to take wind into consideration. The speeds indicated in the previous steps are all true air speed (TAS). In order to relate the results to the ground speed, the wind speed then needs to be incorporated.

To characterize wind speeds encountered by flights in cruise (near the freeze horizon) going into ATL, a large sample of ATL arrivals and the along-track wind for each trajectory were analyzed. The set of trajectories is shown in Figure 1. Figure 2 shows the histogram of along-track wind speed for the altitude bin [30,000 - 40,000] feet. Since Figure 1 shows clustering of trajectories, it is not surprising that there is also some bunching of speeds in Figure 2.

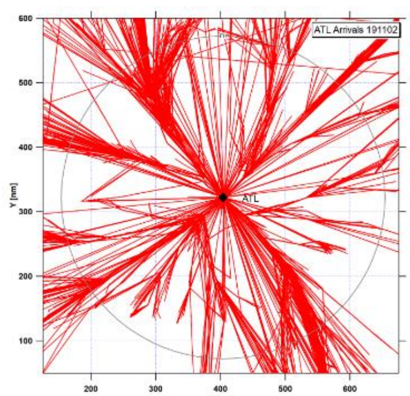

Figure 1. Trajectory Footprint for Selected ATL Arrivals (11/02/2019)

In an effort to keep the fidelity of the simulation at a manageable level, it was decided not to try to correlate trajectories to speed distributions, particularly because while these things were true for this particular historical day, something else would be true on other days. For example, there is a seasonal variation, with lower wind speed values expected in warmer air. The simplest way to simulate the wind distribution was to use a uniform distribution. We chose U(-70, 70) after observing the empirical data, but of course this could be changed. Wind speed was generated for each flight following this distribution and added to the three TAS columns in the flights table to represent the ground speed.

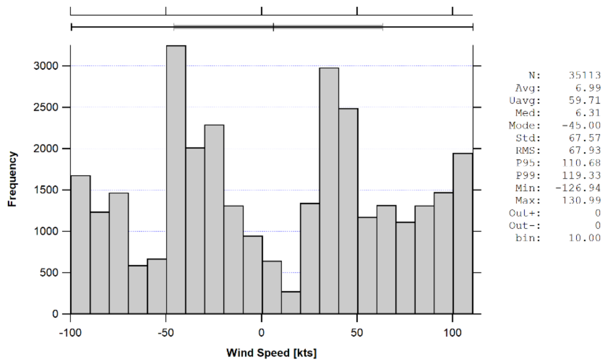

Figure 2. Along-track wind speed distribution for altitude bin [30000,40000] feet

Since the distance (along the queue) allowed for swapping is restricted by ground speeds, it is expected that the results of optimization swapping are sensitive to the wind speed.

4) Airport Arrival Rate (AAR)

The Airport Arrival Rate (AAR) is a dynamic parameter specifying the number of aircraft that an airport, in conjunction with terminal airspace, can accept under specific conditions throughout any consecutive 60-minute period [21]. The AAR for a given airport varies in different weather conditions and runway configurations. A full table for the ATL AAR was obtained from the FAA website [22] (a sample is shown in TABLE II.). Among all weather and runway configuration combinations, the maximum AAR of all conditions is 132 and the minimum AAR of all conditions is 18. We thus use a range of [20, 140] ops/hour with a step size of 10 ops/hour to simulate a full range of operational conditions at ATL.

TABLE II. ATL AAR INFORMATION

Runway configuration		***Meteorological Condition***			
arrival	***departure***	***VMC (3600/7)***	***LOW VMC***	***IMC***	***LOW IMC***
26R|27L|28	26L|27R	132	124	110	98
26R|27R|28	26L|27L	122	116	104	96
…	…	…	…	…	…
09L		50	45	40	36
Data source: https://www.fly.faa.gov/Information/east/ztl/atl/atl_aar.htm					

There were no extreme weather conditions in Atlanta on 10/07/2019; hence, our traffic data represent a high traffic scenario.

B. Simulation Environment

The starting point for developing the simulation software was a package called QSim, which was developed by Leidos, Inc. [1]. Qsim is, at its heart, a queueing engine designed to model traffic arriving to an airport, subject to a variety of forms of stochastic digressions from their schedule (this could be used to represent the effects of en route winds, departure delays, routing choices, etc.), and then to mimic the first-come-first-served (FCFS) process currently in TBFM. Once aircraft are metered into their appropriate slots, various performance metrics are generated.

