Distributed Feedback Control for an Eulerian Model of the National Airspace System

Jerome Le Ny and Hamsa Balakrishnan

Abstract— This paper proposes an Eulerian model of traffic flows in the National Airspace System (NAS) and presents a distributed feedback control approach for managing the flows. The main contribution is the development of a model that reflects the way that air traffic controllers regulate traffic flows into their facilities, and of a feedback control approach that provides them with decision support, in a distributed manner that is consistent with the existing communication structure. The focus is on developing techniques that guarantee that the aircraft queues in each airspace sector, which are an indicator of air traffic controller workload, are kept small. It is shown that the problems of scheduling and routing aircraft flows in the NAS can be posed as the fluid approximation of a network of queues, and that under appropriate conditions, a MaxWeight policy can be used to determine a distributed feedback control law that stabilizes the system. Extensions to the problem of airport arrival/departure management, and to traffic flows that are driven by demand are also described.

I. INTRODUCTION

The current levels of congestion in the National Airspace System (NAS) with the associated delays, and the predicted increases in demand, motivate the more efficient utilization of system resources and capacities. Scheduling flights in order to best balance the capacity of and demand for resources, also known as traffic flow management, involves the coordination of operations in a large number of facilities, such as airports (there are about 650 airports with at least one commercial domestic operation per day) and airspace sectors (541 low, 490 high, and 231 super-high altitude sectors in 2007). Traffic flow management in today's NAS is typically performed by air traffic controllers (ATCs) using a complex set of ad-hoc rules, which results in a significant increase in ATC workload as the traffic demand increases, and which unfortunately does not provide performance guarantees. Consequently, there has been an increasing trend towards introducing a greater degree of automation and decision support for air traffic management, using optimization methods.

Most optimization-based methods proposed for air traffic management have traditionally focused on open-loop policies and periodic reoptimization [1], [2]. In these methods, given a fixed time horizon and based on available information, a schedule is decided for each aircraft traveling through the system. By solving large-scale integer programs, the optimal solution is determined that specifies the position of every aircraft at each instant. Besides the inherent complexity
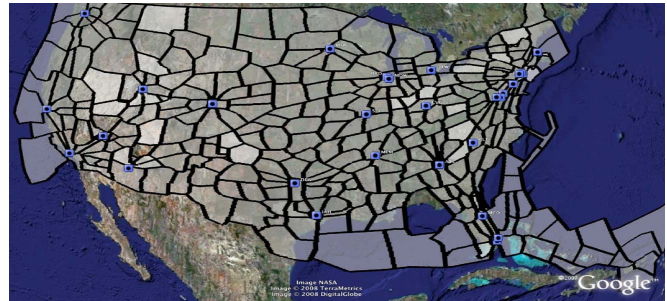
J. Le Ny is with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA jeromel@seas.upenn.edu. H. Balakrishnan is with the Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139-4307, USA hamsa@mit.edu.

Fig. 1. Sector boundaries for enroute sectors in the continental US. The markers denote the top 30 airports.

of scheduling approximately $40,000$ flights a day in this manner, controlling such as system using open-loop policies has potential disadvantages. Weather in particular is a major source of disturbances. Due to the typical travel times of cross-country flights, such traditional traffic flow management techniques require planning horizons of 5-6 hours, which are arguably beyond the limits of even state-of-the-art weather forecasting tools. Moreover, using these centralized algorithms at a more tactical level (say 5-20 min or perhaps to an hour ahead of operations) would require a complete transformation of the current air traffic control system, where local control authority is distributed between air traffic facilities such as airports and enroute sectors (which each have an ATC responsible for sector traffic planning). As an aircraft passes from its origin airport to the departure terminal-area, and through a series of sectors to the arrival terminal-area to its destination airport, responsibility is transferred from one ATC to another through a series of handoffs as they cross consecutive sector boundaries (Fig. 1). Neighboring sector controllers communicate with each other to coordinate handoffs. ATCs control the rates at which aircraft enter their airspace by regulating handoffs into their sector from the neighboring sectors. The approach presented in this paper models these local interactions between neighboring facilities, and uses feedback control to provide decision support for ATCs, within the existing control architecture of the NAS [3].

The many advantages 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 that are not as detailed in specifying trajectories in both space and time for each aircraft, and

that are therefore more tractable for the purpose of control. These aggregate models, called *Eulerian models*, are gaining popularity [5], [6], [7], [8], [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], and in a decentralized setting for networks with a single origin and destination [10].

