

Optimal Dynamic Scheduling of Aircraft Arrivals at Congested Airports

Aditya P. Saraf* and Gary L. Slater†
University of Cincinnati, Cincinnati, Ohio 45221

DOI: 10.2514/1.29464

An optimal airport arrival scheduling algorithm, which works within a hierarchical scheduling structure, is presented. This structure consists of schedulers at multiple points along the arrival route. Schedulers are linked through acceptance rate constraints, which are passed up from downstream metering points. The innovation in this algorithm is that these constraints are computed by using an Eulerian model-based optimization scheme. In this scheme an aggregate airspace model is derived online. Optimization based on predictions of this model is used to compute the optimal acceptance rates at all metering points, so that the airport throughput is maximized while keeping sector counts within limits. This rate computation removes inefficiencies introduced in the schedule through ad hoc acceptance rate computations. The scheduling process at every metering point uses its optimal acceptance rate as a constraint and computes optimal arrival sequences by using a combinatorial search algorithm. We test this algorithm in a dynamic air-traffic environment, which can be customized to emulate different arrival scenarios. Results from Monte Carlo simulations show improved performance over first-come, first-served ordering for different values of key parameters such as arrival aircraft mix, average arrival rate, maximum position-shift constraint, and level of uncertainty in estimated time of arrival prediction.

I. Introduction

BECAUSE of increases in air-traffic density, limited-capacity resources in the National Airspace System (NAS) are currently under strain. This is particularly true for airports and Terminal Radar Approach Control (TRACON) areas. Especially for those airports that lie close to the boundaries of Air-Route Traffic Control Centers (ARTCCs or “centers” for short) and are surrounded by short path-length sectors, containing complex traffic flows, a need for coordinated air-traffic flow management across multiple centers has been identified [1]. The short sectors surrounding these airports have small controllable times. Hence, a large number of sectors are required to absorb a given delay and the aircraft have to be delayed when they are far away from the airport, maybe even across multiple centers. When aircraft are at such large distances from the airport, expected time of arrival (ETA) predictions can be erroneous and, if used in scheduling algorithms, may lead to congestion, unnecessary delays, and missed slots. Currently, a time-based decision support tool called Traffic Management Advisor—Single Center (TMASC) is in use at eight ARTCCs in the United States [2] and expansion to all the 20 centers is planned [3] in 2007. TMASC helps the controllers within a single center in scheduling the aircraft arriving at airports within their center. TMASC creates scheduled times of arrival (STAs) for the aircraft by retaining the first-come, first-served (FCFS) order of arrival according to their estimated ETAs, but delaying some aircraft to maintain the mandatory separations. TMASC ignores uncertainty in the ETA estimates while computing the STAs. The application of such a time-based metering scheme at constrained TRACON spaces discussed earlier can lead to inefficient schedules with unnecessarily large delays and congestion [4].

To eliminate reliance on long-distance trajectory estimates, in this paper we propose a hierarchical scheduling structure, wherein aircraft are scheduled successively to the next downstream

scheduling point on their route. The scheduler at each point schedules only those aircraft which are present in a short-length window situated upstream of it. The schedulers at multiple points are linked with each other by passing acceptance rate constraints upwards from downstream schedulers. Thus, the big scheduling problem is broken down into smaller, loosely coupled scheduling problems. ETA of an aircraft at only the immediate downstream metering point is deconflicted with other aircraft ETAs. ETAs at further downstream points are only used to check acceptance rate conformance. This allows for swapping of aircraft in the arrival sequences at downstream metering points in successive scheduling runs, if their ETAs were erroneous during the previous runs. A tightly coupled scheduler would, on the other hand, deconflict aircraft at all metering points and its schedules would be affected the most by the ETAs that have the tightest restrictions, namely ETAs at the runway, which being the most distant estimates are the most erroneous.

This hierarchical scheduling structure is similar to the one proposed by NASA for the decision support tool called Multi-Center Traffic Management Advisor (McTMA) [1], which is currently under development. McTMA implements a coordinated multicenter, time-based metering program by using a hierarchical multipoint scheduling structure to schedule arrival aircraft landings at constrained airports [3–5]. The scheduling algorithm that we present here was developed as a general approach to time-based metering and is not restricted to the scheduling environments for which McTMA was designed. The algorithm is flexible and reconfigurable; it can be modified easily to work within any environment, terminal or en route, wherever a loosely coupled hierarchy of schedulers has to be used. The scheduling algorithm consists of two interconnected parts. The first part is the computation of optimal acceptance rate constraints at all metering points by an Eulerian model-based optimization scheme. Schedulers at each metering point in the hierarchical structure are subject to these optimal acceptance rate constraints. The second part is an optimal scheduling and sequencing algorithm (a combinatorial A search-based algorithm), which works separately for each metering point and computes the optimal arrival sequence at that metering point.

In the past, there have been a number of efforts to develop algorithms for optimal aircraft arrival scheduling. Brinton [6] used a branch and bound tree-search algorithm to find the optimum arrival sequence. This algorithm was found to be unacceptable by controllers [7], because it was slow and produced unstable aircraft sequences as a result of ETA fluctuations. Beasley et al. [8] used a

Received 24 December 2006; revision received 15 May 2007; accepted for publication 13 June 2007. Copyright © 2007 by Aditya Saraf and Gary Slater. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0731-5090/08 \$10.00 in correspondence with the CCC.

*Graduate Research Assistant, Department of Aerospace Engineering and Engineering Mechanics. Student Member AIAA.

†Professor Emeritus, Department of Aerospace Engineering and Engineering Mechanics. Associate Fellow AIAA.

mixed integer zero-one formulation for the scheduling problem and solved the static case by using a linear programming-based tree-search similar to the one in Brinton [6]. In another work [9], Beasley et al. used a genetic search algorithm to find the optimum arrival order for a group of aircraft. In a later paper [10], Beasley et al. extended their previous work to the dynamic case, where as time progresses new aircraft are added to the scheduling process and some old aircraft land and leave the scheduling process. Also, there have been some efforts to incorporate airline prioritization into the scheduling process, most notably by Carr et al. [11,12]. In our prior work, Saraf and Slater [13], we presented an optimal static scheduling algorithm, which takes advantage of differing aircraft-type-dependent minimum separation constraints to save delay by commanding sequence changes from the FCFS order. In this paper, we further develop this algorithm and adapt it to a hierarchical scheduling structure, by adding the computation and application of acceptance rate constraints to the previous algorithm. Additionally, we carry out extensive testing of this algorithm in a dynamic air-traffic environment.

The acceptance rates to be used as constraints in the scheduling process are computed by using a spatially discretized Eulerian airspace model in an optimization scheme. The idea of using such models of the airspace and using aggregate quantities like average speed to characterize flows has been used in the past, most notably by Menon et al. [14–17]. The basic idea behind these so-called Eulerian models is to divide the airspace into interconnected control volumes and obtain an airspace model in terms of the aggregate flow properties of aircraft within each control volume. Using steady state assumptions, Menon et al. [14] model the air-traffic flow as a linear time invariant dynamic discrete-time system. In another research initiative, Bayen et al. [18–20] derive an Eulerian network model for the air-traffic flow in the NAS, based on the continuous Lighthill–Whitham–Richards partial differential equation [21,22]. In Saraf and Slater [23], the authors combine the Eulerian and Lagrangian approaches to eliminate the drawbacks inherent in either approach when used on its own, and develop a two level control system for optimal air-traffic flow management. It was this work that motivated us to explore the use of Eulerian models for optimal acceptance rate computation.

The main innovation in our work is the computation of acceptance rates from an Eulerian model-based optimization scheme. These acceptance rates are used as inputs to the scheduling algorithm. In this optimization scheme, the arrival airspace is divided into control volumes and an aggregate airspace model is derived online, in terms of conservation equations for each control volume. Optimization based on predictions of this model is used to compute the optimal acceptance rates at all metering points, so that the airport throughput is maximized while keeping the number of aircraft in each sector below a maximum allowable limit. In McTMA, acceptance rate constraints are applied in the form of rate profiles; a rate profile coming from a downstream metering point represents gaps in the STAs of aircraft that have been already scheduled to the downstream metering point. Aircraft are scheduled to an upstream metering point, only if they would fit onto the gaps in the rate profile when they reach a downstream metering point. This acceptance rate computation depends directly on downstream STAs, which seems counter-intuitive. Rather, the acceptance rate at each metering point must represent the maximum air-traffic flow that can pass through that point without causing sector overload, irrespective of the STAs. The novel rate computation method in our work removes inefficiencies introduced in the schedule through ad hoc acceptance rate computations, because the acceptance rates computed by this algorithm are exactly the highest flow rates that do not violate maximum allowable sector-count constraints. So if scheduling at each metering point is constrained to agree with the maximum acceptance rate computed for that point, then it is guaranteed that acceptance rates at further downstream points would not be violated by the resulting schedule.

Another feature of this paper is the formulation of a succinct, linear, discrete-space model of the air-traffic flow, which is tailored for optimal schedule computation with separate constraints on each

individual sector delay, with the facility of allocating different priorities (weights) to delays assigned to different aircraft or in different sectors, and with the freedom of choosing the cost function from a wide range of options, so as to achieve a number of different air-traffic management (ATM) objectives. The scheduling algorithm uses this linear model for the computation of delay costs and the assignment of delays. It applies the optimal acceptance rates at metering points as constraints on the scheduling. The algorithm computes the optimal aircraft sequence and the optimal delays to be taken by each aircraft in each sector of its route by doing a tree search, based on the A-search [24] algorithm, over a tree of arrival orders. A quadratic function of the delays is used as a measure of the “goodness” of an arrival order.

We test this algorithm in a dynamic air-traffic environment, with aircraft coming in and leaving the scheduling horizon as time progresses. This dynamic environment can be customized to emulate a variety of arrival scenarios, with varying average arrival rates and aircraft-type mixes. Results showing improved performance over first-come, first-served scheduling are presented and analyzed for different values of key parameters like aircraft mix, average arrival rate, maximum position-shift constraint, hold-delay assignment procedure, and maximum position-shift constraint. Average performance characteristics are obtained by carrying out Monte Carlo simulations. Robustness of this scheduling algorithm to expected time of arrival prediction errors is also demonstrated by carrying out Monte Carlo simulations with random ETA prediction uncertainties.