It was determined that this package would be the appropriate starting point for a simulation capable of studying CDM-enhanced TBFM, because the data structures and queueing engine could be modified readily to reflect the novel circumstances under investigation. The package was ported to MATLAB because of its universal usage.

The components of Qsim used in this research are the Scheduler, Flight Propagator, Arrival Service, and Queue Analyzer. The Scheduler takes as its input the set of flight arrival times and assigns each flight an arrival slot with a corresponding CTA. When the CTA assigned is later than the ETA, a controlled arrival delay occurs. The scheduling process in the initial version of Qsim followed the FCFS rule, and it is in this module that modified CDM-based assignment strategies were implemented. The Flight Propagation module simulates the flights' arrival queue after the flights cross the freeze horizon by adding random noise to the target CTAs. The Arrival Service models the arrival service based on given AARs. It puts one flight at a time in the queue in the order it arrives and serves the next one in the queue when the server is freed up from the previous flight in the queue. When the time of service is later than the ATA, a service delay occurs. The Queue Analyzer collects the information from the previous steps and computes performance metrics such as average arrival delay caused by assigning CTAs, average service delay, inter-arrival time variance, unused capacity, and queue lag time. The Qsim process loop includes all these previous components and was repeated for 5000 realizations to generate the performance metrics.

We designed our simulations in this study based on the original Qsim variant, but with changes in the following aspects (shown in Figure 3):

1) We started our simulation with a flight arrival queue containing a list of individual flight ATAs instead of aggregated data. Time shifts and noise were added to the ATAs to mimic the ETAs following the procedures described in III.A.
2) In addition to the FCFS scheduling scheme (baseline model), we also developed a 2-opt-swap CTA scheduling scheme to reduce the total controlled arrival delay cost, the cost associated with delays due to flow management measures during arrival (illustrated in III.B. 1). The 2-opt-swap takes account of information including carrier information and aircraft number of seats and allows flights from different carriers to swap CTA slots to achieve a lower controlled arrival delay cost.
3) We present the monetary values for all the delays in the Queue Analyzer, which gives the CDM and TBFM participants a clearer view of how much they can save by adopting our schemes.
4) Our simulation incorporates more details to mimic more complicated and realistic scenarios. For example, speed limits and current speeds were used when defining the available CTA slots for each flight, and wind speed were generated for each flight to mimic the real-world scenarios.

A simulation system flow chart is shown in Figure 3.

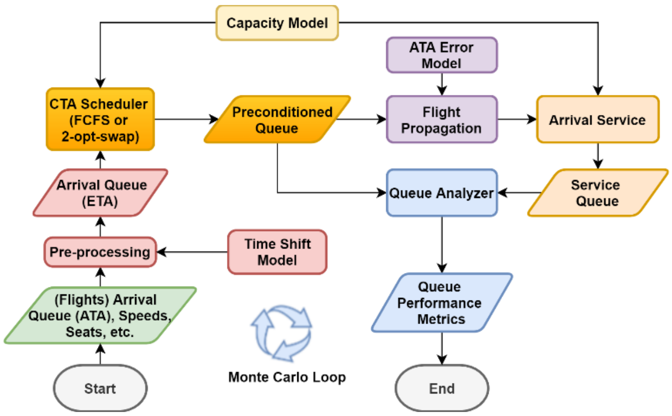

Figure 3. Simulation System Flow Chart

1) Flight CTA Scheduling

The flight scheduling process takes a flight table containing flight ETAs and speeds as input, then provides a column of CTAs in the flights table as output. We tested two flight scheduling schemes in our simulation:

1) FCFS

The first-come-first-served (FCFS) scheduling scheme is used in TBFM today and serves as the baseline model for this study. This scheduling scheme assigns CTA slots to flights on an order-of-arrival basis. In our simulation, the number of slots and slot length are calculated using the airport arrival rate (AAR). Each slot represents the CTA that was assigned. Given a flight in the list, we first try to assign the slot that corresponds to the flight's ETA. If this slot is occupied, we consider advancing the flight to a slot earlier than its ETA, starting with the closest slot available, and moving incrementally up to earlier ones if previous checked slots are not available, up to a limit of the maximum advance time ($\mathrm{ADV_{MAX}}$). Each flight's maximum advance time was calculated using Equation 1. In this way, our scheduling scheme will not violate the speed range limit for controllers.