These Eulerian models suggest strong parallels between approaches to air traffic flow management and the research on the control of stochastic networks. A survey of control approaches for other complex networks, such as semiconductor manufacturing systems or the internet [11], shows that discrete formulations based on deterministic integer programs or stochastic controlled Markov chains have been considered intractable and too detailed for the purpose of controlling realistic networks. As a first approximation, the discrete effects are usually neglected and continuous traffic flows are considered instead, much like Eulerian models of the NAS. For stochastic networks, these models used for control purposes are also called fluid models. However, unlike many Eulerian models of the NAS that involve Partial Differential Equations [7], fluid models for stochastic networks yield Ordinary Differential Equations (ODEs), thereby simplifying the analysis and control synthesis tasks significantly. These fluid models aim at capturing the average behavior of the system and are usually sufficiently tractable to design a first control policy. Unmodeled components and variability in the dynamics must be accounted for by modifying this basic policy adequately, similar to a robust control approach.

In this paper, we propose a lumped parameter model of the NAS in the spirit of the two-dimensional Eulerian model of Menon et al. [6]. However, our model is adapted to using standard network control tools, whereas they focused on linear systems theory. Our main objective is to develop a model that reflects the control structure in the current system, where facility (airport or sector) level traffic planning is done locally through coordination with neighboring facilities. In particular, the constraints on the system state (each sector contains a positive number of aircraft) and allocation rates (e.g. departure and arrival rates at airports [12], or limits on the rate of handoffs between sectors) play a major role in network control, yet they are neglected in most previous approaches [5], [6], [7]. We also note that queuing network models of the NAS have been proposed before, but generally with a view towards performance prediction and analysis, and not control [13], [14], [15], [8].

A tractable control design procedure for the NAS will most likely involve a hierarchical design, which is also the structure currently adopted [16]. Using a MaxWeight policy [11], we propose distributed feedback control laws for traffic flow planning at different facilities. The resulting flow rates will provide guidance to ATCs for implementation at the tactical scale of less than 5 min (which places emphasis on conflict detection and maintaining adequate separation between aircraft), although an increased level of automation at this time-scale is also possible. For example, a control
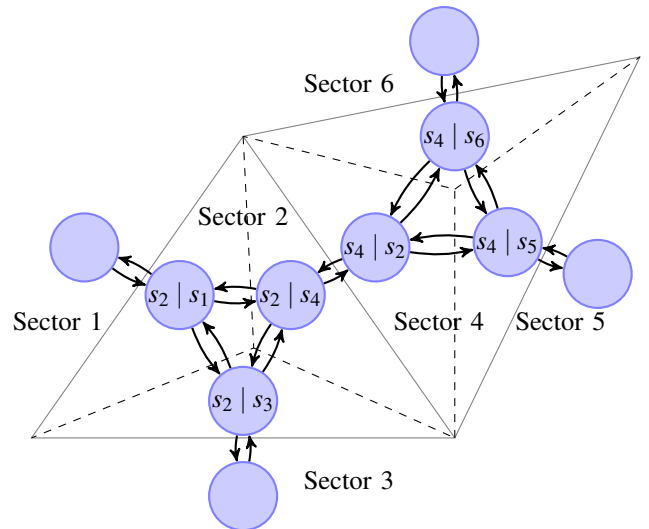

Fig. 2. Two triangular sectors of the NAS. Each node contains several queues, one for each path (section II-B) or each destination (section II-C). The links between nodes or queues are directed.

policy can suggest that an ATC decreases the flow rate of aircraft through a particular metering point. The actual choice of action (such as a speed change, vector for spacing, a holding pattern, etc.) required to achieve this objective is left to the discretion of the ATC. We believe that the flexibility of our model and its conformance to current inter-facility coordination procedures would allow it to be implemented on the existing system, instead of requiring a complete redesign of the system. The ability to handle uncertainty in our model is also important since it is not possible to account deterministically for all details and disturbances in a tractable model, which is necessary for control purposes.

The rest of the paper is organized as follows. In Section II, we present a new Eulerian model of the National Airspace System that reflects the way that flows in the current system are controlled by ATCs. We consider two variants of the model, one that addresses the scheduling of flights with no rerouting, and the second that considers the problem of simultaneous scheduling and routing. In Section III, we adapt well-known network control techniques to the models proposed in Section II. This allows us to develop control policies that can be easily used to provide decision support to ATCs, since they are consistent with the distributed manner in which flow rates into air traffic control facilities are controlled today, although in an ad hoc manner. In Section IV, we present extensions of these models, and in Section V, we present simulations of representative scenarios to illustrate the proposed approaches.

