Feedback Control of the National Airspace System to Mitigate Weather Disruptions

Jerome Le Ny and Hamsa Balakrishnan

Abstract—This paper proposes a general modeling framework adapted to the feedback control of traffic flows in Eulerian models of the National Airspace System (NAS). It is shown that the problems of scheduling and routing aircraft flows in the NAS can be posed as the control of a network of queues with load-dependent service rates. We can then focus on developing techniques to ensure that the aircraft queues in each airspace sector, which are an indicator of the air traffic controller workloads, are kept small. This paper uses the proposed framework to develop control laws that help manage the NAS during a weather event, given a probabilistic forecast of capacities. We also address the management of airport arrivals and departures subject to runway capacity constraints, which are highly sensitive to weather disruptions.

I. Introduction

The frequent occurrence of air traffic delays in the National Airspace System (NAS), along with the projected increase in demand, motivate the scheduling of flight operations to better utilize available system resources. The process of planning operations a few hours ahead of time in order to balance the available capacity and the demand for resources is known as Traffic Flow Management (TFM). This task is currently conducted by air traffic controllers (ATC) manually, and contributes significantly to their workload. In order to meet the increasing traffic demand, there is a desire to introduce a greater level of automation and decision support for air traffic management.

Research on the TFM problem initially focused on developing open-loop policies prescribing the position of each aircraft in the system at each instant, and obtained by solving large-scale integer programs [1], [2]. This approach is clearly difficult to scale to the scheduling of approximately 40,000 flights a day, and generally pays only little attention to the many sources of uncertainty

present in the system. Weather in particular is a major source of disturbances. For instance, 66% of all NAS delays in 2009 were attributed to weather [3]. Due to the typical travel times of cross-country flights, open-loop traffic flow management requires planning horizons of 5-6 hours, which are arguably beyond the limits of even state-of-the-art weather forecasting tools.

The disturbance attenuation properties of feedback control make closed-loop control policies for the NAS very attractive. Early attempts to introduce a limited amount of feedback can be found in [4], building on the previous integer programming-based formulations. More recently, researchers have started developing new models more tractable for the purpose of control, which only record aircraft counts in specific control volumes of airspace rather than follow each aircraft trajectory. These aggregate models, called *Eulerian models*, are gaining popularity [5]–[9] and have been shown to have reasonable predictive capabilities [7]. Some first attempts at feedback control using Eulerian models have also been made, both in the context of centralized traffic flow management [6], [10], and in a decentralized setting for networks with a single origin and destination [11]. In this paper, we extend our previous work [12] and present an Eulerian model for TFM that can be used to control all resources of the NAS, rather than focusing on high-altitude traffic as was considered in [5]–[7]. The inclusion of airports is particularly important because they are typically the bottlenecks of the system. Moreover, it can be shown that our general model, presented in section II-C, captures as a special cases other recent Eulerian models such as the CTM(L) model [7], while offering additional modeling flexibility.

At the TFM level, our model provides high-level directives to the ATCs prescribing the desired flow rates of aircraft traveling through restricted resources of the NAS. These resources can be runways, metering points at certain airspace fixes, or flow control lines at the boundaries of sectors or of flow constrained areas during Airspace Flow Programs (AFPs) [13]. ATCs can implement these directives by issuing orders at the path planning level, such as aircraft speed changes, vector

for spacing or holding patterns, which modify the time it takes for aircraft to travel between control boundaries. The specific choices adopted by the ATC at the tactical level require the precise knowledge of the current spatial configuration of the aircraft and are not specified at the TFM level.

The ability to integrate weather forecasts is perhaps the single most desirable feature of TFM procedures. Local weather events can be greatly amplified by network effects without proper congestion control. We can use our general control model to develop a tractable model-predictive control strategy that can plan using probabilistic weather forecasts.

The rest of the paper is organized as follows. Section II describes the general Eulerian model that we use for TFM. Section III describes some natural control policies for this model, including a distributed network control policy and a model predictive control strategy. In section IV we see how the model can be used to mitigate the impact of weather disruptions, given probabilistic forecasts. We also discuss simulation results for the TFM problem over part of the Western United States.

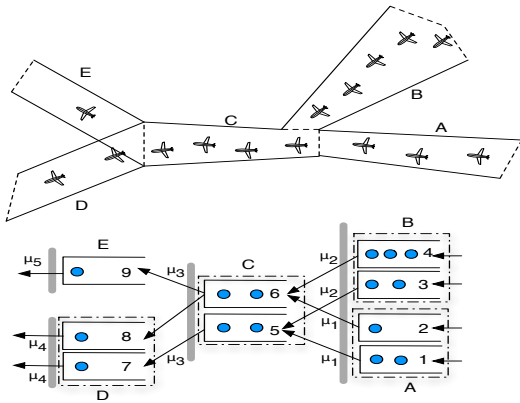## II. EULERIAN MODEL OF THE NAS ADAPTED TO NETWORK CONTROL

Fig. 1. Five control volumes and their corresponding abstract buffer model. The situation depicted here corresponds to two routes merging, and all aircraft leaving through volume C. Each control volume except E supports two flows with distinct characteristics (e.g., different destinations, see section II-B.1), and each flow is associated with one buffer. Note that the flow type of buffers 8 and 9 is the same, which allows us to include routing decisions at the output boundary of volume C. Each buffer has an associated maximum throughput function. Thick gray lines represent additional resource constraints coupling the control vectors of different buffers, see sections II-B.1 and II-B.2. For example, in this case, the buffers associated with the same control volume have the same maximum throughput function.