$$\mathrm{ADV_{MAX}} = \frac{D_{FH}}{s_{curr}} - \frac{D_{FH}}{s_{max}}, \qquad (1)$$

where D_{FH} is the distance between the freeze horizon and the metering fix/runway. For short distance flights with planned distance L shorter than 250 nm, $(L - 50)$ nm was used instead of D_{FH}, where 50 nm is the nominal distance to the top of descent;
s_{curr} is the flight's current speed;
s_{max} is the flight's maximum speed allowed.

If a flight cannot either be assigned to the slot corresponding to its ETA or be advanced up to the limit of its maximum advance time, we delay the flight by assigning the earliest available CTA slot later than its ETA slot. Because we want to assign all the flights CTAs, there is no maximum delay in this step. Practically, moving a flight to a later slot corresponds to a speed decrease en route, and if this cannot be accommodated, then some physical vectoring must occur.

2) 2-opt-swap

We proposed a 2-opt-swap scheduling scheme, inspired by a local search algorithm first proposed in 1958 for solving the traveling salesman problem [23]. The idea is to improve the results by swapping the position of every possible pair in the solution set (CTA slot assignments) and update the solution set if the swapping makes improvements.

Our 2-opt-swap starts with a FCFS scheduling scheme and adds a swapping step after FCFS to reduce the total controlled arrival delay cost. Each feasible pair of flights is first taken out from the flights table as the "before swap table". Then, an "after swap table" is created by swapping their CTAs. A set of feasibility checks are then conducted to examine if the swapping is allowed: the *maximum advance time* check, the *maximum delay time* check, the *carrier preference* check, and the *controlled arrival delay cost* check.

The *maximum advance time* check examines if the swapping requires the controller to speed up the flight to where the maximum speed limit is violated. The calculation can be done following Eq. (1) in the FCFS section.

The *maximum delay time* check examines if the swapping generates any controlled arrival delay that exceeds the maximum value. We present two methods to compute this maximum delay value. The first method is fairly simple: we set a fixed value for all flights' maximum delay time, e.g., 10 min, and only allow the swapping if neither of the two flights' controlled arrival delay time exceeds this value. This may require the flights to do path stretching and vectoring to meet the assigned CTAs. The second method does not allow any path stretching and vectoring and only allows the flights to slow down and remain on their original path. Each flight's maximum advance time can be calculated using Eq. (2). By completing this maximum delay time check, our scheduling scheme will not violate the speed range limit for controllers and allow flights to remain on their original path. In this study, the first method (10 mins) was chosen for its relaxed limits. Note this value can be relaxed when arrival demand exceeds AAR/capacity, and controlled arrival delay is generally large, to achieve more feasible swaps.

$$\mathrm{DLY_{MAX}} = \frac{D_{FH}}{s_{min}} - \frac{D_{FH}}{s_{curr}} \qquad (2)$$

where D_{FH} is the distance between the freeze horizon and the metering fix/runway;
s_{curr} is the flight's current speed;
s_{min} is the flight's minimum speed allowed.

The *carrier preference* check only allows swapping if a) the two flights are from the same carrier, or b) the two flights are from different carriers, but the flight being swapped back has lower or equal priority level than the flight being swapped forward. Each carrier is required to provide priority labels for their flights. The priority labels range from 1 (lowest priority) to 5 (highest priority), where each priority label group contains 20% flights of the whole fleet. Some values we recommend airlines to take consideration into when setting their own priority level list can be given, such as number of connecting passengers within a time range, aircraft types and fuel consumption, number of seats and passengers, deviation from their scheduled times of arrival, but airlines can remain confidential on how they construct their own priority list as long as they provide the final labels in desired format.