II. EULERIAN MODEL OF THE NAS

A. Graphical Model

We describe a two-dimensional lumped-parameter model of a sector of the NAS, as depicted on Fig. 2, which can be used to build a NAS-wide model for traffic flow management. Extension to a three-dimensional model is straightforward.

Before we proceed, we should emphasize that it is possible to easily modify the described framework and add other elements of interest. For example, we describe in section IV-A how to add traffic sources and sinks at airports, following the airport model of Gilbo [12]. The simulations in section V are performed using a simplified variation of the model, the main idea being to model the airspace as a network of queues that can be used for control.

We assume for simplicity on Fig. 2 that the two-dimensional model of the NAS has been triangulated into sectors. Menon et al. [6] have proposed a similar idea of controlling aircraft count in airspace "surface elements", in contrast to some other approaches based on one-dimensional models and focusing on jet routes and Victor airways. The particular shape of the sector is not important for our purposes and using the current geometry of the sectors is straightforward, as done in our simulation in section V. For the problem of model parameter identification, which we do not address in this work, but necessary to integrate with the tactical control level, it is useful to have sectors matching the existing geometry. We note that the complexity of the parameter identification problem and of the interactions between different traffic flows within a sector increases with the complexity of the sector geometry and the size of the sector. On the other hand, having numerous smaller sectors simplifies the identification problem and limits the number of coupling constraints between traffic flows, but increases the number of variables in the control problem. Therefore, a right balance must be found in the modeling process, a point which we do not address in this paper.

We model the NAS as a network. In the model shown on Fig. 2, there are 2 nodes per sector boundary, one in each sector. A node in a sector represents a portion of the airspace of that sector. Formally, if we start with the sector graph where edges are sector boundaries and vertices are the points where two or more of these sector boundaries join, then the network that we construct is the associated *line graph* [17], with every node doubled. This line graph has one node per edge of the sector graph and one edge between two nodes representing adjacent sector graph edges. To node i we can potentially associate a maximum capacity C_i, which is the maximum number of aircraft that can be present in the associated region of the airspace. The capacity of a node depends on the experience of the ATC, the weather, and can be time varying and evolve randomly. For tractability in the control design, and especially since capacity is often not a hard constraint, it might also be preferable not to impose a node capacity and instead to simply verify a posteriori that a given control law respects these capacities. Ensuring that airspace capacities are respected in critical regions can then be done in our framework by adding a higher cost per plane transiting through this region (via the matrix D in (8) below). The links between the nodes represent abstractly the airspace as a shared resource. To each node we associate a set of queues, each containing a subset of the aircraft currently in the corresponding region, and aircraft move between queues at different nodes as dictated by the flow rates on the links,

under air traffic control. We describe two models, without and with aircraft rerouting respectively.

B. A Model for Scheduling with No Routing

If the flight plans are predecided and cannot be changed by the ATCs, the remaining task is the scheduling of traffic via flow rate control. We denote by $\mathscr{P}$ the set of travel paths or flight plans. A given flight chooses one of these paths once and for all when filing its flight plan before departure. On the network model, the flight path corresponds to a path through a set of links. *For a given link (i,j), there is one buffer at its origin node i for each flight path that uses the link.* Let $m \in \mathscr{P}$ be a flight path that uses the link between nodes i^- and i followed by the link between nodes i and i^+. At node i, the content of the buffer for m varies according to the controlled random walk (CRW) on $\mathbb{N} := \mathbb{Z}_{\geq 0}$

$$Q_i^m(t+1) = Q_i^m(t) + A_i^m(t+1) \qquad (1)$$
$$+ M_{i^-i}^m(t+1; Q_{i^-}^m(t))U_{i^-i}^m(t) - M_{ii^+}^m(t+1; Q_i^m(t))U_{ii^+}^m(t),$$

where $Q_i^m(t)$ is the number of aircraft in region i on path m at time t, $M_{ij}^m(t;X) \in \mathbb{N}$ are random variables depending on the random variable X modeling the transitions of aircraft between regions and conditionally independent given X, $A_i^m(t) \in \mathbb{N}$ are i.i.d. random variables that are 0 except when i represents the departure region for path m and that model new aircraft taking off having filed path m in their flight plan. $U_{ij}^m(t) \in \{0,1\}$ are variables under the control of the ATCs, subject to certain constraints assumed to be linear and written

$$U(t) \in \{0,1\}, \ CU(t) \leq 1, \ \forall t \geq 1, \qquad (2)$$

where C is called the constituency matrix [11]. For instance $U_{ki}^m(t) = 1$ means that the ATC responsible of node k should let the aircraft present at that node with destination m progress toward node i. Note that implicit in (1) is an absorbing boundary condition at 0 which imposes that Q_i^m remains non-negative. Also we must have $U_{ij}^m(t) = 0$ if $Q_i^m(t) = 0$. More details on the control constraints are given in paragraph II-E.