To construct the network model that we use for TFM, we start by deciding the points or lines through which

traffic flow rates need to be determined. These boundaries, called hereafter *control boundaries*, can consist of sector boundaries, runways, airspace fixes, intersections of major jet routes, or other metering points. Adding more control boundaries provides more decision support, but decreases the flexibility to adapt (at the tactical level) to factors not modeled precisely at the TFM level. Each control boundary has an associated traffic flow direction.

In Eulerian models, we are interested in controlling the aircraft counts in certain *control volumes* rather than individual aircraft trajectories. A control volume is delimited by an input and output control boundary, that is, all the traffic associated with it enters through the same control boundary and exits through the same one. Control volumes can overlap, for example, due to intersecting traffic flows. Within a control volume, we have one or more queues, also called buffers. Using several queues in the same control volume allows us to separate the traffic based on distinct characteristics, such as destinations. Fig. 1 shows an example of model with five control volumes and two distinct flows.

A. Maximum Throughput of a Single Buffer

We consider discrete-time models, and assume that a suitable time period T is chosen to model the system dynamics. This period also dictates the frequency of updates to flow rate directives. The number of aircraft in buffer i at time kT is denoted $Q_i(k)$. Note that all buffers are associated with a control volume (see Fig. 1), and hence to an input and output boundary. The *maximum* rate at which traffic can flow out of a buffer during any time period depends on the number of aircraft that it contains. The aircraft count associated with buffer i follows the dynamics

$$Q_i(k+1) = Q_i(k) + A_i(k) + \sum_{j \in \mathcal{J}} U_j(k) - U_i(k), \quad (1)$$
$$0 \leq U_i(k) \leq D_i(k), \ \forall k \geq 0. \quad (2)$$

Here, $A_i(k)$ is the number of *external* arrivals in buffer i during the time interval $[kT, (k+1)T)$, originating from unmodeled parts of the system (e.g., pushbacks from airport gates). $U_i(k)$ is the number of departures from buffer i during the same period (the $(k+1)^{\text{th}}$ period), and is controlled by the ATC. $\mathcal{J}$ is the set of buffers sending aircraft to buffer i. These terms will be discussed in more detail later. In this section, we focus on the quantity $D_i(k)$, which is the *maximum* possible number of departures from buffer i during the $(k+1)^{\text{th}}$ period. Note that there is an additional non-negativity constraint, $Q_i(k) \geq 0$, which will generally be automatically satisfied by imposing the condition $D_i(k) \leq Q_i(k) + A_i(k)$. We assume that $U_i(k)$ can

depend on $Q_i(k)$ as well as $A_i(k)$ and $D_i(k)$, therefore these quantities must be known prior to determining $U_i(k)$. Knowing $D_i(k)$ requires the accurate prediction of at most how many aircraft from buffer i can cross the exit boundary of the control volume during the $(k+1)^{\text{th}}$ period, a simple trajectory prediction problem for the typical values of T used in flow control (1-15 min).

In general, the travel times of aircraft through a control volume vary due to differences in speeds, trajectories and environmental factors such as wind speed. Since it is not tractable to keep track of all these variations exactly, we treat them as disturbances on a nominal aggregate model. We assume a stochastic model of arrivals and departures in (1), (2). Moreover, we assume that the sequence of bounds $\{D_i(k)\}_{k \geq \mathbb{N}}$ can be random, but such that for a given load level $Q_i(k) = Q_i$, the variables $\{D_i(k)\}$ are identically distributed, with mean denoted $\mu_i(Q_i) = \mathbb{E}[D_i(k)|Q_i(k) = Q_i]$. We call the function μ_i the *maximum throughput function* for buffer i. Usually, the same function μ_i can be used for all buffers in the same control volume, unless traffic flows are separated into different buffers based on characteristics such as aggregate trajectories or velocities. In order for this to be a reasonable model of the maximum number of aircraft departures in a period, the control volume should not be too small relative to the sampling period, otherwise several aircraft can enter and leave the volume during the same period. A significant number of such aircraft would require that the function μ depend on $A(k)$.