The *controlled arrival delay cost* check guarantees that we only accept the swapping if it results in reduced total controlled arrival delay cost (calculated using Eq. (3)). This can be easily changed to other metrics we are interested in, such as the delay cost related to the scheduled times of arrival. We use a delay coefficient of 0.4$/min/seat to calculate the controlled arrival delay cost. This value is used in existing studies for airline passenger delay cost (in cost per passenger per minute) [24], [25]. In real implementations, a load factor can also be multiplied if it is provided by the carriers. More comprehensive cost schemes could also be applied.

$$\text{controlled\_arrival\_delay \_cost} = c \cdot n \cdot t \tag{3}$$

where c is the delay coefficient, n is the number of seats, and t is the controlled arrival delay, calculated by subtracting the ETA from the CTA, and only positive values are kept.

Due to swapping feasibilities and calculation time consideration, we do not conduct these checks between every possible pair of flights in the dataset. Instead, we conduct these checks for every flight and its following 4 flights in the list, ranked by the simulated ETAs. In the simulation, we also tested setting this number to 3, 5, 6, and 7. A larger number is computationally expensive but does not necessarily bring a significant amount of additional swapping or controlled delay cost saving, because eventually the swaps start to violate the above time checks.

The algorithm is given below:

```
2-opt-swap (flightsTable) {
  calculate MAX_DLY for each flight and save in flightsTable
  do until no improvement is made {
    for every flight i in flightsTable except the last one {
      for the following 4 flights j or until the end of flightsTable {
        if swapping the CTA of i and j passes the MAX_ADV
        and MAX_DLY check {
          if swapping the CTA of i and j passes the carrier
          preference check {
            if swapping the CTA of i and j reduces
            controlled arrival delay cost {
              swap the CTA of i and j
            }
          }
        }
      }
    }
  }
}
```

2) Flight Propagation

The aircraft receive these assigned CTAs and maneuver to meet them. In reality, it is not likely that aircraft will be able to meet their CTAs with zero error. A service delay occurs when a flight arrives when the server is busy (with previous flights) and has to wait in the queue for the next available service slot. The service delay is calculated by substracting the ATA from the time of service (only positive values are kept). The service delay cost can then be calculated using the same parameters as in Eq. (3), similar to the controlled arrival delay cost.

To model these en route delays, we added Gaussian noise to the assigned CTAs to represent their actual time of arrival (ATA).

$$\mathrm{ATA}_i = \mathrm{CTA}_\mathrm{i} + t_i \tag{4}$$

where $t_i \sim N(\mu, \sigma^2)$. In Run 1 (the single realization), one Gaussian noise model, $N(0, 15^2)$ was used. In Run 2 (the 5000 iteration realization), 5 zero-mean Gaussian noise models with standard deviations of 0, 10, 20, 30, and 40 seconds were used to simulate the different navigation capabilities of aircraft. All aircraft in each model were treated to follow the same distribution. In real life, the exact form of the noise should depend on the phase of flight and on different aircraft's navigation capabilities.

IV. Analysis and Results

A first simulation run with a single realization was conducted to illustrate the simulation process and metrics. The parameters we used are as follows: AAR = 84 h^{-1}; flight meeting noise model: Gaussian noise among all flights with mean = 0s and standard deviation = 15s. The flights priority labels were simulated by ranking the number of seats within each carrier and allowing 20% flights in each priority label group (higher number of seats corresponds with higher priority level). We conducted the simulation with the two different scheduling schemes separately. Figure 4 shows that the total controlled arrival delay is reduced with each swapping iteration in the 2-opt-swap scheme. A comparison of the results between the two schemes is given in TABLE III.

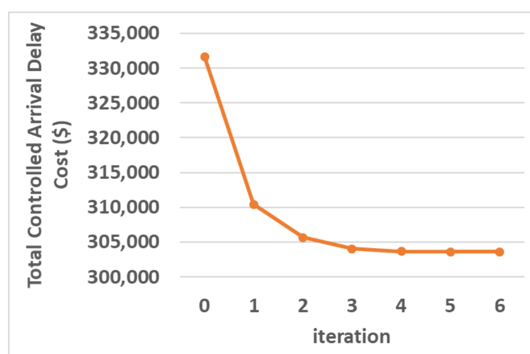

Figure 4 Total Controlled Arrival Delay Cost vs. iteration (2-opt-swap)

From TABLE III. , we can observe that there is a reduction of 8.5% for the total controlled arrival delay cost by conducting 2-opt-swap, compared to the FCFS only scheme. The total service delay between the two scheduling schemes only varies by 0.5%, which is as expected since the only source of this delay comes from the random meet time noise generated after scheduling. There is still a reduction of 7.8% for total delay cost after meet time noise models in the service. Note that these results are from one realization and are for illustrative purposes only.