C. A Model for Simultaneous Scheduling and Routing

It is possible to add routing control to the previous model. For this purpose, we now differentiate the aircraft only by their destination. We denote the destination by m, with $1 \leq m \leq l_f$ where l_f is the number of destination airports. At node i, there is now *one buffer for each possible destination* instead of one buffer for each path. The number of aircraft at node i with destination m evolves according to the CRW

$$Q_i^m(t+1) = Q_i^m(t) + A_i^m(t+1) + \sum_{k \in I(i)} M_{ki}^m(t+1; Q_k^m(t))U_{ki}^m(t)$$
$$- \sum_{j \in O(i)} M_{ij}^m(t+1; Q_i^m(t))U_{ij}^m(t), \qquad (3)$$

where $I(i)$ and $O(i)$ are the in- and out- neighbors for node i, respectively, and the variables A_i^m now model an external source of aircraft at node i, such as an airport, with destination m. Clearly, (3) is a generalization of (1). The

constraints on the control variables are again assumed to be of the form (2). We also impose $Q_i^m(t) = 0$ and $U_{ji}^m(t) = 0$ for all nodes $j \in I(i)$ if starting from node i there is no route to the desired destination m.

D. Fluid Model

For the purpose of analysis and control synthesis, the following so-called *fluid approximation* of the CRW (1) or (3) is useful. To fix ideas, we work with (3) in the following, which is the more general form. We associate to this CRW the constrained *deterministic* ODE

$$\frac{d^+}{dt}q_i^m(t) = \alpha_i^m + \sum_{k \in I(i)} \mu_{ki}^m(q_k^m)\zeta_{ki}^m - \sum_{j \in O(i)} \mu_{ij}^m(q_i^m)\zeta_{ij}^m, \quad (4)$$

where $\frac{d^+}{dt}$ denotes the right-derivative, and the number of aircraft $q_i^m(t) \in \mathbb{R}_+$ with destination m and present at node i at time t is now approximated by a continuous quantity. The control variable ζ_{ij}^m is called the rate allocation through the link (i, j) for path m and the vector of these rates is subject to the linear constraints

$$0 \leq \zeta(t) \leq 1, \ C\zeta(t) \leq 1, \forall t, \quad (5)$$

corresponding to (2). This approximating ODE can be postulated for our purposes, in the spirit of the previous Eulerian approximation of Menon et al. [5], and aims at capturing the mean behavior of the stochastic model (3). Note however that the connection can be made more formal, and an interpretation of (4) is in terms of a functional law of large numbers for a sequence of stochastic networks, see [18]. To connect (4) with (3) in simulations, we take α_i^m to be the mean of $A_i^m(t)$ and $\mu_{ij}^m(q_i^m)$ to be the mean of $M_{ij}^m(t; X)$ given $X = q_{ij}^m$.

E. Flow Rates and Constraints

To a link between nodes i and j and for destination m (or path m) is associated a maximum flow rate $\mu_{ij}^m(q_i^m)$ in (4), which depends on the number of aircraft q_i^m with destination m in the portion of airspace corresponding to node i, although other forms of dependencies could be used. Typically, if q_i^m is small, $\mu_{ij}^m(q_i^m)$ is approximately the inverse of the expected time for an aircraft to travel between the regions represented by nodes i and j, that is, the expected travel time in the absence of any delaying control issued by the ATC. As q_i^m increases, the rate μ_{ij}^m increases but remains bounded due to the required separation distance between aircraft. Assume for example that a single route is used by the aircraft traveling from i to j. Then μ_{ij}^m is bounded by the inverse of the time necessary for an aircraft to travel the separation distance between itself and the previous aircraft. For example, with a separation distance of 5 nm (nautical miles) and an aircraft speed of 450 knots, this gives a rate of $1/40$ aircraft/sec at high densities. Typically we expect the rate curve to have a shape similar to the one shown on Fig. 3. We aggregate the trajectories of all aircraft traveling through the region of the airspace represented by a particular link, and we can determine $\mu_{ij}^m(q_i^m)$ empirically based on historical data. The variations in traveling times among aircraft are treated as
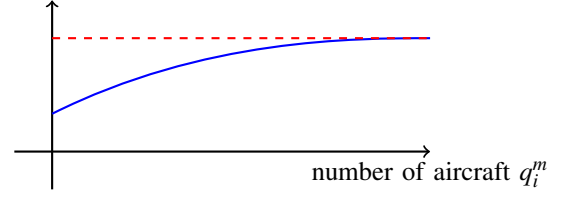