Intuitively, μ_i should increase with Q_i, since an increase in traffic along a route reduces the separation between departures. However, μ_i remains bounded due to the minimum required separation distance between aircraft, which limits the rate at which aircraft can cross the exit boundary of the control volume. In general, we expect μ_i to be a concave saturating function, as depicted on Fig. 2. The exact values of the function depend on the geometry of the control volume and the typical aircraft trajectories between its boundaries. The maximum throughput curve for a given buffer or control volume can be obtained via simulation, by fitting empirical data, or perhaps a combination of both. Note that $\mu(0) = 0$, and $\mu(1)$ is roughly inversely proportional to the typical minimum travel time of an aircraft through the region (measured in number of periods of length T). For the simulation results shown on Fig. 2 we assumed that aircraft could travel at up to 500 knots, that the control volume length is 100 nm with a simple narrow linear geometry where all planes strictly follow each other, and used ATC directives asking that aircraft entering the volume set their velocity to the maximum
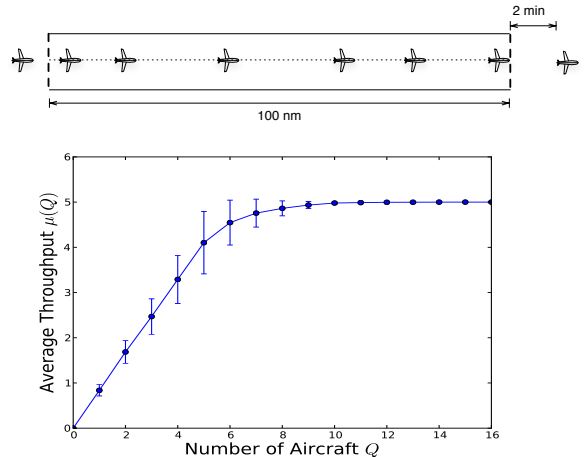

Fig. 2. Maximum throughput for the exit boundary of a simple linear control volume. The empirical curve of the mean number of departures per period was obtained via discrete-event simulation, with random arrivals. The error bars show the empirical variance of the number of departures per period.

possible while respecting the separation constraint with the previous aircraft. The sampling period T was 10 min. Note that for the purpose of evaluating μ_i, we can consider the situation where the control volume contains only buffer i. The interaction among flows of different buffers within the same control volume are modeled as additional constraints on the control variables U_i, as discussed in section II-B.1.

In Fig. 2, we also assumed that successive aircraft crossing the exit boundary of this control volume are separated by at least 2 min. This results in the saturation of the curve at 5 departures per period. For control volumes that are not subject to such explicit metering constraints, the saturation phenomenon still persists due to the mandatory separation between aircraft, currently set to 5 nm in enroute airspace. In our example, this would result in a curve saturating at about 16 aircraft per period instead of 5. Finally, we note that if the length of the time period (T) is changed from T_1 to T_2, the resulting curve $\mu_i(Q_i)$ can be obtained by scaling the curve for T_1 by T_2/T_1.

B. Additional Resource Constraints

We construct a network model of the NAS using control volumes containing buffers following the dynamics (1), (2) as building blocks. Traffic flows in different buffers compete for limited airspace resources, resulting in additional linear constraints on the control variables, which are described in the following paragraphs.

1) Shared Buffers within a Control Volume: For a control volume containing m buffers, the control

vector for the exit boundary is denoted by $U(k) = [U_1(k), \ldots, U_m(k)]^T$. If several distinct flows carry a significant number of aircraft, the bounds $D_i(k)$ can be large for several values of the index i and the exit boundary of the volume might not be able to accommodate $\sum_{i=1}^m D_i(k)$ departures in a single period. It may therefore be necessary to give priority to certain flows over others, i.e., to *schedule* the flows at the control boundaries. We add scheduling constraints of the form

$$c^T U(k) = \sum_{i=1}^m c_i U_i(k) \leq r(k), \quad \forall k \geq 0. \qquad (3)$$

In general we take $c_i = 1$ for all i, and $r(k)$ is simply the maximum number of aircraft that can cross the boundary at period k. However, we may prioritize certain flows by varying c_i, and adjust r appropriately.

2) Intersecting and Merging Flows: Consider the situation depicted on Fig. 1. The input boundary of sector C coincides with the output boundaries of sectors A and B. U^1 and U^2 denote the control vectors associated with the latter two regions, which support two flows each. At the merge point, suppose we cannot accommodate the sum of the maximum flow rates of A and B, and increasing the flow rate out of one volume requires reducing the flow rate out of the other. This aspect can be incorporated by imposing linear constraints on the control vectors of the form

$$c_1^T U^1(k) + c_2^T U^2(k) \leq r(k), \qquad (4)$$

for some vectors c_1, c_2, and some scalar r. Here c_1, c_2 can be all-one vectors and r represents the maximum number of aircraft that can enter volume C per period. The contraint (4) is not active if the volumes are only lightly loaded and the resulting bounds in (2) are small.