This paper is organized as follows: Section II gives a problem statement, which gives a clear understanding of the scheduling problem that we intend to solve in this paper. It also presents a detailed description of the elements of the scheduling process and where the different parts fit into the big picture. Because optimal acceptance rate computation is the main innovation in this work, we describe this process first, in Sec. III. Section IV will then present an overview of the linear discrete-space air-traffic model that we use in the optimal scheduling algorithm and give details of the scheduling algorithm. Section V presents results of using the optimal scheduling algorithm to solve some static scheduling problems. Section VI will give an overview of the dynamic scheduling environment, which we use for dynamic testing of our scheduling algorithms. Section VII presents comparative results obtained by running simulations with both the optimal scheduling algorithm and the FCFS scheduling algorithm in the dynamic scheduling environment. The paper will end by summarizing the conclusions in Sec. VIII.

II. Problem Statement

In this section, we first define the scheduling problem to be solved. A detailed description of all the elements of this problem follows in Sec. II.A. The scheduling process in a hierarchical scheduling structure, being different than the traditional scheduling process, is discussed separately in Sec. II.B. Because the main innovation in our work is the acceptance rate computation, we allocate the explanation of this process to a separate subsection, Sec. II.C.

Problem definition:

Given the following data:

- 1) airspace geometry, consisting of the lengths and directions of air routes coming into the airport area with scheduling points defined on them;
- 2) aircraft specific information, like the aircraft’s wake-vortex category (aircraft type), its arrival route, and its ETA estimates to different metering points, for each aircraft;
- 3) freeze-region and scheduling-window lengths;
- 4) predicted sector transit times for each aircraft and for each sector along its route;
- 5) minimum separation constraints, in the form of temporal separations to be maintained between successive aircraft for all combinations of leading and trailing aircraft types (weight classes);
- 6) aircraft acceptance rates from downstream metering points;
- 7) maximum sector delay absorption capacity constraints for all sectors in the airspace under consideration; and

8) maximum position-shift constraint in the form of maximum shifts allowed from the FCFS order.

Compute:

- 1) the optimal arrival sequence at the scheduling point for all the aircraft that are currently present in the scheduling window;
- 2) the optimal delay to be taken by each of these aircraft in each sector of the freeze region, when it reaches that sector in the future (refer to Sec. II.B for details on the scheduling process); and
- 3) the optimal STAs to the scheduling point for each of these aircraft.

Such that the following constraints are satisfied:

- 1) aircraft-type dependent minimum separation constraints at the scheduling point;
- 2) sector delay capacity constraints for all the sectors in the freeze region;
- 3) acceptance rate constraints coming from the downstream reference points; and
- 4) constraint on maximum position shifts (MPS) allowed from the FCFS sequence.

And a measure (to be defined) of the total system delay is minimized.

A. Detailed Description of the Elements of the Optimal Scheduling Problem

1. Airspace Geometry

The scheduling algorithm uses information about the geometry of the air routes merging onto the airport, with scheduling points and reference points defined on them. A scheduling point is a point at

which a scheduler is placed. The section of the air route between two scheduling points consists of multiple sectors. The end points of a part of the air route lying inside a single sector are called reference points (RP). A simple example arrival scenario is shown in Fig. 1. Aircraft are coming along routes A, B, C, and D, named according to the reference point that they pass through. Routes A and B merge at the scheduling point RP0 and routes C and D merge at RP0'. These two merged streams then proceed to meet at the scheduling point RP5. The converged aircraft then proceed from RP5 to the final scheduling point RP10, which may be the runway threshold or a point on the boundary of the TRACON. Our algorithm is flexible in that it can handle any airspace, terminal or en route. For simplicity, it is assumed that each of the segments RP0–RP5, RP0'–RP5, and RP5–RP10 consist of five equal length sectors. With the equal length assumption, there are four reference points each, between RP0 and RP5, and between RP0' and RP5 (these reference points are not shown in the figure). However, for dealing with a real-world problem, our algorithm could handle sectors differing in their lengths.

2. Freeze-Region and Scheduling-Window Lengths

It is assumed that the freeze region for each scheduling point reaches out up to 200 n mile upstream of the scheduling point. The scheduling window starts at 200 n mile from the scheduling point and continues until 400 n mile from it. For the arrival geometry in Fig. 1, the scheduling window for RP5 consists of the 200-n-mile-long sections upstream of points RP0 and RP0' and the freeze region consists of the regions RP0–RP5 and RP0'–RP5.

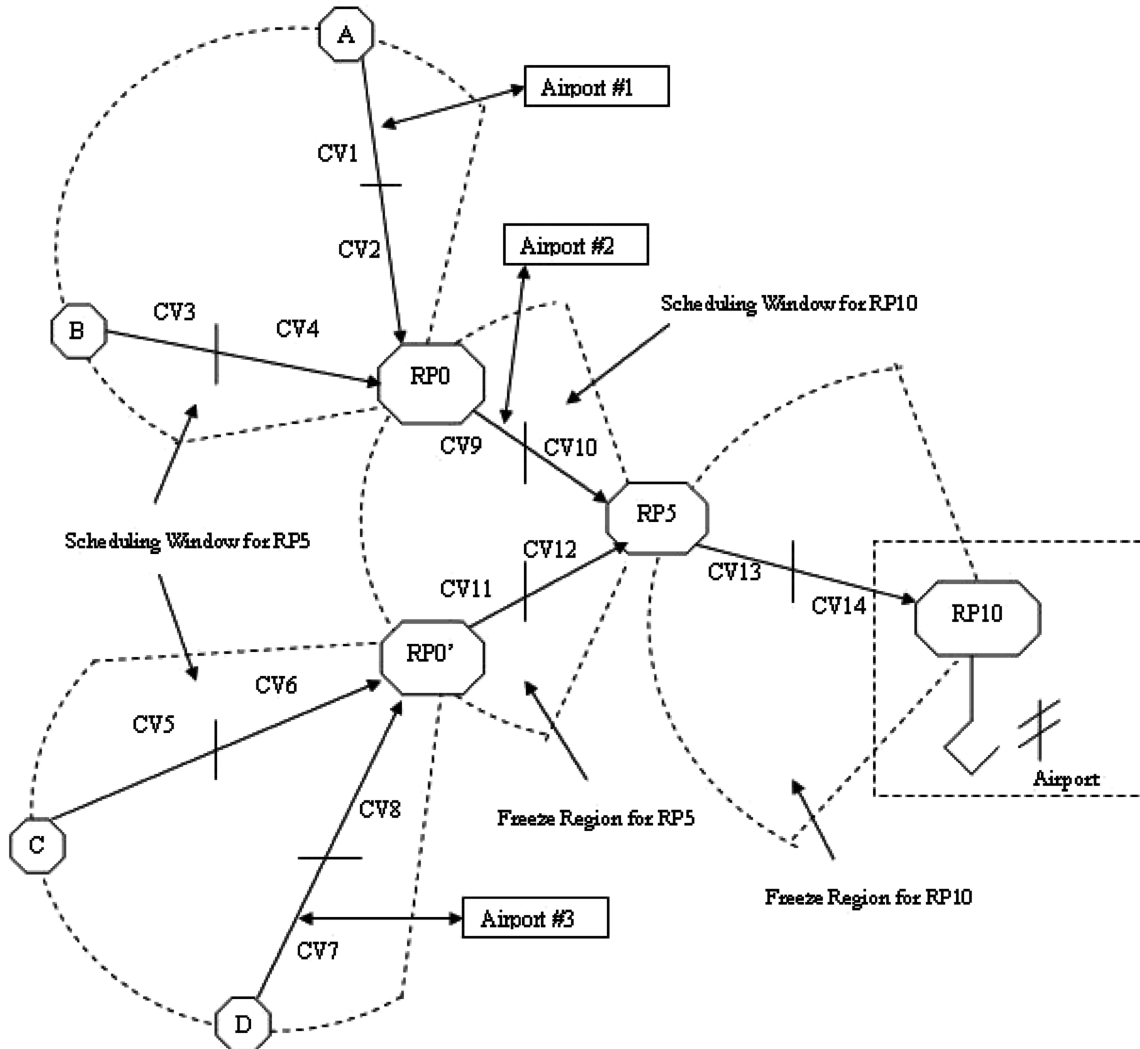


Fig. 1 Arrival airspace geometry.

Table 1 Minimum separation constraints for different aircraft types at 200 n mile from the airport

Trailing aircraft type →	Heavy	Large	Small
Average speed when 200 n mile away from DFW airport, nmi/h	396.57	337.48	124.56
Minimum separation assuming a 5 MIT restriction, s	45.39	53.34	144.51

Table 2 Minimum separation constraints in seconds for different pairs of aircraft types at the runway

Trailing aircraft type → Leading aircraft type	Heavy	Large	Small
↓			
Heavy	60	90	120
Large	60	90	90
Small	60	60	60

3. Sector Transit Times

For the purpose of the simulation, sector transit times are assigned to aircraft depending on their aircraft types; for a heavy aircraft, the nominal sector transit times are computed by assuming a nominal speed of 0.8 Mach (which is approximately the cruise speed for most heavy aircraft). However, to depict small aircraft-to-aircraft speed differences within the same weight class, the nominal speed for each aircraft of the heavy type is computed as a random number uniformly distributed between 0.76 and 0.84 Mach (i.e., $\pm 5\%$ from 0.8 Mach). Nominal speeds of 0.7 Mach ($\pm 5\%$) and 0.6 Mach ($\pm 5\%$) are used for the large and small aircraft types, respectively.

4. Minimum Separation Constraints

Minimum separation requirements at all scheduling points, other than the runway threshold, are computed by assuming a 5 miles-in-trail distance separation and converting it into a time separation by using average speeds for the three types of aircraft. The aircraft-type dependent average speeds were computed as follows: recorded Aircraft Situational Display to the Industry (ASDI) format data for Dallas/Fort Worth (DFW) airport arrivals on 17 March 2005 were first processed through NASA's Future ATM Concepts Evaluation Tool (FACET) to get formatted track data. FACET [25] is a simulation and analysis tool which provides a simulation environment for development and analysis of new air-traffic management concepts. The formatted data were then analyzed to obtain the average speeds of aircraft when they were 200 n mile away from the DFW airport. Average speeds were computed separately for the three weight classes: heavy, large, and small. Table 1 shows these average ground speeds and the time separation obtained assuming a 5 miles-in-trail distance separation constraint. The values of minimum separation requirements in Table 1 were used as a constraint on the scheduling at point RP5 (which we have assumed to be 200 n mile away from the airport in our simulations).