TABLE III. RUN 1 RESULTS: COMPARISON BETWEEN THE TWO CTA SCHEDULING SCHEMES

	FCFS	2-opt-swap	Reduction	Reduction %
Number of service delayed flights:	1082	1082	0	0.0%
Total service delay cost ($):	34734	34916	-182	-0.5%
Number of controlled arrival delayed flights:	1000	951	49	4.9%
Total controlled arrival delay cost ($):	331636	303614	28022	8.5%
Total delay cost ($): [1]	366370	338529	27841	7.6%
Total delay cost (Method 2) ($): [2]	365125	336599	28527	7.8%
Number of Delayed Flights	1076	1070	6	0.6%

Notes:

1. The total delay cost here is simply the summation of total service delay cost and total controlled arrival delay cost;
2. The delay here is defined as the difference between the time a flight gets service and its ETA.

To further test and validate the 2-opt-swap scheduling scheme, a second simulation was conducted by repeating a Monte Carlo loop (shown in Figure 3) for 5000 realizations, for each of 13 different AAR values (20 to 140), 5 aircraft meet time performance models, and the 2 scheduling schemes. We tracked the total delay cost (using Method 2 in TABLE III.), the total controlled arrival delay cost, and the total service delay cost.

Similar to Run 1, it is not surprising that we do not observe a significant difference for the total service delay cost between the two scheduling schemes, given the same AAR and the same standard deviation of the meet time noise σ (shown in Figure 5). In fact, the absolute difference percentage of the total service delay between the two schemes stayed within 0.6%. This is because the service delay occurs due to the meet time error of the aircraft, which happens after flights being assigned their CTAs in the scheduling process and should not be affected by the differences between the two scheduling schemes. Within each scheduling scheme, we observe some trends as expected: as σ increases, the total service delay cost increases; as AAR increases, the total service delay cost decreases.

On the contrary, we can observe a significant reduction for the total controlled arrival delay cost when switching from FCFS to 2-opt-swap (summarized in TABLE IV. and Figure 6). In the interest of brevity, only 5 selected AAR values are shown in TABLE IV. Since there were no differences in the CTA scheduling part among the 5 noise models, we summarized the results (mean and standard deviation) for the total controlled arrival delay cost over these 5 noise models; that is, a total of 25,000 simulation runs for each AAR and scheduling scheme.

We can observe that compared to the FCFS only scheme, our 2-opt-swap scheme can reduce the total controlled arrival delay cost by $18,000 when AAR = 80 h^{-1} and 20% when AAR = 110 h^{-1}. The total controlled arrival delay cost drops significantly when the AAR increases from below 80 h^{-1} to above 80 h^{-1}, which is about the peak hour demand, for both scheduling schemes. This is one of the reasons why the reduction percentages remain low when the AAR is below 80h^{-1} (the denominators are large). Another reason is because the swaps between any two CTAs are more likely to violate the maximum time checks when the AAR is low, and thus less likely to be feasible. In reality, a ground delay program will kick in when AAR is at such low levels and the incoming traffic will be adjusted. On the other hand, when the AAR is very large, the total controlled arrival delay cost reductions are also relatively low because there is not much controlled arrival delay in the system and not many swaps are needed (passing the *controlled arrival delay cost* check). The 2-opt-swap scheduling scheme generates the most savings when the AAR is within [80, 90] in this simulation setting (shown in Figure 6). The total delay cost reduction shows a similar trend to the total controlled arrival delay cost reduction, because the scale of the total service delay cost is much smaller than the total controlled arrival delay cost.