Intersections of major jet routes can be handled similarly, and competition between flows for passage through limited airspace resources can be modeled by additional linear constraints of the form

$$C(k)U(k) \leq R(k), \qquad (5)$$

where $C(k)$ is a matrix, U is the vector of all control variables for the problem, and $R(k)$ is a vector.

3) Routing: We can easily can incorporate routing decisions into the scheduling model presented so far, by adding control variables. Allowing ATCs to make tactical routing decisions can help accommodate dynamically changing conditions in the network, such as the impact of weather on capacities. In the model of Fig. 1, aircraft from buffer 6 in control volume C can enter buffer 8 or

9. The dynamics of the buffer 6 are then

$$\begin{aligned} Q_6(k+1) = \quad & Q_6(k) + A_6(k) + U_2(k) + U_4(k) \\ & -(U_{6,8}(k) + U_{6,9}(k)) \end{aligned}$$

$$U_{6,8}(k) \geq 0, U_{6,9}(k) \geq 0, U_{6,8}(k) + U_{6,9}(k) \leq D_6(k).$$

Here $U_{6,8}$ (resp. $U_{6,9}$) represents the number of aircraft routed by the ATC from buffer 6 to buffer 8 (resp. 9).

4) Sector Load Capacities: We can add bounds on the vectors $Q(k)$ to impose limits on the capacity of the sectors. In general, these constraints take the linear form

$$MQ(k) \leq S(k), \qquad (6)$$

where M is a matrix and $S(k)$ a vector.

5) Airport Resources: To model an airport, arrival queues of aircraft waiting to land are associated with air traffic flows in the close vicinity of the airport, and one or several departure queues contain the aircraft on the ground waiting to take-off. All queues have dynamics of the form (1), (2). Arrivals and departures at an airport share ground resources. Consequently, the arrival and departure control vectors are also subject to resource constraints. If U_d is the control vector for the departure queues and U_a is the control vector for the arrival queues, the global vector $U = [U_a^T, U_d^T]^T$ is again subject to linear constraints of the form (5), as discussed by Gilbo [14], [15]. These constraints depend on runway configuration and can be determined empirically or analytically [14], [16]. Note again that these linear constraints (5) are not necessarily active in low traffic conditions due to the more constraining conditions (2). For the maximal departure throughput μ_d, we can consider a model incorporating aircraft from the time they pushback from the gate [17], [18]. In this model, we have a nonlinear maximal throughput curve for departures that depends on the number of aircraft in transit between the gate and the runway. The airport transit zone then constitutes the control volume. The throughput curve obtained is again of the same form as the one shown on Fig. 2 [17, chap. 5].

6) Building Larger Models: Using the basic building blocks presented in the previous sections, it is straightforward to construct larger networks adapted to the control of traffic flows in the NAS. An example is described in section IV. Consider, for example, a model incorporating the scheduling and routing of aircraft between origin-destination (O-D) pairs along a set of possible preferred routes. Within control volumes, we separate flows into distinct buffers based on their O-D pair, indexed by m. A flight trajectory corresponds to a path through the set of control boundaries. We can index buffers by their associated input and output boundaries

as well as O-D pair, hence Q_{ij}^m is the load for O-D pair m in the control volume going from boundary i to boundary j. The dynamics of this buffer are then

$$Q_{ij}^m(k+1) = Q_{ij}^m(k) + A_{ij}^m(k) \qquad (7)$$
$$+ \sum_{s \in I(i,m)} U_{sij}^m(k) - \sum_{t \in O(j,m)} U_{ijt}^m(k),$$

$$\sum_{t \in O(j,m)} U_{ijt}^m(k) \le D_{ij}^m(k),\ 0 \le U_{ijt}^m(k),\ \forall t \in O(j,m),$$

where $I(i,m)$ and $O(j,m)$ are the set of all control boundaries preceding i, and the set of all control boundaries following j respectively, that are on an allowed route between the O-D pair m. The variable U_{ijt}^m corresponds to aircraft of O-D pair m routed through the control boundaries i,j and then through t. Other choices of models are discussed in [12].

C. General Discrete Model

In the previous sections we have seen that, after discretization of the airspace into control volumes carrying flows separated based on characteristics relevant to the TFM problem, we obtain a network model with dynamics which can be written in matrix form as

$$Q(k+1) = Q(k) + B\,U(k) + A(k), \quad \forall k \ge 0 \quad (8)$$
$$M\,Q(k) \le S(k), \quad \forall k \ge 0 \qquad\qquad (9)$$
$$C_1 U(k) \le D(k), \quad \forall k \ge 0 \qquad\qquad (10)$$
$$C_2(k)U(k) \le R(k), \quad \forall k \ge 0 \qquad\quad (11)$$
$$U(k) \ge 0, \quad \forall k \ge 0. \qquad\qquad\qquad (12)$$

The matrix B consists of $+1$'s and -1's and is essentially the incidence matrix [19] for a graph whose nodes are the buffers and edges are present between successive buffers. In addition, there are nonnegativity constraints $Q(k) \ge 0$, which can usually be automatically enforced through the constraints (10), see subsection II-A. The difference between model (8) and queueing network models studied in the literature on communication or manufacturing networks [20] is the addition of the load dependent constraint (10), modeling the fact that the number of aircraft in a sector influences the maximum rate at which aircraft leave the sector. However, we note that if all components of $Q(k)$ are large, then $D(k)$ tends to a constant vector, hence for high loads and when considering stability issues, we expect the analysis to be close to the one developed for standard queueing network models.