At the runway, the Federal Aviation Administration has mandated aircraft-type dependent minimum separations to minimize the effect of wake vortices. We assume three aircraft types: heavy, large, and small. The aircraft-type dependent minimum distance separations are converted into minimum time separations that we impose at the runway (RP10) in our scheduling algorithm. These separations are shown in Table 2.

5. Maximum Position-Shift Constraint

For a set of N aircraft there are $N!$ possible orders of arrival, but we look at only those orders which allow a maximum of m position changes from the FCFS order for any aircraft. While making the choice of the maximum limit on position changes, we have to achieve a trade-off between decreasing the delay cost associated with the arrival sequence and keeping the controller workload manageable. A good choice for this constraint is two or three position shifts, because a larger limit on position changes could increase controller workload

without achieving much delay reduction. We studied the effect of using different values for the MPS constraint, and the results bear out that, for a typical traffic mix, little or no advantage can be gained by using an MPS constraint bigger than three. Results are presented in Sec. VII.D.

6. Maximum Sector Delay Absorption Capacity Constraints

The maximum delay that can be absorbed in a particular sector is a complex function of factors, such as the size and shape of the sector, the complexity of the traffic flows present in the sector, and winds. Sector controllers take into consideration these factors and decide the amount of delay that they can absorb in their sector at any particular time. A narrow sector with little space for vectoring aircraft will have a lower delay absorption capacity than a wide sector of the same length, because aircraft can be delayed longer in a wider sector by vectoring. In our study, we have used a maximum delay absorption limit of 1 min for a sector of length 40 n mile. However, the algorithm can handle sectors with differing maximum allowable delay constraints, and we tested its sensitivity to this limit by using different sector delay absorption limit values ranging between one and three minutes. The results of these simulations and analysis thereof will be presented in Sec. VII.C.

7. Hold-Delay Assignment

If the delay assigned to a particular aircraft is greater than the delay capacity of all the sectors in the freeze region, then the aircraft is commanded to fly around a "holding fix" situated at the starting point of the freeze region (points RP0 and RP0' represent the holding fixes for the scheduling process at RP5) for an integer number of cycles, so that the excess delay can be absorbed. The remaining delay (which is within the capacity of the freeze region) is assigned to the freeze-region sectors. The time required by an aircraft to complete one hold cycle (i.e., go around the holding fix once) is assumed to be 2 min for all aircraft; the total hold delay will be in multiples of 2 min. So if an aircraft is held, then it may end up taking more delay than what is required. However, in some airspaces, the procedure for holding an aircraft specifies an expected further clearance time, which is an expected time to exit the hold or get further holding-delay command. If the delay time to be absorbed by holding is already known, then the size of the "race track" hold pattern can be varied to absorb only that much delay. The hold delay in such a case can be set exactly equal to the required value. We tested our algorithm with both types of hold-delay assignment, continuous-valued delay and delay in integer multiples of a fixed cycle time. Results are presented in Sec. VII.B.

B. Scheduling Process

At any point in time, the scheduling process for a given scheduling point will only involve those aircraft which are present in its scheduling window at that time. For example, looking at Fig. 1, we note that the scheduling process at RP5 will only consider aircraft in the 200-n-mile-long sections upstream of RP0 and RP0'. For these aircraft, the scheduling algorithm computes the optimal arrival sequence at RP5 and the optimal delays to be taken by the aircraft in the regions RP0–RP5 and RP0'–RP5. These optimal delays, when added to the current predicted ETAs, give the optimal STAs (delays are assumed to be nonnegative). For aircraft in the scheduling window, the freeze-region sectors are downstream of their current positions. So in the scheduling algorithm we are computing the optimal delays that these aircraft will have to take when they reach these downstream sectors in the future. Delay is assigned to an aircraft (by adding the delay to its freeze-region transit time) only when it reaches the freeze horizon. Until then its STA and its position

in the arrival sequence keep getting updated at each scheduling cycle. The scheduling will be subject to acceptance rate constraints coming from downstream metering points.

C. Role of the Eulerian Model-Based Optimization Scheme

As explained in the Introduction, we subject the scheduling process to acceptance rate constraints coming from downstream metering points to avoid congestion in the future, when the aircraft reach downstream metering points. The role of the Eulerian model-based optimization (EMBO) scheme is to compute the optimal acceptance rates at all metering points. This scheme first divides the arrival airspace into control volumes and derives an Eulerian airspace model in terms of conservation equations for each control volume. Then, optimization based on predictions of this model is used to compute the optimal flow rates at all metering points so that the rate of arrival is maximized (but within the allowed airport acceptance rate) while keeping all the sector aircraft counts within their allowed limits. A detailed explanation of the Eulerian model derivation and the model predictive control scheme appears in Sec. III.

Another problem with scheduling arrivals to busy airports is the presence of tower en route control (TEC) aircraft. These are aircraft which take off from nearby airports and land at the main airport, without ever entering the center airspace. Track data for such aircraft is currently not available to the center host computer; it is available through the enhanced traffic management system (ETMS). An algorithm for scheduling arrivals to busy airports should be able to input ETMS data and assign proper places for such aircraft in the arrival sequence. The EMBO scheme takes care of this problem by prescribing the optimal number of aircraft that can take off from each nearby airport and enter the scheduling process for the main airport during a given time interval, subject to constraints on the minimum and maximum number of takeoffs during that time interval.

Thus, the innovative EMBO scheme computes the optimal acceptance rates at all metering points and the optimal departure times for TEC aircraft taking off from nearby airports. The optimal acceptance rate computed by the EMBO scheme for a scheduling point is used as an upper limit on the number of aircraft that can be scheduled to arrive at the scheduling point per time window (more on this in Sec. IV.C). In this way, the EMBO scheme avoids the generation of schedules that would cause congestion and missed slots at downstream metering points in the future and involves TEC aircraft into the scheduling process.

III. Eulerian Model-Based Optimization Scheme

This section describes the Eulerian model derivation and the use of this airspace model in a model-based predictive control (MPC) scheme to compute optimal acceptance rates at all scheduling points.

A. Eulerian Model of the Airspace Around a Busy Airport

The basic idea behind the Eulerian modeling technique is to divide the airspace into smaller interconnected control volumes (CVs) and then describe the traffic flow dynamics by equations involving aggregate flow properties of all the aircraft within each CV [14]. For steady state assumptions (i.e., assuming that the average flow properties like average velocity of the aircraft within each CV remain constant), these equations lead to a linear time invariant model of the airspace. In our work, we do not assume that the average flow properties in each CV remain constant. Instead, we derive a new model for the airspace at each time step using updated radar data. This leads to a linear time-varying (LTV) model for the airspace. In this section, we describe the derivation of an Eulerian model specialized to the route topology shown in Fig. 1. This figure shows the airspace around an airport divided into control volumes numbered 1–14. Each control volume represents a 100-n-mile-length part of the arrival route. There are three airports near the main airport and TEC aircraft taking off from these nearby airports head to the main airport's TRACON without entering the center airspace. For each of the nearby airports we prescribe minimum and maximum takeoff rates, and within these limits the MPC algorithm computes

the optimum number of TEC aircraft that should take off from each airport.

Let $n_i(k)$ be the total number of aircraft in the CV i at the time instant $k\tau$, where τ is a discrete time step. We make the following modeling assumptions.

In the time interval $k\tau$ to $(k+1)\tau$, aircraft can enter the CV i in one of the following ways:

- 1) From one of the adjacent CVs i_{adjacent} with rate $y_{i_{\text{adjacent}}}(k)$ aircraft per minute.
- 2) From an airport within the same CV at the rate $y_i^{\text{depart}}(k)$ aircraft per minute.
- 3) If the CV is on the boundary of the modeled region and gets input from outside the region, then aircraft enter the CV at the rate $y_j^{\text{in}}(k)$ aircraft per minute. Here j stands for the external input stream number. (In this example $j = 1, 2, 3$, and 4 , respectively, for the streams coming in through points A, B, C, and D).

In the time interval $k\tau$ to $(k+1)\tau$, aircraft can leave a particular CV i in one of the following ways:

- 1) Going into the corresponding adjacent CV at the rate $y_i(k)$ aircraft per minute.
- 2) If $i = 14$, then the aircraft leave this CV to arrive at the main airport at the rate $y_{14}(k)$.

Now, consider a single CV j of length L_j . Let v_j be the average speed of all aircraft present within the CV. Then, the average time taken by an aircraft to transit through CV j is $\tau_j = L_j/v_j$. If there are $n_j(k)$ aircraft in CV j at time $k\tau$, then on an average we expect that the number of aircraft that leave CV j after a time interval of τ is $(\tau/\tau_j)n_j(k)$. So the expected ratio of aircraft that stay back in the CV after one time step of τ seconds is $\alpha_j = (1 - (\tau v_j)/L_j)$. This ratio is an important parameter of the Eulerian model. We model the air traffic control (ATC) control actions as an additional number of aircraft per τ seconds held back in the CV beyond the number $\alpha_j n_j(k)$ that is expected to stay back. This additional number of aircraft $u_j^{\text{ATC}}(k)$ is held back per minute so as to maintain a smooth flow of the air traffic and avoid congestion in the entire airspace. Thus, the dynamics of the traffic flow can be written as

$$n_i(k+1) = \alpha_i n_i(k) + \tau y_{i_{\text{adjacent}}}(k) + \tau y_j^{\text{in}}(k) + \tau y_i^{\text{depart}}(k) + \tau u_i^{\text{ATC}}(k), \quad i = 1, 2, \dots, 14; \quad j = 1, 2, 3, 4 \quad (1)$$

$$y_i(k) = \frac{1}{\tau} (1 - \alpha_i) n_i(k) - u_i^{\text{ATC}}(k) \quad (2)$$

The inflow $y_{i_{\text{adjacent}}}(k)$ from an adjacent CV i_{adjacent} is in turn dependent on the number of aircraft in CV i_{adjacent} at time $k\tau$, $n_{i_{\text{adjacent}}}(k)$. This gives a coupling between the state equations for $n_i(k)$ and $n_{i_{\text{adjacent}}}(k)$.