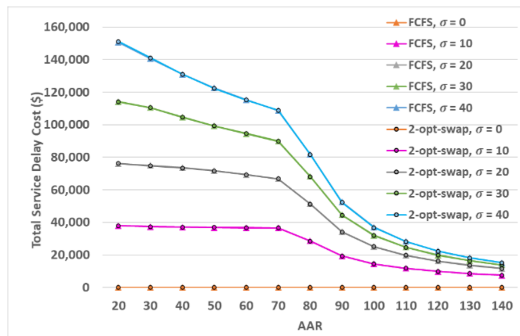

Figure 5 Total Service Delay Cost vs. AAR for different noise models and scheduling schemes

TABLE IV. RUN 2 RESULTS: TOTAL CONTROLLED ARRIVAL DELAY COST FOR SELECTED AAR VALUES

AAR	70	80	90	110	130
Baseline (FCFS) ($)	3896085	554814	113242	18164	5474
2-opt-swap ($)	3892945	536145	96745	14554	4505
Reduction Value ($)	3055	18733	16777	3772	1041
Reduction Percentage	0.1%	3.4%	14.8%	20.8%	19.0%

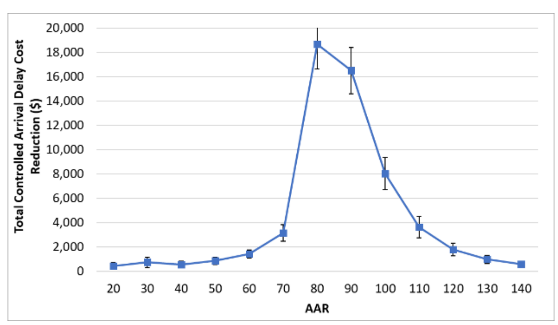

Figure 6. Total Controlled Arrival Delay Cost Reduction ($) vs. AAR

V. Findings and conclusions

This paper proposed a scheduling scheme – 2-opt-swap – for incorporating CDM concepts into TBFM. The scheme allows flight pairs (including those from different carriers) to swap their CTAs under certain circumstances in order to reduce the total controlled arrival delay cost. Simulations were designed for a FCFS baseline model and the 2-opt-swap model for a full day period in ATL. Real historical flight data were used in the simulation. Aircraft information including the speed and the number of seats were also used in the swapping process and the delay calculations. A simulation run with one realization was first done to illustrate how these schemes work, followed by a simulation run of 5000 Monte Carlo realizations and summarized averages from those realizations. Results show that compared to the FCFS only scheduling scheme, the 2-opt-swap scheduling scheme can reduce the total controlled arrival delay cost up to $18,000 and 20% for the 1293 flights on the day of simulation under certain parameter settings.

VI. Future work

A significant contribution of this project is the refinement of the simulation platform to be able to handle a wider variety of CTA scheduling schemes. In this vein, other scheduling schemes could be developed and tested. Other goals and objectives could be adopted, such as reducing the controlled arrival delay cost related to the scheduled time of arrival instead of ETA. More comprehensive cost schemes could be applied. The inter-carrier swapping idea could be adjusted to adapt different versions of carrier-expressed preferences. Optimization models such as integer programming could also be applied. We could compare our current heuristic scheduling algorithms with the optimal solution from the optimization models to check the performance of the algorithms.

As stated in the introduction, we have not addressed the challenges of obtaining airline cost/preference information, while ensuring that the airlines are encouraged to provide truthful information and also ensuring airlines are treated fairly. Related to this, we are only conducting scheduling for the metering fix freeze horizon. In the future, we could test CDM concepts whereby airlines could swap in and out of their "slots" at other locations such as the Coupled Metering Point (CMP) freeze horizon and the Extended Metering Point (XMP) freeze horizon. An advantage to this kind of scheme is that a more robust re-organization of the CTA scheduling can be accomplished over a greater distance, with the understanding that the CTA order need not stay the same between these different metering points. Thus, more real-time information on flight progress (due, for example, to wind conditions) could be accommodated, as could a multi-stage slot-swapping strategy.

References

[1] Torres, S. and Nagle, G. (2016). Analysis of prediction uncertainty in interval management. 16th AIAA Aviation Technology, Integration, and Operations Conference (ATIO), Washington, DC.

[2] Wambsganss, M. (1997). Collaborative decision making through dynamic information transfer. *Air Traffic Control Quarterly* 4(2), 107-123.