III. Control Strategies

We now need to develop algorithms to determine appropriate decisions $U(k)$ for the closed loop control of the system (8). Below, we describe two natural control strategies, with very different computation and implementation requirements, that can be used for scheduling and routing aircraft through the NAS. The first strategy, the MaxWeight policy, is a distributed policy with very few implementation requirements that could be used under normal conditions, while the second approach is a more demanding MPC controller that could be used in rapidly-changing capacity scenarios, such as bad weather days, as discussed in section IV.

A. Distributed MaxWeight Policy

The celebrated MaxWeight or maximum back pressure policy for network control [21], proposed for air traffic control in [12], can be obtained as follows. At time period k, consider some control variables $U_{i_1}, U_{i_2}, \ldots, U_{i_b}$ that are coupled by resource constraints from (10) and (11). Starting from $U_i = 0, i = i_1, \ldots, i_b$, we increase the variables by unit increments, until none of them can be increased any further because some constraint would then be violated. At each step, we increment the control variable j for which the back pressure $Q_{\sigma(j)} - Q_{\sigma(j)_+}$ is maximum, where $\sigma(j)$ is the index of the buffer associated with j, and $\sigma(j)_+$ is the index of the buffer following $\sigma(j)$ for control j (recall that several control variables can be associated to the same buffer if rerouting is allowed). Weighted versions of this policy can be used to give preferences to certain routes over other. This policy is appealing because it can be implemented locally by the ATC using communication only with the neighboring sectors. See [12] for a more detailed discussion of this control law.

B. Model Predictive Control

Model Predictive control (MPC) [22] is a general tool that is well suited for the feedback control of constrained systems of the form (8). Menon et al. previously proposed to use MPC to control a deterministic Eulerian model of the NAS [6]. We use MPC for the weather scenario considered in section IV. To obtain a feedback control law via MPC for the general problem (8), we proceed as follows. At period k_0, we can observe the state $Q(k_0)$, and the maximum possible number of arrival and departures $A(k_0)$ and $D(k_0)$ for the period. We fix a horizon length $K \ge 0$, and choose $U(k_0)$ by solving the following convex program with variables

$$\mathbf{Q} = \{Q_k\}_{k_0+1 \leq k \leq k_0+K+1}, \mathbf{U} = \{U_k\}_{k_0 \leq k \leq k_0+K}$$

$$\min_{\mathbf{Q},\mathbf{U}} \quad f(\mathbf{Q}) \tag{13}$$

subject to $\qquad\qquad\qquad\qquad\qquad\qquad$ (14)

$$Q_{k_0+1} = Q_{k_0} + BU_{k_0} + A(k_0)$$

$$Q_{k+1} = Q_k + BU_k + \alpha(k), k_0 + 1 \leq k \leq k_0 + K$$

$$MQ_k \leq S(k), \quad k_0 + 1 \leq k \leq k_0 + K + 1 \tag{15}$$

$$C_1 U_{k_0} \leq D(k_0) \tag{16}$$

$$C_1 U_k \leq \mu(Q_k), \quad k_0 + 1 \leq k \leq k_0 + K \tag{17}$$

$$C_2(k) U_k \leq R(k), \quad k_0 \leq k \leq k_0 + K, \tag{18}$$

$$U_k \geq 0, \quad k_0 \leq k \leq k_0 + K,$$

where f is chosen to be a convex function of $\mathbf{Q}$, for example a linear objective

$$f(\mathbf{Q}) := \sum_{k=k_0+1}^{K} c_k^T Q_k + c_{K+1}^T Q_{K+1}, \tag{19}$$

for some vectors $\{c_k\}_k$. We use the certainty-equivalence heuristic [23] which consists of replacing $A(k)$ and $D(k)$ of (8) by their average values $\alpha(k)$ and $\mu(Q_k)$. The program (13) is convex if we assume that μ is concave, which is generally true, as seen in Fig. 2.