If we aggregate all the state variables into a single vector \mathbf{X} , $\mathbf{X}(k) := [n_1(k) n_2(k) \dots n_{14}(k)]^T$, all the control inputs into a vector \mathbf{U} , $\mathbf{U}(k) := [u_1^{\text{ATC}}(k), \dots, u_{14}^{\text{ATC}}(k); y_1^{\text{depart}}(k), \dots, y_{14}^{\text{depart}}(k)]^T$, and all the uncontrolled inputs into a vector \mathbf{Y}_d , $\mathbf{Y}_d(k) := [y_1^{\text{in}}(k), \dots, y_4^{\text{in}}(k)]^T$, then we get a linear model for the air-traffic flow:

$$\mathbf{X}(k+1) = \mathbf{A}_k \mathbf{X}(k) + \mathbf{B}_k \mathbf{U}(k) + \mathbf{E}_k \mathbf{Y}_d(k) \quad (3)$$

$$\mathbf{Y}(k) = \mathbf{C}_k \mathbf{X}(k) + \mathbf{D}_k \mathbf{U}(k) \quad (4)$$

where the entries of matrices \mathbf{A}_k , \mathbf{B}_k , \mathbf{C}_k , \mathbf{D}_k , and \mathbf{E}_k are obtained by looking at equations for individual CVs.

1. Obtaining the Coefficients of the Linear Time-Varying Model

The ratio α_i , in the LTV airspace model (1), represents the expected fraction of the number of aircraft currently present in CV i that will still be present in the same CV after one Eulerian time-step interval (τ). To predict this ratio, we model all the aircraft in a CV as one entity characterized by one speed, namely the average speed of the aircraft currently present in the CV. Using this average speed we compute the average CV transit time for an aircraft as

$$\text{avg time to transit CV}_i = \frac{\text{CV length}_i}{\text{avg speed}_i} \quad (5)$$

where CV length_i is the length of the air route belonging to CV i . Then the fraction α_i is computed as

$$\alpha_i = 1 - \frac{\tau}{\text{avg time to transit CV}_i} \quad (6)$$

2. Prediction of Future Input Airflows $y_1^{\text{in}}-y_4^{\text{in}}$ and the Ratios α_i

Predictions of future input airflows $y_1^{\text{in}}-y_4^{\text{in}}$, which enter the airspace via points A–D, can be obtained from the ETAs of aircraft upstream of the respective points A–D. To predict the flow rate y_1^{in} for example, we put all the aircraft upstream of point A into one-minute bins depending on their ETAs at point A; bin 1 will contain the aircraft expected to arrive at point A within the next minute, bin 2 will contain all the aircraft which are expected to arrive at A within the next 60–120 s, and so on. Then the inflow rate y_1^{in} is predicted as the average number of aircraft per one-minute bin. This average is computed over a time window of half-hour length.

B. MPC Using the Eulerian Model

Now, we apply MPC [26] using this linear model of the air-traffic flow to solve the following problem: Given an initial configuration $X(0)$ and predicted values of the uncontrolled input vector $Y_d(k)$ over a certain future horizon $k\tau$ to $(k+N_p)\tau$, find the optimum ATC control inputs (rates of ATC hold backs in all CVs), so that the outflow rate from the last CV, which is equal to the airport landing rate, is maximized subject to constraints on the maximum number of aircraft in each CV during each time step and the airport acceptance rate constraint. Note that a prediction horizon of 1 h is used in the simulations that we carried out. Updated acceptance rates are computed by solving the MPC every 5 min (i.e., $\tau = 5$ min).

Objective Function: We iterate the linear airspace model of Eq. (3) to predict the states $\mathbf{X}(k)$, $k = 1, 2, \dots, N_p$, and the outputs $\mathbf{Y}(k)$, $k = 1, 2, \dots, N_p$, where $N_p\tau$ is the prediction horizon. The outputs $\mathbf{Y}(k)$ represent the flow rates (number of aircraft per minute) coming out of each of the CVs, and so $Y_{14}(k)$ represents the airport arrival rate during the k th time step. We form a vector of N_p future airport arrival-rate predictions $\hat{\mathbf{Y}}$ from $Y_{14}(k)$, $k = 1, 2, \dots, N_p$, and define the objective function (which we seek to maximize) as

$$J[u_i^{\text{ATC}}(k), u_i^{\text{depart}}(k): i = 1, 2, \dots, 14; k = 1, 2, \dots, N_p] := \hat{\mathbf{Y}}^T \hat{\mathbf{Y}} \quad (7)$$

Constraints: Maximization of the above objective function is subject to the following constraints:

1) The first constraint comes from the natural fact that the number, $\tau u_i^{\text{ATC}}(k)$, which is the number of aircraft that are held back in a particular CV during a time step $k\tau$ to $(k+1)\tau$, should be a nonnegative number and it should not exceed the number of aircraft already present in the CV during that time step, $n_i(k)$. This gives the first constraint:

$$\text{Constraint 1: } 0 \leq (\tau u_i^{\text{ATC}}) \leq n_i(k): \quad i = 1, 2, \dots, N; k = 0, 1, \dots, N_p.$$

2) The aircraft flow rate from the last CV should not exceed the prescribed airport acceptance rate. This gives the second constraint:

$$\text{Constraint 2: } 0 \leq y_{14}(k) \leq AAR(k): k = 1, 2, \dots, N_p.$$

3) The number of aircraft in each CV should not exceed the maximum allowed limit. This limit is a parameter of the scheduling process. For the simulations carried out here, we have set it equal to CV length/miles-in-trail requirement (i.e., a 6 miles-in-trail constraint is assumed for all the simulations done). This gives the third constraint:

$$\text{Constraint 3: } n_i(k) \leq n_{(\text{max},i)}: i = 1, 2, \dots, N; k = 0, 1, \dots, N_p, \text{ where}$$

$$n_{(\text{max},i)} = \frac{\text{CV length}}{\text{MIT restriction}}$$

4) Number of TEC aircraft taking off from a nearby airport in each time step is also constrained to comply with the minimum and the maximum allowed takeoff rates at the nearby airports. These limits are parameters of the scheduling process. This gives the fourth constraint:

$$\text{Constraint 4: } \min \text{ take-off rate}_i \leq y_i^{\text{depart}}(k) \leq \max \text{ take-off rate}_i: i = 1, 2, \dots, N, k = 0, 1, \dots, N_p.$$

The MPC solves the following optimization problem:

$$\max_{\{u_i^{\text{ATC}}(k), u_i^{\text{depart}}(k): i=1,2,\dots,14; k=1,2,\dots,N_p\}} \hat{\mathbf{Y}}^T \hat{\mathbf{Y}} \quad (8)$$

subject to the state equation (3), and the constraints 1, 2, 3, and 4.

Solution to the quadratic program (8) gives the optimal hold-back rates and the optimal flows through all the CVs. The optimal flow rate for the CV containing the scheduling point represents the maximum aircraft flow rate that can pass through that point. This flow rate will be used as a constraint on the scheduling process. Section IV.C will explain how this acceptance rate is enforced on the scheduling process.

IV. Optimal Scheduling Algorithm

From a scheduling perspective, we model the air-traffic flow as a dynamic discrete-space system. Let $t_k^{(i)}$ be the time of arrival of the i th aircraft at the k th reference point along the trajectory of the aircraft. At any time, we index the first reference point that the aircraft is expected to reach by zero and the scheduling point is, say, the M th reference point. Then the air-traffic flow is modeled by the following equation:

$$t_{k+1}^{(i)} = t_k^{(i)} + \Delta_k^{(i)} + \tilde{u}_k^{(i)} + n_{\text{hold}}^{(i)} \times T_{\text{hold}}, \quad k = 0, 1, \dots, M-1; i = 1, 2, \dots, N \quad (9)$$

where

$\Delta_k^{(i)}$: transit time of the i th aircraft from the k th to the $k+1$ th reference point;

$\tilde{u}_k^{(i)}$: nonhold delay (always nonnegative) assigned to the i th aircraft, to be taken between the k th and $k+1$ th reference point;

T_{hold} : single-hold cycle time (assumed to be the same for all aircraft types);

$n_{\text{hold}}^{(i)}$: number of hold cycles assigned to the i th aircraft.

Transit times $\Delta_k^{(i)}$ are assumed to be obtained from aircraft trajectory predictions and are known at each scheduling run. When an aircraft reaches the freeze horizon, its freeze-region sector transit times are updated by adding the delays assigned to it. $n_{\text{hold}}^{(i)} \times T_{\text{hold}}$ is the hold delay assigned to the i th aircraft, and $\tilde{u}_k^{(i)} + n_{\text{hold}}^{(i)} \times T_{\text{hold}}$ is the total delay (hold delay plus nonhold sector delay) assigned to it. So the optimization variables are the nonhold delays, $\tilde{u}_k^{(i)}$, for $i = 1, 2, \dots, N; k = 0, 1, 2, \dots, M-1$; and the number of cycles for which the aircraft has to be held at the starting point, $n_{\text{hold}}^{(i)}$, for $i = 1, 2, \dots, N$. Note that the optimization variables $n_{\text{hold}}^{(i)}$, $i = 1, 2, \dots, N$, are integer valued.

By putting together terms for all the aircraft into a single vector, we get the following vector form of the discrete-time air-traffic flow system:

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \mathbf{\Delta}_k + \mathbf{\tilde{u}}_k + T_{\text{hold}} \mathbf{n}_h, \quad k = 0, 1, \dots, M-1 \quad (10)$$

where $\mathbf{x}_k := [t_k^{(1)}, t_k^{(2)}, \dots, t_k^{(N)}]^T$, $\mathbf{\Delta}_k := [\Delta_k^{(1)}, \Delta_k^{(2)}, \dots, \Delta_k^{(N)}]^T$, $\mathbf{\tilde{u}}_k := [\tilde{u}_k^{(1)}, \tilde{u}_k^{(2)}, \dots, \tilde{u}_k^{(N)}]^T$, and $\mathbf{n}_h := [n_{\text{hold}}^{(1)}, n_{\text{hold}}^{(2)}, \dots, n_{\text{hold}}^{(N)}]^T$.