[3] Ball, M.O., Hoffman, R.L., Knorr, D., Wetherly, J., & Wambsganss, M. (2000). Assessing the benefits of collaborative decision making in air traffic management. In: Proceedings of the 3rd USA/Europe Air Traffic Management R&D Seminar, Napoli, Italy.

[4] Chang, K., Howard, K. Oiesen, R., Shisler, L., Tanino, M., & Wambsganss, M.C. (2001). Enhancements to the FAA ground-delay program under collaborative decision making. *Interfaces* 31(1), 57-76.

[5] Vlachou, K. & Lovell, D. J. (2013). Mechanisms for equitable resource allocation when airspace capacity is reduced. *Transportation Research Record* 2325, 97-102.

[6] Vakili, N. (2009). Preference Based Fair Allocation of Limited Airspace Resources. Ph.D. Dissertation, University of Maryland, College Park.

[7] Raj Mohanavelu Umamagesh, P. (2018). Mechanisms for Trajectory Options Allocation in Collaborative Air Traffic Flow Management. M.S. Thesis, University of Maryland, College Park.

[8] Carlson, P.M. (2000). Exploiting the opportunities of collaborative decision making: A model and efficient solution algorithm for airline use. *Transportation Science* 34(4), 381-393.

[9] Ball, M.O., Hoffman, R., Odoni, A., & Rifkin, R. (2003). A stochastic integer program with dual network structure and its application to the ground-holding problem. *Operations Research* 51, 167 – 171.

[10] Sherali, H.D., Staats, R.W., & Trani, A.A. (2003). An airspace planning and collaborative decision-making model: Part I – Probabilistic conflicts, workload, and equity considerations. *Transportation Science* 37(4), 434-456.

[11] Jones, J.C., Lovell, D.J., & Ball, M.O. (2018). Stochastic optimization models for transferring delay along flight trajectories in order to reduce fuel usage. *Transportation Science* 52(1), 134-149.

[12] Vossen, T. & Ball, M.O. (2006). Slot trading opportunities in collaborative ground delay programs. *Transportation Science* 40, 29 - 43.

[13] Idris, H., Chin, C., & Evans, A. (2019). Accrued delay application in trajectory-based operations. Proceedings of the 13th USA/Europe Air Traffic Management R&D Seminar, Vienna, Austria.

[14] EUROCONTROL. (2018). Delays – three questions and many answers. Retrieved from https://www.eurocontrol.int/news/delays-three-questions-and-many-answers.

[15] Schwartz, B., Benjamin, S., Green, S., & Jardin, M. (2000). Accuracy of ruc-1 and ruc-2 wind and aircraft trajectory forecasts by comparison with acars observations. *Weather and Forecasting* 15, 313-326.

[16] DVB Bank SE (2017). An Overview of Commercial Aircraft 2018-2019. Aviation Research Department, DVB Bank SE, Schiphol, The Netherlands.

[17] FAA (2019). Aircraft Registry. Retrieved from https://www.faa.gov/licenses_certificates/aircraft_certification/aircraft_registry/releasable_aircraft_download/.

[18] FlyRadius (2017). Aircraft Information. Retrieved from https://www.flyradius.com/aircraft-information.

[19] Controller (2019). Aircrafts for Sale. Retrieved from https://www.controller.com/.

[20] ICAO (2020). DOC 8643 – Aircraft Type Designators. Retrieved January 14, 2020 from https://www.icao.int/publications/DOC8643/Pages/Search.aspx.

[21] FAA (2011). Air Traffic Organization Policy, Facility Operation and Administration. Retrieved from http://tfmlearning.faa.gov/Publications/atpubs/FAC/1007.html.

[22] FAA (2019). ATL. Retrieved from https://www.fly.faa.gov/ Information/east/ztl/atl/atl_aar.htm.

[23] Croes, G.A. (1958). A method for solving traveling-salesman problems. *Operations Research* 6(6), 791-812.

[24] Ferguson, J., Qadar Kara, A., Hoffman, K., & Sherry, L. (2013). Estimating domestic U.S. airline cost of delay based on European model. *Transportation Research Part C: Emerging Technologies* 33, 311-323.

[25] Pita, J.P., Barnhart, C., & Antunes, A.P. (2013). Integrated flight scheduling and fleet assignment under airport congestion. *Transportation Science* 47(4), 477-492.