After solving (13), we the obtain a sequence of vectors $U_{k_0}, \ldots, U_{k_0+K}$, which are real-valued. We round the first vector U_{k_0} and use it as a control directive $U(k_0) = U_{k_0}$ for the current period. We discard the other vectors $U_{k_0+1}, \ldots, U_{k_0+K}$. At the next period, we repeat the procedure to obtain a new control, after observing the new values of Q, A, D. Solving the convex program (13) can be done via efficient interior point methods for various choices of objective and throughput functions [24]. For example, we can consider the linear objective function (19) and an approximation of the maximal throughput functions of the piecewise linear form

$$\mu_i(q_i) = \min_{1 \leq j \leq m_i} \{a_{ij}^T q_i + b_{ij}, \mu_{i,sat}\}.$$

Then the constraint $c_i^T \zeta \leq \mu_i(q_i)$ can be rewritten as m_i affine constraints

$$c_i^T \zeta \leq a_{ij}^T q_i + b_{ij}, \; j = 1, \ldots, m_i \tag{20}$$

$$c_i^T \zeta \leq \mu_{i,sat}.$$

Then the program (13) becomes a linear program. The hard constraint (15) is usually removed to avoid infeasibility issues, in which case we can try to respect capacity constraints by only penalizing sector loads in the objective function.

We consider a naive rounding procedure to obtain an integer-valued control vector U at each period (only making sure that the constraints (16) are satisfied, which is easy). The feasibility of MPC approach requires that the ATCs have enough time to implement the directives, and therefore needs the computation of (13) to be possible in a time much shorter than the time period T. This also requires a model that is not too large, hence a discretization that is not too fine, as our model allows. MPC can also benefit from recent advances in convex optimization in solving large-scale problems [25].

IV. WEATHER MANAGEMENT

Our discussion so far assumes that parameters such as the average capacities at control boundaries are known constants. However, if a region experiences bad weather, its capacity to accommodate traffic flows can be greatly reduced. Weather forecasts useful for detailed aircraft route planning are becoming available for TFM [26]. In this section we consider the systematic integration of probabilistic weather forecasts within the TFM system for the NAS. Currently, some limited forms of flow management procedures taking weather into account are Ground Delay Programs and Airspace Flow Programs [13], but we would like to generalize and coordinate such programs in a more systematic way, while incorporating real-time information in a feedback loop. Preparing a region of the NAS for fast recovery from a weather event is a problem similar to those encountered in production systems subject to machine breakdown [27], [28]. Here we extend the MPC approach of section III-B to the situation where the capacity constraints evolve randomly in time according to a probabilistic weather forecast.

A. A Model Integrating Weather Uncertainty

The weather state evolves as a Markov chain $\{w(k)\}_{k \geq 0}$ on a finite state-space W, with time-varying transition kernel $\mathbb{P}(w(k + 1) = w'|w(k) = w) = [P(k)]_{ww'}$. The transition kernel is assumed known, determined from probabilistic weather forecasts. The weather state influences some of the parameters appearing in the model of section II-C, namely the capacity vectors S, R and the matrix C_2. We continue to ignore the hard state constraint (9) in our algorithms, and focus instead on the control constraints (11). The vector R generally describes the number or aircraft that control boundaries can accomodate, hence clearly depends on the weather state, and is fixed for example by Airspace Flow Programs [13]. The fact the the matrix C_2 changes with the weather state is evident from the fact that this matrix includes the models of the capacity envelopes of airports (see section II-B.5). These envelopes vary in shape according to the weather state [16]. Hence

we replace the constraints (11) by $C_2(w(k))U(k) \leq R(w(k))$. Note that one could easily include weather-related changes in other parameters, for example the throughput functions μ.

B. Certainty-Equivalent MPC

We can now develop a certainty-equivalent model predictive controller (CE-MPC) similar to the one described in section III-B. Namely, at period k_0, we observe the weather state $w(k_0)$, and we replace constraint (18) by

$$C_2(w(k_0))U_{k_0} \leq R(w(k_0))$$
$$\mathbb{E}[C_2(w(k))|w(k_0)]U_k \leq \mathbb{E}[R(k)|w(k_0)],$$
$$k_0 + 1 \leq k \leq k_0 + K, \quad (21)$$

where the expectations in the constraints (21) can be computed recursively for the next K stages at the cost of essentially K matrix multiplications of size $|W| \times |W|$. The optimization problem is still convex, and a linear program under the assumptions stated in section III-B.

C. Simulation Results

We consider the TFM problem over a portion of the western part of the United States. The airspace modeled is primarily within the Oakland Air Route Traffic Control Center (ZOA), and we consider the airports at Los Angeles (LAX), San Francisco (SFO), Seattle (SEA), Portland (PDX), and Las Vegas (LAS). External traffic also enters the system, mainly via the major routes coming from the east towards SFO, LAX and SEA. After identifying the major routes supporting most of the traffic, we define control boundaries at which we would like to regulate traffic rates, and the corresponding control volumes. The resulting queueing network is shown on Fig. 3. It consists of 50 queues, with most of the control consisting of scheduling decisions, with a few routing decisions. The capacity envelopes for the airports were obtained from [16]. The throughput functions for the control volumes were only approximately identified, based on the length of the control volumes, the required separation distance (5 nm), and assuming simple traffic geometries in the volumes as in Fig. 2. All buffers have a limited capacity of 20 aircraft, and there are coupled capacity constraints for the buffers at the intersection East of SFO and in the arrival area of SFO.