The aim of the optimal scheduling algorithm is to maintain a conflict-free flow of aircraft, while minimizing some function of the total delay undergone by the aircraft as a group and at the same time keeping the commanded position changes in the arrival sequence to a

minimum. To achieve this aim, we define the following cost function, which we intend to minimize:

$$J(\mathbf{x}_0, \{\tilde{\mathbf{u}}_k\}_{k=0}^{M-1}, \mathbf{n}_h) := \sum_{k=0}^{M-1} \tilde{\mathbf{u}}_k^T R_k \tilde{\mathbf{u}}_k + \mathbf{x}_M^T Q_f \mathbf{x}_M + \mathbf{n}_h^T W \mathbf{n}_h \quad (11)$$

The quadratic function $\sum_{k=0}^{M-1} \tilde{\mathbf{u}}_k^T R_k \tilde{\mathbf{u}}_k$ penalizes the deviations away from the nominal transit times, while the quadratic function $\mathbf{x}_M^T Q_f \mathbf{x}_M$ puts a penalty on the total delay over the N aircraft. The remaining part of the cost, $\mathbf{n}_h^T W \mathbf{n}_h$, penalizes hold delays. The weights in the matrix W are assigned much larger values than the weights in the matrices R_k , because a nonhold sector delay is more acceptable than a hold delay. With the flexible weights in the matrices R_k , delays in different sectors can be weighted differently, for example, high altitude delays may be assigned lower penalties as compared to lower altitude delays. Entries of the weight matrix Q_f may be set up to assign higher priority to a particular aircraft or a group of aircraft as compared to others. This flexibility in the scheduling algorithm can be used for airline prioritization, that is to say, for a group of aircraft belonging to the same airline, the airline can assign different priorities to different aircraft by allocating high or low weights to their corresponding entries in the Q_f matrix.

Scheduling is subject to minimum separation constraints, constraints on maximum allowable delay in each sector, and the MPS constraint. We can also put constraints on the maximum number of hold cycles allowed for an aircraft by constraining the integer variables $n_{\text{hold}}^{(i)}$. So the resulting optimization problem is

$$\min_{\{\tilde{\mathbf{u}}_0, \tilde{\mathbf{u}}_1, \tilde{\mathbf{u}}_2, \dots, \tilde{\mathbf{u}}_{M-1}, \mathbf{n}_h\}} J(\mathbf{x}_0, \{\tilde{\mathbf{u}}_k\}_{k=0}^{M-1}, \mathbf{n}_h) \quad (12)$$

subject to

$$\begin{aligned} \mathbf{x}_{k+1} &= \mathbf{x}_k + \Delta_k + \tilde{\mathbf{u}}_k + \mathbf{T}_{\text{hold}} \mathbf{n}_h, \quad k = 0, 1, 2, \dots, M-1 \\ \tilde{\mathbf{u}}_k \min_k &\leq \tilde{\mathbf{u}}_k \leq \tilde{\mathbf{u}}_k \max_k, \quad k = 0, 1, 2, \dots, M-1 \\ |x_M^{(i)} - x_M^{(i+1)}| &\geq s_{\min}^{(i,i+1)}, \quad i \neq j, \quad i, j = 1, 2, \dots, N \\ \mathbf{n}_h^{(\min)} &\leq \mathbf{n}_h \leq \mathbf{n}_h^{(\max)} \end{aligned} \quad (13)$$

where $\tilde{\mathbf{u}} \min_k$ and $\tilde{\mathbf{u}} \max_k$ are the minimum and maximum allowable delays that can be absorbed in the k th sector of the air route and $s_{\min}^{(i,i+1)}$ is the minimum allowable temporal separation between consecutive aircraft i and $i+1$. In addition, $s_{\min}^{(i,i+1)}$ depends on the type of the aircraft i and $i+1$. Table 2 gives the runway separation constraints while Table 1 gives the separation constraints used at scheduling points 200 n mile away from the airport. Note that we have not added the maximum position-shift constraint to this optimization problem. This is because we take care of that constraint in the combinatorial tree-search algorithm (see Sec. IV.A) by limiting the search to only those arrival orders which satisfy the maximum position-shift constraint. Additionally, precedence constraints (constraints like aircraft “A” should arrive before aircraft “B”) can also be applied by limiting this tree search to only those arrival orders which satisfy precedence constraints. Acceptance rate constraints are also applied separately as explained in Sec. IV.C.

We note that the minimum separation constraints in this optimization problem (OP) (12) depend on the order of aircraft arrival at the M th reference point, which is the scheduling point. To find a possibly optimal sequence we have designed an intelligent search algorithm for searching a tree of arrival orders. The MPS constraint is taken care of during this tree search by eliminating those arrival orders that violate it. A detailed discussion of the tree-search algorithm and a fast algorithm for solving the OP (12) for each arrival order, is presented in Sec. IV.A. At each step of the tree search, after computation of the optimum STAs, the STAs are run through an acceptance rate violation check algorithm. This algorithm applies the optimal acceptance rate, computed by the EMBO scheme, as a constraint on the number of aircraft that can be scheduled to the scheduling point in a given interval of time and computes the additional delays required to satisfy this constraint.

A. Combinatorial Search Algorithm for Finding the Optimum Arrival Order

We use a combinatorial search algorithm based on the general principles of the A-search algorithm [24] to search a tree of arrival orders and find a possibly optimum arrival order by solving the OP (12) for only a small fraction of the total possible orders. Because of its special structure, the OP (12) is solved by a fast algorithm, to be explained in Sec. IV.B. The arrival-order search is done only over those arrival orders that satisfy the maximum position-shift constraint. The cost function for the tree-search algorithm can be designed in such a way that the algorithm could be guaranteed to end by finding the optimal solution if it is given enough time [24]. However, because we are interested in finding the best possible answer as quickly as possible, we design the cost differently so that after a limited number of node expansions the order with which we end up is as close to the optimum as possible, if not the optimum. In separate experiments with a small (≤ 10) number of aircraft, the optimal order was found using an exhaustive search over all possible arrival orders. In these cases it was found that the optimum order always matched with the order given by the A-search algorithm.

The basic idea behind this algorithm is to form a tree, wherein each node is associated with an aircraft. The aircraft are numbered 1– N in the FCFS order. Nodes at the j th level of the tree are associated with those aircraft that can occupy the j th position in the aircraft sequence. At each level of the tree, the decision that we need to make is which node to expand first. This decision is made on the basis of a cost assigned to each node. As per the general guidelines of the A search, this cost must be a sum of a cost associated with the path from origin node to the current node, plus an estimate of the cost associated with the path from the current node to a goal node. In our algorithm, the cost assigned to each node is the cost of the OP (12) associated with the arrival order formed by the parents of the current node, followed by the current node, followed by the remaining nodes in FCFS order. So by the time we reach an N th level node, we expect that we would end up finding the optimal arrival order.

B. Fast Algorithm for Solving the Optimization Problem (12)

The optimization problem (12) has a special structure because it is associated with a fixed arrival order. Once the arrival order is fixed, a delay assignment, which minimizes any weighted sum of squares of the delays (of course with all positive weights), is the one which assigns to each aircraft its earliest possible STA. The earliest possible STA for any aircraft in our scheduling structure is either its ETA (if that is not in violation of the minimum separation constraints) or the minimum STA such that the separation constraints are not violated. The fast algorithm does delay assignment in just this same way. The fast algorithm needs just $N-1$ additions, $N-1$ subtractions, $2(N-1)$ comparisons and a maximum $N-1$ divisions to solve the OP (12), where N is the number of aircraft involved in the scheduling process. Once the optimal STA assignment associated with a given arrival order is computed by the fast algorithm, these STAs are processed through an acceptance conformance check, which is explained in the next section.

C. Interaction Between the EMBO Scheme and the Optimal Scheduling Algorithm

The solution to the Eulerian model-based optimization problem (8) of Sec. III.B gives the optimum desired acceptance rates at all metering points along the air route. We apply these acceptance rate constraints in the scheduling algorithm as follows: at each step of the arrival-order tree search, after solving the OP (12) to get the optimal STAs for a given arrival order, we further subject these STAs to an acceptance rate conformance check using the computed optimal flow rate for the scheduling point. This acceptance rate conformance check is done by running the STAs through a bin constraint application (BCA) algorithm. This algorithm is similar to the acceptance rate algorithm in Wong [27]. The cost associated with that arrival order is then computed by including the additional delays that the BCA algorithm assigns to conform with acceptance rate constraints. The cost is given in Eq. (11).

1. Bin Constraint Application Algorithm

The purpose of the BCA algorithm is to schedule the aircraft to arrive in a given order and as close to their computed STAs as possible, without exceeding the optimum aircraft flow rate at the scheduling point mandated by the EMBO scheme. For a given arrival order, solving the OP (12) gives deconflicted STAs for all aircraft in RP5's scheduling window. After that, the BCA algorithm will check if the computed STAs violate the optimal acceptance rate at RP5. In [27], the BCA algorithm also uses data structures called acceptance rate bins. Each acceptance rate bin represents an interval of 30 s starting at the bin's start time and contains a number which represents the number of aircraft scheduled to arrive within that 30 s time interval. Arrival rate is expressed in terms of the number of aircraft arriving in a fraction of an hour. This fraction is known as arrival-rate interval (ARI), and its length is an adjustable parameter of the BCA algorithm. The length of the ARI is generally a whole number multiple of the bin length. In our simulations we use a 10 min ARI. To check violation of the arrival rate, we compare the number of aircraft arriving in a window of the size of the ARI with the allowed arrival rate. This rate conformance is applied as follows.

Before starting the BCA algorithm, all aircraft which have entered the freeze region within the past "ARI-length" minutes are assigned to their respective bins. All aircraft which had entered the freeze region before this time are ignored because they would not affect the bin constraints. The BCA algorithm starts with the first aircraft in the optimum arrival sequence. The 30 s bin corresponding to the STA of this aircraft is identified and a window of the length of the ARI is extended from the STA's bin into the past. The number of aircraft in all the bins contained in this window is counted. If this total number is less than the maximum number of aircraft allowed to arrive in one ARI-length interval, then this means that scheduling the aircraft at its current STA will not violate the arrival rate in this window. So now the window is advanced in time by one bin, and the total number of aircraft in the bins contained in the new window is again computed. If this total number is still less than the allowed arrival rate, then the window is again advanced one bin into the future. This process is continued until either the window has moved far enough into the future that it no longer contains the current aircraft's STA's bin or the total number of aircraft contained in some window exceeds the allowed arrival rate. If the total exceeds the arrival rate, then the STA of the aircraft is assigned to the first bin lying outside the window, because assigning it to any bin in the current window will still exceed the arrival rate. After the STA is delayed, the entire process of sliding

the window is restarted from the newly assigned bin. On the other hand, if the window moves past the STA's bin without exceeding any arrival rates, then the aircraft is assigned its current STA and the number of aircraft in the corresponding bin is increased by one. Thereafter, we continue with the same process for the next aircraft in the given arrival order. In this way, we find STAs that satisfy the downstream arrival-rate constraint at the scheduling point for all the aircraft in the scheduling window. The additional delays, assigned to aircraft due to acceptance rate violations, are also taken into consideration while computing the delay cost associated with a given arrival order. In this way, the optimum sequence computation will take into consideration the acceptance rate constraints. These constraints are not just postapplied to the computed optimum sequence, because that may destroy optimality.

V. Static Scheduling Results

In this section, we present the results of some static scheduling runs that demonstrate the capabilities of our scheduling algorithm. In Sec. V.A, we test the scheduling algorithm by running the algorithm on a set of aircraft arrivals for 17 March 2005, for the DFW airport. The purpose of this test was to verify the practical application of our algorithm. Section V.B discusses how flexibility in the choice of cost function can be used to achieve a number of ATC objectives and presents an example of optimizing a "passenger cost."

A. Scheduling of DFW Arrivals Using Simulated Track Data

Track data required for this test were obtained by running NASA's FACET software in its "simulation" mode, with the recorded ASDI data file for 17 March 2005 as its input. ETAs of the aircraft to the DFW airport and to circles of radii 200 and 400 n mile around the DFW airport were obtained from this track data. An optimal schedule was generated for the aircraft present in the scheduling window of the DFW airport (i.e., aircraft which lie in between the 400 and 200 n mile circles surrounding DFW airport). Figure 2 gives a pictorial view of the order changes commanded by our algorithm from the FCFS order. In this figure, a solid line joins the FCFS STA and the optimal STA of each aircraft. The identifiers of the aircraft, along with aircraft type (H: heavy, L: large, S: small), are displayed just above or below the line corresponding to that aircraft. An MPS constraint of two was used here. Minimum separation constraints in Table 2 were imposed. Out of the whole set of aircraft present in the scheduling window, only those aircraft that show position changes




Fig. 2 ETAs and optimal STAs of DFW arrivals.

from FCFS, plus one aircraft each at the beginning and the end of this partial sequence, are shown in the figure. All other aircraft were assigned the same delays and positions by both the FCFS and optimal schedulers (these aircraft are not shown in the figure).

The optimal scheduler saved 14.17% of total delay over the FCFS scheduler by commanding some order changes. Particularly, the heavy aircraft AAL1994 was assigned to a slot two places ahead in the sequence to avoid “a small trailing a heavy” sequence which would require a 2 min separation. The small aircraft EGF874 was placed in between two large aircraft, AAL2434 and AAY6181, to minimize the overall delay. One reason behind presenting this example here is to demonstrate the scheduling algorithm’s applicability to real-world problems. Another reason is that this example brings out a key factor that determines the performance of the optimal scheduling algorithm: the optimal scheduler will produce bigger savings over the FCFS scheduler if the FCFS sequence contains a larger proportion of a small trailing a heavy sequences.

B. Results with Passenger Cost

Because the optimization in our scheduling algorithm does not depend on the specification of the cost, the cost can be customized to represent a wide variety of ATC objectives (i.e., the cost could be a linear or quadratic or any other complex function of the optimization variables). For example, the delay cost could be customized to put heavier penalties on delays assigned to heavy aircraft, like commercial airline jets that carry hundreds of passengers, and smaller penalties on small aircraft, like private jets that carry only a few passengers. Alternatively, if a particular sector along the arrival route is congested, then we could impose a heavier cost on the delay taken by all aircraft in that sector to avoid the assignment of big delays in that sector. A heavier cost on a single aircraft can be used in case that aircraft needs to be expedited.

As an example, we carried out a scheduling run with a delay cost that puts a 20-times heavier weight on delays assigned to heavy aircraft and a 5-times heavier weight on delays assigned to large aircraft as compared to small aircraft delays. In other words, this delay cost is a passenger cost, because costs on delays assigned to aircraft are commensurate with the number of passengers that they carry. Static scheduling runs carried out with this passenger cost show that the optimal sequence produced large (about 35%) delay cost savings over the FCFS sequence; as expected, optimal schedules generated with the passenger cost show a tendency to slow down small and large aircraft, while assigning very small delays to heavy aircraft. On the other hand, the optimal sequence, with equal weights on all aircraft types, was more than 50% costlier than the FCFS sequence in terms of the passenger cost.

VI. Dynamic Air-Traffic Environment and Simulation

A. Dynamic Air-Traffic Environment

In Sec. V, we presented results of some static scheduling runs using the optimal algorithm. Now we wish to test the algorithm in a dynamic environment, with aircraft entering and leaving the scheduling window and acceptance rates getting updated as time progresses. For this purpose, we have developed a dynamic air-traffic simulation environment at the University of Cincinnati using MATLAB, within which we can test different scheduling algorithms.

In this section, we explain this algorithm testing environment. The environment models the scheduling scenario shown in Fig. 1. The incoming aircraft are generated as Poisson arrivals with a specified average arrival rate, that is to say, we generate ETAs of the aircraft at the reference point RP10 by assuming that the interarrival time between two consecutive flights is an exponential random variable with its mean being equal to the reciprocal of the specified arrival rate. As explained in Sec. II.A.3, the aircraft are assigned speeds/transit times depending upon their weight classes plus a small random perturbation to depict aircraft-to-aircraft differences. The ETAs of the aircraft at upstream scheduling points RP5 and RP0 are computed by subtracting the corresponding transit times from their

ETAs at RP10. The simulation is dynamic in the sense that as time progresses, aircraft come in and move out of the scheduling window; also, acceptance rates at the scheduling points get modified because the EMBO scheme uses a time-varying Eulerian model, which is derived in real time. Delays are added to aircraft transit times only when the aircraft reach the freeze horizon. We looked at three average arrival rates: one aircraft arriving at RP10 per 60, 90, and 120 s. We performed the simulation for a period of 30 min. A new schedule was computed every minute, and so there were 30 scheduling runs during the simulation period. The optimal acceptance rates were updated every 5 min by running the EMBO scheme, each time using an updated Eulerian model derived in real time.

For each average arrival rate, we looked at three different mixes of aircraft: the first mix consisted of 15% heavy aircraft, 55% large aircraft, and 30% small aircraft; the second mix was a 25–50–25 distribution; and the third mix was a 40–40–20 distribution. For each combination of arrival-rate aircraft mix, we generated 100 arrival data sets by randomly assigning arrival times and aircraft types to the aircraft. Note that in each of these data sets not only the arrival times but also the arrival order is randomly generated; only the long-term average arrival rate and the long-term percentage of H–L–S aircraft remain constant for a given data set. For each arrival-rate aircraft mix combination, we applied the FCFS scheduling algorithm and optimal scheduling algorithm to each of the 100 data sets corresponding to that combination by running the dynamic simulation for a period of 30 min. We computed the average delay per aircraft and total number of holds assigned by the FCFS and the optimal schedulers as well as the total number of order changes over FCFS commanded by the optimal scheduler for each such data set. The average was taken over the total number of aircraft involved in the scheduling process during the simulation period. Then, we took a mean of these values over the 100 data sets corresponding to each arrival rate and mix. Table 3 lists the mean average values of some important parameters of the FCFS and optimal schedulers. Section VII.A presents the related analysis of these results. Note that in the simulations the scheduling process of both the FCFS and the optimal schedulers was subject to the downstream acceptance rate constraints given by the EMBO scheme. Separate tests were carried out to study the effect of the hold-delay assignment process, the sector delay limit, and the MPS constraint on the scheduling results. Results of these tests are listed in Tables 4 and 5, and the related analysis is presented in Secs. VII.B–VII.D.

B. Sequence Jumps Due to Trajectory Prediction Errors

As mentioned in the preceding sections, we work within a hierarchical scheduling structure to eliminate reliance on unreliable long-distance ETA estimates. However, as a result of ETA prediction uncertainty, the sequence computed at a given scheduling run may vary greatly from that computed at the previous run. This was the pitfall of some past attempts at computing optimal sequences, most notably Brinton [6]. We wanted to test our scheduling algorithm for such undesirable “sequence jumps” by incorporating ETA prediction uncertainty in our simulations. For this purpose, a separate test was run on the data sets associated with the average arrival rate of one aircraft per minute and the 15–55–30 H–L–S aircraft mix. In this separate test, to simulate this uncertainty, a random perturbation was added to the nominal transit time for each aircraft each time a new scheduling run was carried out. This random perturbation was assumed to be uniformly distributed between $-\Omega/N_{\text{steps}}$ and $+\Omega/N_{\text{steps}}$ percent of the freeze-region transit time. N_{steps} represents the number of scheduling runs carried out during the simulation ($N_{\text{steps}} = 30$). In past studies, the maximum ETA prediction uncertainty, when the aircraft is 200 n mile from an airport, was found to be of the order of 1 min [4], which is 3.33% of the average freeze-region transit time. To take care of different levels of uncertainty within the 1 min upper limit, we look at three values of Ω : 1, 2, and 3.3333 in our tests.