For this system, we simulate a bad weather event around San Francisco that changes the capacity envelope at SFO airport according to the data given in [16]. We assume that after some initial period with optimal capacity, the weather changes to one of 3 states: low (L), medium (M) and high (H) weather impact (in the H state, SFO allows no departures and accepts only 10

arrivals per hour; M corresponds to IFR conditions and L to Marginal conditions in [16]). The chain can jump between these 3 states for some time, but eventually reaches a final absorbing state with optimal capacity. The transition structure of the chain is shown in Fig. 3.

We use the CE-MPC in the linear programming form described in section IV-B. The period T is set to 4 min, simulations are run for an 8hr-duration, i.e., 120 periods. The horizon K for MPC is set to 32min, since no significant performance improvement was obtained with a larger horizon. At each step, solving the linear program takes about 2s on a standard laptop with a generic Matlab implementation (we used CVX [29], without trying to improve the computational performance). The rest of the period would be left to the ATC to implement the directives. The simulation results, obtained from 100 simulations, are shown on Fig. 4. The MPC policy performs better than the MaxWeight policy, for example due to its better ability to manage busy intersections such as the one East of SFO.
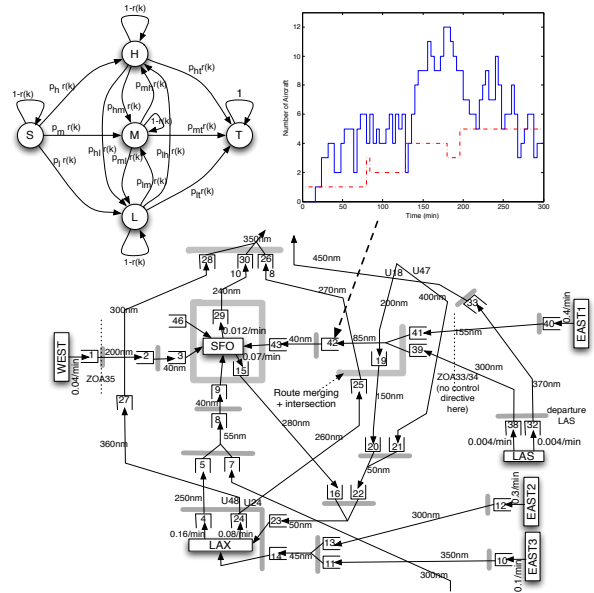

Fig. 3. Portion of the network model for TFM in the simulated system. Thick gray lines correspond to control variables coupled via linear constraints. All queues start initially empty. We also show the automaton describing the probabilistic weather evolution as well as a sample trajectory of a queue in the system (number of aircraft on the y-axis vs. time). The red dashed line shows the evolution of the weather state for that sample (namely $S - M - L - H - M - T$).

V. Conclusion

This paper discusses an improved Eulerian model that can be used to develop closed-loop control policies for the NAS and takes into account all air traffic resources,
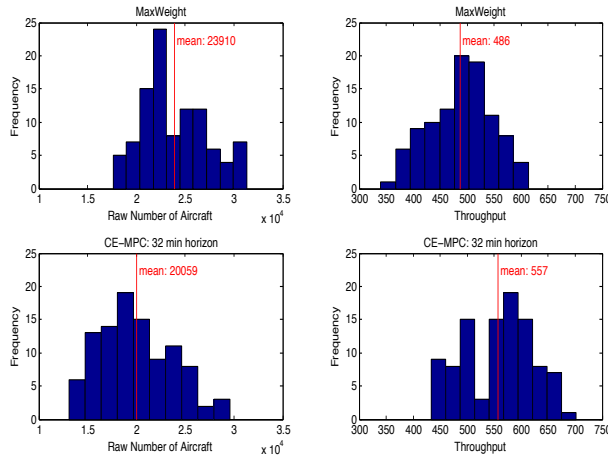
Fig. 4. Histogram of the Simulation Results over 100 simulations. The cost on the left is the cumulative number of aircraft in the system during the 8 hour period of the simulation (an aircraft is counted once per period where it is present in the system). The throughput on the right is the total number of aircraft that landed during the simulation.

including airport capacity envelopes. The model is very flexible and provides decision support to the ATC to control traffic flow rates at the control boundaries of their choice. We believe that the proposed model is particularly useful in developing planning strategies during extreme weather events.

REFERENCES

[1] M. Terrab and A. R. Odoni, "Strategic flow management for air traffic control," *Operations Research*, vol. 41, no. 1, pp. 138–152, Jan.-Feb. 1993.

[2] D. Bertsimas, G. Lulli, and A. Odoni, "The air traffic flow management problem: An integer optimization approach," in *Integer Programming and Combinatorial Optimization*, 2008, pp. 34–46.