For this separate test, we used 100 randomly generated arrival data sets corresponding to an average arrival rate of one aircraft per minute and a 15–55–30 H–L–S aircraft mix to study the effect of

Table 3 Mean delay savings over FCFS scheduling and mean FCFS delays (0% ETA prediction uncertainty) (savings expressed as percentages of the mean FCFS delay)

Aircraft-mix arrival rate (AR) combination	Overall savings (average FCFS delay per aircraft in seconds)	Savings for heavy aircraft	Savings for large aircraft	Savings for small aircraft
AR = 1 aircraft per 60 s Mix 1 (15% H, 55% L, 30% S)	26.83% (252)	0.92% —	6.08% —	35.31% —
AR = 1 aircraft per 60 s Mix 2 (25% H, 50% L, 25% S)	28.21% (200)	−0.05% —	7.59% —	38.82% —
AR = 1 aircraft per 60 s Mix 3 (40% H, 40% L, 20% S)	27.51% (146)	3.18% —	14.77% —	41.22% —
AR = 1 aircraft per 90 s Mix 1 (15% H, 55% L, 30% S)	29.64% (178)	17.30% —	13.56% —	37.60% —
AR = 1 aircraft per 90 s Mix 2 (25% H, 50% L, 25% S)	33.77% (133)	23.49% —	17.48% —	46.07% —
AR = 1 aircraft per 90 s Mix 3 (40% H, 40% L, 20% S)	32.43% (109)	17.33% —	28.50% —	44.53% —
AR = 1 aircraft per 120 s Mix 1 (15% H, 55% L, 30% S)	30.59% (102)	30.12% —	21.36% —	38.61% —
AR = 1 aircraft per 120 s Mix 2 (25% H, 50% L, 25% S)	36.12% (73)	42.05% —	39.33% —	44.49% —
AR = 1 aircraft per 120 s Mix 3 (40% H, 40% L, 20% S)	36.88% (65)	42.23% —	35.81% —	44.95% —

Table 4 Effect of using holding for an integer number of cycles vs assigning just a single hold time (arrival scenario: average arrival rate =1 aircraft per min, aircraft mix=mix 1)

Hold cycle time	FCFS		Optimal		
	Mean average delay per aircraft, s	Mean no. of holds	Mean average delay per aircraft, s	Mean no. of holds	Mean average delay savings (as a percentage of FCFS delay)
2 min	252	52.52	186	27.73	26.83%
1.5 min	223	54	172	31.70	23.54%
Hold delay assigned as just one delay time	211	—	184	—	14.86%

Table 5 Effect of using different limits on the maximum allowable sector delay (arrival scenario: average arrival rate =1 aircraft per min, aircraft mix=mix 1)

Sector delay absorption limit, min	FCFS		Optimal		
	Mean delay per aircraft, s	Mean no. of holds	Mean delay per aircraft, s	Mean no. of holds	Mean average delay savings (as a percentage of FCFS delay)
1	252	52.52	186	27.73	26.83%
2	231	27.77	160	13.50	31.26%
3	221	16.50	153	7.84	31.64%

ETA prediction uncertainty. The same 100 data sets were used for testing scheduling at all uncertainty levels. Because randomness is introduced in the predictions of sector transit times, it is necessary to carry out the simulation over and over for a large number of times to get a general picture of the characteristics and advantages of the optimal scheduling algorithm as compared with the FCFS scheduling algorithm. To this end, 30-min simulations were carried out 10 times for each of these 100 data sets and for each of the perturbation percentages (1, 2, and 3.3333%). The data were averaged over the all simulations to get a general picture. Section VII.A gives some of the results of these simulations.

VII. Dynamic Scheduling Results

In this section, we present some results obtained from the tests that we ran in the dynamic simulation environment described in Sec. VI.A. Section VII.A presents comparison of mean average delays assigned by the FCFS and optimal schedulers and related analysis. Sections VII.B–VII.D present comparison and analysis of the results obtained by using different values for key parameters of

the scheduling process: the hold-pattern transit time, maximum sector delay absorption capacity, and MPS constraint, respectively. Section VII.E presents the results of the separate test carried out to check variability of computed sequences from one scheduling run to the next under ETA prediction uncertainty.

A. Comparison of Optimal and FCFS Scheduling

The analysis in this subsection is based on results shown in Table 3. Table 3 shows the overall percentage mean average delay savings achieved by the optimal scheduling algorithm over FCFS scheduling and their breakup for each aircraft type. The numbers in parentheses are the mean average FCFS delays per aircraft in seconds.

1. Distribution of Delays and Delay Savings for Different Arrival-Rate and Aircraft-Mix Combinations

From Table 3, it is observed that, on average, FCFS delays go up with increasing average arrival rate, which is expected (optimal delays also show similar behavior). The overall delay savings over

FCFS are substantial for all three average arrival rates and it is seen that the savings increase with decreasing arrival rate. Analysis showed that the reason behind this is that although the average FCFS delay increased with increasing average arrival rate, average optimal delay increased at a bigger rate. This was because the percentage of a small behind a small/large/heavy sequences in the arrival batch, which is an indicator of the potential for delay savings, remained constant in spite of the increasing arrival rate and it was not possible to achieve delay savings beyond a certain point.

For each arrival rate, the delays assigned by both the FCFS and optimal schedulers are maximum for mix 1 and minimum for mix 3. An analysis of the aircraft sequences present within the scheduling window shows that the mean percentage of aircraft pairs having a small behind a small/large/heavy sequence is an indicator of the size of average delays assigned as well as an indicator of the potential for delay saving. This indicator is maximum for mix 1 (28.41%), smaller for mix 2 (25.33%), and minimum for mix 3 (21.41%). Because a small behind a small/large/heavy sequence requires the maximum temporal separation (refer to Table 1), it is natural that the mix having a larger percentage of such sequences be assigned the maximum delays. The delay savings for all arrival rates and mixes are a sizeable proportion of the respective FCFS delays. From the numbers, it appears that for each average arrival rate, the delay savings go up with the percentage of small aircraft until a point of maximum delay savings is reached, thereafter delay savings go down with the increasing percentage of small aircraft.

2. Distribution of Delay Savings Among Different Aircraft Types

From Table 3, it is seen that the optimal scheduler favors small aircraft and in some cases penalizes heavy aircraft unjustly, but this is a manifestation of the cost function that we define. This cost function penalizes delays assigned to all aircraft equally and hence our algorithm minimizes the overall delay by slowing down heavy and large aircraft to avoid a small behind a small/large/heavy sequences. But this effect could be remedied by using heavier weights for large and heavy aircraft as compared to the weights on the small aircraft in the cost function, or by using a passenger cost function as in Sec. V.B. It also seen that the percentage delay savings achieved by running the simulation for longer periods remain within the 25–27% range.

B. Effect of Using Different Values for the Hold-Pattern Transit Time

As mentioned in Sec. II.A.7, the scheduling algorithm can be customized to handle continuous-valued hold-delay assignment as well as hold-delay assignment in terms of integer multiples of a fixed hold-pattern cycle time. We studied the effect of using different values (2 and 1.5 min) for the hold-pattern transit time as well as of using a continuous hold time on the performance of the FCFS and optimal scheduling algorithms. Separate tests were carried out for this purpose. Table 4 shows the results obtained from these tests. The purpose of these tests is to verify if the optimal scheduler produces substantial benefits in terms of delay savings in all hold-delay assignment scenarios.

It was observed that a larger number of holds are commanded by both algorithms if the hold-pattern transit time is decreased from 2 to 1.5 min. However, larger delay savings are observed in the 2 min hold-pattern transit time case. When the hold delay is assigned as a

single continuous number there is smaller percentage delay saving over the FCFS scheduler. These numbers show the tendency of the optimal scheduler to avoid the larger than necessary delays produced by the airborne holds by keeping the number of airborne holds as low as possible via order changes. That is why the optimal scheduler gives maximum delay savings when the hold-pattern transit time is the largest and gives minimum delay savings when aircraft are held for a single continuous time.

C. Effect of Using Different Limits on the Maximum Allowable Sector Delay Constraint

In Sec. II.A.6, we noted that the amount of delay that can be absorbed in a sector is a function of the sector size and shape as well as the complexity of traffic flows through the sector. We studied the effect of using different values (1, 2, and 3 min) for the maximum delay absorption limit [also called the allowed maximum delay time (AMDT)] in each sector on the performance of both the FCFS and the optimal scheduling algorithms. Separate tests were carried out for this purpose. Table 5 shows the results of these tests. It is observed that with a larger AMDT the number of holds commanded by both the algorithms goes down and the delay savings go up, because with a larger AMDT more advantage can be gained by making order changes without causing those aircraft which are assigned to a lower position than their FCFS position to be held.

D. Effect of Using Different Values for the Maximum Position-Shift Constraint

Another parameter of the optimal scheduling algorithm is the maximum position shifts allowed from the FCFS sequence. As discussed in Sec. II.A.5, the MPS constraint is kept at a small number (2 or 3) to avoid causing controller overload. We tested the performance of the optimal scheduling algorithm for different values of the MPS constraint to see of any advantage can be gained by commanding larger position shifts. Table 6 shows the results of these tests. It is seen that very little benefit can be obtained by using a higher value than two for the MPS constraint. For the aircraft mixes used here, there was no advantage to be gained by using more than three maximum position shifts from the FCFS order.

E. Effect of ETA Prediction Uncertainty

1. Comparison of Mean Assigned Delays per Aircraft

Table 7 shows the delay savings achieved by the optimal algorithm over FCFS scheduling for different values of ETA prediction uncertainty. It is observed that both the mean average FCFS delays and the mean average optimal delays per aircraft go up with percentage uncertainty. This is due to the fact that additional delay is required each time an aircraft's updated ETA estimate puts it in violation of minimum separation constraint with another aircraft. It is also observed that the percentage delay savings over FCFS scheduling go down with uncertainty but still remain at a substantial level. This is an indicator that the optimal scheduling algorithm is robust in the sense that despite the ETA prediction uncertainty the algorithm still produces substantial delay savings over FCFS scheduling.

Table 6 Effect of using different maximum position-shift constraints (arrival scenario: average arrival rate =1 aircraft per min, aircraft mix=mix 1)

Maximum position shifts from FCFS allowed	FCFS		Optimal		
	Mean delay per aircraft, s	Mean no. of holds	Mean delay per aircraft, s	Mean no. of holds	Mean average delay savings (as a percentage of FCFS delay)
2	252	52.52	186	27.73	26.83%
3	252	52.52	184	26.93	27.21%
4	252	52.52	184	26.93	27.21%

Table 7 FCFS and optimal mean average delays and savings vs ETA prediction uncertainty (arrival scenario: average arrival rate of one aircraft per 60 s, mix 1) (all delays in seconds)

% Uncertainty	0	1	2	3.3333
Overall FCFS delay	252	254	273	277
Overall optimal delay	186	202	221	235
Overall % delay saving	26.83%	22.29%	20.45%	17.03%
FCFS delay (heavy aircraft)	57	59	63	70
Optimal delay (heavy aircraft)	57	63	74	80
% delay saving (heavy aircraft)	0.92	-6.85	-17.04	-15.46
FCFS delay (large aircraft)	107	108	120	128
Optimal delay (large aircraft)	100	112	127	144
% delay saving (large aircraft)	6.08%	-3.39%	-6.22%	-12.55%
FCFS delay (small aircraft)	580	580	612	614
Optimal delay (small aircraft)	375	402	427	441
% delay saving (small aircraft)	35.31%	30.71%	30.19%	28.19%

Table 8 Mean number of position shifts from the previously computed optimal sequence (sequence “jumps”) (arrival scenario: average arrival rate of one aircraft per 60 s)

Uncertainty level = 1%	Uncertainty level = 2%	Uncertainty level = 3.33%
Sequence jumps mean (max) 5.87 (58)	Sequence jumps mean (max) 7.92 (60)	Sequence jumps mean (max) 10.46 (60)

2. Sequence “Jumps” from the Optimal Sequence Computed at the Previous Scheduling Run

Table 8 shows the mean number of position shifts from the previously computed optimal sequence commanded by the scheduling algorithm over 30 scheduling runs. As is expected, the instability of the schedule in terms of unwanted sequence “jumps” goes up with the level of uncertainty. Still, the maximum number of position shifts over 30 scheduling runs (shown in parentheses) is 60 in the presence of ± 1 min (3.33% of freeze-region transit time) uncertainty in ETA predictions, which translates to two position shifts per scheduling run. It was observed that around 40 aircraft were involved in the scheduling process, on an average, over the simulation period. It is expected that there are only two position jumps per scheduling run over this set of aircraft, in the presence of ± 1 ETA prediction uncertainty.

VIII. Conclusions

An optimal scheduling algorithm was developed. This algorithm consists of two interconnected modules. One module computes optimal acceptance rates at all metering points in the airspace. The main innovation in this work is the use of an Eulerian model-based optimization scheme to compute these acceptance rates. The second module is an optimal scheduling and sequencing algorithm. This algorithm applies the acceptance rate constraints computed by the first module and computes optimum order changes from the first-come, first-served order to minimize overall delays. It includes a provision for “holding” aircraft at a point upstream of the freeze horizon for the optimum number of hold cycles if the total delay required is beyond the capacity of the freeze-region sectors. The algorithm is flexible in that it can handle a variety of airspace geometries, both terminal and en route, and it can deal with airspace sectors differing in their lengths and in their delay absorption capacities. The flexibility in defining the cost of the optimization makes this algorithm suitable for achieving a large variety of air-traffic optimization objectives, such as minimizing passenger cost or fuel cost and including user preferences.

The algorithm, with its two modules put together, was tested in a dynamic scheduling environment. The algorithm was able to achieve sizeable delay savings over the FCFS scheduling algorithm for the different aircraft arrival scenarios tested. A key performance indicator, the number of “jumps” in consecutive optimum aircraft sequence computations, is also low, thus demonstrating that our algorithm produces stable aircraft sequences. It was also observed

that optimal scheduling required a much smaller number of aircraft to be assigned to holding patterns as compared to FCFS scheduling. The effect of expected time of arrival prediction uncertainty on the performance of the scheduling algorithm was also tested. To the authors’ knowledge, this aspect of scheduling algorithms has not been tested before. It was observed that even in the presence of ETA prediction uncertainties, aircraft sequences remained stable and delay savings over FCFS scheduling remained sizeable.

Because of the innovative idea of using an Eulerian model-based optimization scheme to include downstream acceptance rate constraints into the scheduling process, scheduling decisions made at the present do not create untenable situations in the future when the aircraft get closer to the airport. Also, computation of optimal acceptance rates coupled with optimal sequencing maximizes the airport landing rates, while enforcing satisfaction of minimum separation constraints at the reference points and sector-density constraints. The dynamic airspace environment created as a part of this work could be used as a “plug and play” environment to test the performance of a variety of scheduling algorithms and to examine scheduling under different optimization objectives and different arrival scenarios.

Acknowledgment

This research was supported by NASA Ames Research Center under grants NAG 2-1581, and NNA06CB04G. Any opinions, findings, or conclusions expressed in this paper are those of the authors and do not necessarily reflect the views of NASA.

References

- [1] Farley, T. C., Foster, J., Hoang, T., and Lee, K., “A Time-Based Approach to Metering Arrival Traffic to Philadelphia,” AIAA Paper 2001-5241, Oct. 2001.
- [2] Kozarsky D., and Wang, J. J., “Final Life-Cycle Cost/Benefit Assessment of Multi-Center TMA,” NASA Ames TR NAS2-03145, 2004.
- [3] Farley, T. C., Hoang, T., Landry, S. J., Nickelson, M., Levin, K. M., Rowe, D., and Welch, J. D., “Multi-Center Traffic Management Advisor: Operations Test Results,” AIAA Paper 2005-7300, Sept. 2005.
- [4] Landry S., and Farley, T., “Distributed Scheduling Architecture for Multi-Center Time-Based Metering,” AIAA Paper 2003-6758, Nov. 2003.
- [5] Hoang, T., Farley, T., Foster, J., and Davis, T., “The Multi-Center TMA System Architecture and its Impact on Inter-facility Collaboration,” AIAA Paper 2002-5813, Oct. 2002.

- [6] Brinton, C. R., "An Implicit Enumeration Algorithm for Arrival Aircraft Scheduling," *The Eleventh Digital Avionics Systems Conference*, IEEE, Piscataway, NJ, Oct. 1992.
- [7] Isaacson, D. R., Davis, T. J., and Robinson, J. E., III., "Knowledge-Based Runway Assignment for Arrival Aircraft in the Terminal Area," *AIAA Guidance, Navigation, and Control Conference*, AIAA, Reston, VA, Aug. 1997.
- [8] Beasley, J. E., Krishnamoorthy, M., Sharaiha, Y. M., and Abramson, D., "Scheduling Aircraft Landings—The Static Case," *Transportation Science*, Vol. 34, No. 2, May 2000, pp. 180–197. doi:10.1287/trsc.34.2.180.12302
- [9] Beasley, J. E., Sonander, J., and Havelock, P., "Scheduling Aircraft Landings at London Heathrow Using a Population Heuristic," *Journal of the Operational Research Society*, Vol. 52, No. 5, 1 May 2001, pp. 483–493.
- [10] Beasley, J. E., Krishnamoorthy, M., Sharaiha, Y. M., and Abramson, D., "Displacement Problem and Dynamically Scheduling Aircraft Landings," *Journal of the Operational Research Society*, Vol. 55, No. 1, Jan. 2004, pp. 54–64.
- [11] Carr, G. C., Erzberger, H., and Neuman, F., "Airline Arrival Prioritization in Sequencing and Scheduling," *Second USA/Europe Air Traffic Management R&D Seminar*, Orlando, FL, Dec. 1998.
- [12] Carr, G. C., Erzberger, H., and Neuman, F., "Delay Exchanges in Arrival Sequencing and Scheduling," *Journal of Aircraft*, Vol. 36, No. 5, 1999, pp. 785–791.
- [13] Saraf A. P., and Slater, G. L., "An Efficient Combinatorial Optimization Algorithm for Optimal Scheduling of Aircraft Arrivals at Congested Airports," *Institute of Electrical and Electronics Engineers Paper 2006-1652*, March 2006.
- [14] Menon, P. K., Sweriduk, G. D., and Bilimoria, K. D., "New Approach for Modeling, Analysis, and Control of Air Traffic Flow," *Journal of Guidance, Control, and Dynamics*, Vol. 27, No. 5, Sept.–Oct. 2004, pp. 737–744.
- [15] Menon, P. K., Sweriduk, G. D., and Bilimoria, K. D., "A New Approach for Modeling, Analysis and Control of Air Traffic Flow," *Proceedings of the AIAA Guidance, Navigation and Control Conference*, AIAA, Reston, VA, Aug. 2002.
- [16] Menon, P. K., Sweriduk, G. D., and Bilimoria, K. D., "Air Traffic Modeling, Analysis and Control," *AIAA Paper 2003-5712*, Aug. 2003.
- [17] Menon, P. K., Sweriduk, G. D., Lam, T., Diaz, G. M., and Bilimoria, K. D., "Computer Aided Eulerian Air Traffic Flow Modeling and Predictive Control," *AIAA Paper 2003-5317*, Aug. 2004.
- [18] Bayen, A. M., Raffard, R. L., and Tomlin, C. J., "Adjoint-Based Control of a New Eulerian Network Model of Air Traffic Flow," *IEEE Transactions on Control Systems Technology*, Vol. 14, No. 5, 2006, pp. 804–814. doi:10.1109/TCST.2006.876904
- [19] Bayen, A. M., Raffard, R. L., and Tomlin, C. J., "Eulerian Network Model of Air Traffic Flow in Congested Areas," *Proceedings of the 2004 American Control Conference*, IEEE, Piscataway, NJ, June–July 2004.
- [20] Bayen, A. M., Raffard, R. L., and Tomlin, C. J., "Adjoint Based Constrained Control of Eulerian Transportation Networks: Application to Arrival Traffic Control," *Proceedings of the 2004 American Control Conference*, IEEE, Piscataway, NJ, June–July 2004.
- [21] Lighthill M. J., and Whitham, G. B., "On Kinematic Waves, 1: Flood Movement in Long Rivers," *Proceedings of the Royal Society of London A*, Vol. 229, 1955, pp. 281–316. doi:10.1098/rspa.1955.0088
- [22] Lighthill M. J., and Whitham, G. B., "On Kinematic Waves, 2: A Theory of Traffic Flow on Long Crowded Roads," *Proceedings of the Royal Society of London*, Vol. 229, 1955, pp. 317–345. doi:10.1098/rspa.1955.0089
- [23] Saraf A. P., and Slater, G. L., "Combined Eulerian Lagrangian Two-Level Control System for Optimal Air Traffic Flow Management," *AIAA Paper 2006-6229*, Aug. 2006.
- [24] Nilsson, N. J., *Principles of Artificial Intelligence*, Morgan Kaufmann, Los Altos, CA, 1980.
- [25] Bilimoria, K., Sridhar, B., Chatterji, G., Sheth, K., and Grabbe, S., "FACET: Future ATM Concepts Evaluation Tool," *Third USA/Europe Air Traffic Management R&D Seminar*, Napoli, Italy, June 2000.
- [26] Maciejowski, J. M., *Predictive Control With Constraints*, Morgan Kaufmann, Los Altos, CA, 1980.
- [27] Wong, G. L., "The Dynamic Planner: The Sequencer, Scheduler, and Runway Allocator for Air Traffic Control Automation," *NASA Ames Research Center TM-2000-209586*, Moffett Field, CA, April 2000.