[3] Bureau of Transportation Statistics. Understanding the reporting of causes of flight delays and cancellations. [Online]. Available: http://www.bts.gov/help/aviation/html/understanding.html

[4] P. B. M. Vranas, D. Bertsimas, and A. R. Odoni, "Dynamic ground-holding policies for a network of airports," *Transportation Science*, vol. 28, no. 4, pp. 273–291, November 1994.

[5] P. K. Menon, G. D. Sweriduk, and K. D. Bilimoria, "New approach for modeling, analysis, and control of air traffic flow," *Journal of Guidance, Control, and Dynamics*, vol. 27, no. 5, pp. 737–744, 2004.

[6] P. K. Menon, G. D. Sweriduk, T. Lam, G. M. Diaz, and K. Bilimoria, "Computer-aided Eulerian air traffic flow modeling and predictive control," *AIAA Journal of Guidance, Control and Dynamics*, vol. 29, pp. 12–19, 2006.

[7] D. Sun, I. S. Strub, and A. Bayen, "Comparison of the performance of four Eulerian network flow models for strategic air traffic management," *Networks and Heterogeneous Media*, vol. 2, no. 4, pp. 569–594, December 2007.

[8] S. Roy, B. Sridhar, and G. C. Verghese, "An aggregate dynamic stochastic model for air traffic control," in *Proceedings of the 5th USA/Europe ATM 2003 R&D Seminar*, Budapest, Hungary, 2003.

[9] B. Sridhar, T. Soni, K. Sheth, and G. Chatterji, "An aggregate flow model for air traffic management," in *AIAA Conference on Guidance, Navigation, and Control*, no. AIAA Paper 2004-5316, Providence, RI, 2004.

[10] A. Mukherjee and M. Hansen, "A dynamic rerouting model for air traffic flow management," *Transportation Research Part B: Methodological*, vol. 43, no. 1, pp. 159 – 171, 2009.

[11] H. Arneson and C. Langbort, "Distributed sliding mode control design for a class of positive compartmental systems," in *Proceedings of the American Control Conference*, 2008.

[12] J. Le Ny and H. Balakrishnan, "Distributed feedback control for an Eulerian model of the national airspace system," in *Proceedings of the American Control Conference*, 2009.

[13] M. Robinson, R. DeLaura, B. Martin, J. E. Evans, and M. E. Weber, "Initial studies of an objective model to forecast achievable airspace flow program throughput from current and forecast weather information," in *Aviation, Range and Aerospace Meteorology Special Symposium on Weather-Air Traffic Management Integration, AMS Annual Meeting*, Phoenix, AZ, January 2009.

[14] E. P. Gilbo, "Airport capacity: Representation, estimation, optimization," *IEEE Transactions on Control Systems Technology*, vol. 1, no. 3, pp. 144–154, September 1993.

[15] ——, "Optimizing airport capacity utilization in air traffic flow management subject to constraints at arrival and departure fixes," *IEEE Transactions on Control Systems Technology*, vol. 5, no. 5, pp. 490–503, September 1997.

[16] "Airport capacity benchmark 2004," U.S. Department of Transportation, Tech. Rep., 2004.

[17] R. Shumsky, "Dynamic statistical models for the prediction of aircraft take-off times," Ph.D. dissertation, Massachusetts Institute of Technology, 1995.

[18] I. Simaiakis and H. Balakrishnan, "Queuing models of airport departure processes for emissions reduction," in *AIAA Guidance, Navigation and Control Conference and Exhibit*, 2009.

[19] R. Diestel, *Graph Theory*, ser. Graduate Texts in Mathematics. Springer, 2005, vol. 173.

[20] S. Meyn, *Control Techniques for Complex Networks*. Cambridge University Press, 2008.

[21] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Transactions on Automatic Control*, vol. 37, no. 12, pp. 1936–1948, December 1992.

[22] C. E. Garca, D. M. Prett, and M. Morari, "Model predictive control: Theory and practice–a survey," *Automatica*, vol. 25, no. 3, pp. 335 – 348, 1989.

[23] D. Bertsekas, *Dynamic Programming and Optimal Control*, 3rd ed. Athena Scientific, 2005.

[24] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.

[25] Y. Wang and S. Boyd, "Fast model predictive control using online optimization," *IEEE Transactions on Control Systems Technology*, 2008, to appear.

[26] D. Michalek and H. Balakrishnan, "Identification of robust routes using convective weather forecasts," in *Eigth USA/Europe Air Traffic Management Research and Development Seminar (ATM2009)*, 2009.

[27] S. Gershwin, *Manufacturing Systems Engineerin*. Prentice-Hall, 1994.

[28] M. Chen, R. Dubrawski, and S. Meyn, "Management of demand-driven production systems," *IEEE Transactions on Automatic Control*, vol. 49, no. 5, pp. 686–697, 2004.

[29] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming. http://stanford.edu/ boyd/cvx.