Fig. 3. Typical shape of the flow rate for a link between two nodes.

a stochastic quantity in the corresponding CRW model (3). Note that usually Eulerian models assume a particular rate curve μ_{ij}^m linear in q_i^m, see e.g. [5], [10]. This assumption is particularly unrealistic when a region contains a large number of aircraft, and we relax it here. Our model and analysis can handle any flow rate curve, as specified by the user. In practice, the maximum rates μ_{ij}^m of the links can also vary in time and stochastically to capture the effect of the weather and variations during the day, but this interesting aspect is not treated in this paper. For a boundary between two sectors the maximum flow rates μ_{ij}^m are also bounded by the maximum number of aircraft that are allowed to cross between boundaries per time unit (handoffs), as fixed by formal agreement between sectors [19].

ATCs can issue orders, such as aircraft speed change, vector for spacing, holding pattern, which modify the time it takes for an aircraft to travel a particular region of the airspace, that is, to travel from one node to the next in our model. In this way, they affect the flow rates of aircraft through the network. These orders correspond to controlling the rate allocation vector ζ, so that the actual flow rate through link (i, j) for destination m at time t is $\zeta_{ij}^m(t)\mu_{ij}^m(q_i^m)$. At the border between two sectors, the allocation rate also controls the rate of handoffs to the next sector. There are linear coupling constraints of the form (5) between the rate allocations. For example, the required separation between aircraft translates into a capacity constraint for link (i, j) of the form $\sum_m \zeta_{ij}^m \leq r_{ij}$. These constraints can also capture the effect of intersecting routes for example, and could be determined empirically, as in the case of airport capacities [12]. Indeed increasing the flow rate on one route is likely to require a decrease in the flow rate on intersecting routes. The same framework can be used to handle traffic flow management around airports, as explained in section IV-A. Finally, we see that having a finer grid partitioning the airspace, although increasing the number of variables in the model, allows for refined controls available to the ATCs. Hence it should be beneficial to divide each sector into smaller regions to further help the task of the controllers. Spatial aircraft separation places a lower bound on how small these regions can be made in practice.

F. General Formulation

In the previous sections we have seen that, after discretization of the airspace, we obtain a fluid network with dynamics

which can be written in matrix form as

$$\frac{d^+}{dt}q(t) = B(q)\zeta(t) + \alpha \tag{6}$$
$$q(t) \geq 0, \ \zeta(t) \in \mathsf{U}, \forall t, \quad \text{with } \mathsf{U} = \{u : u \geq 0, \ Cu \leq 1\},$$

where $B(q)$ is obtained from (4). Note that we are writing informally here a constrained differential equation, with q_i remaining 0 if we have $q_i = 0$ and $\dot{q}_i < 0$. We also denote by $\mathsf{U}(x)$ the state dependent control constraint set such that ζ respects the state constraint at $q(t) = x$, i.e. $\zeta_{ij}^m = 0$ if $q_i^m = 0$ (or if $\sum_m q_j^m = C_j$, if a capacity constraint is enforced at node j). Clearly $\mathsf{U}(x) \subset \mathsf{U}$. System (6) is then the fluid model associated to a discrete stochastic model of the form

$$Q(t+1) = B(t+1; Q(t))U(t) + A(t+1), \tag{7}$$
$$Q(t) \geq 0, U(t) \in \mathbb{N}, \ CU(t) \leq 1, \ \forall t.$$

Starting with the network of paragraph II-A, additional nodes can be added to model additional features without changing the general form of the equations. For example, to model transition through certain fixes in the airspace where traffic is metered, we just add two nodes and a link between them, with the capacity of the link determined by this metering. More examples are provided in section IV.

III. DISTRIBUTED CONTROL USING MAXWEIGHT

A number of control techniques for the network model (6) and (7) have been investigated in the past decades [11]. We remark that the state and control constraints, neglected in the original work of Menon et al. [5], in fact play a fundamental role in the analysis of a network's stability and performance. For example, simple scheduling problems typically correspond to a square invertible constant matrix B, which implies controllability of a version of model (6) neglecting the control constraint U, yet in practice not all such models are stabilizable. Their subsequent work used model predictive control to cope with these constraints [6], but result in centralized policies that are potentially hard to implement in the current distributed form of the NAS traffic flow management system. We now present a well-known and widely-used distributed control policy for the network models (6) and (7). For this section, we assume that the only allocation rate constraints are of the form

$$U := \{\zeta : \zeta \geq 0, \ \sum_m \zeta_{ij}^m \leq r_{ij}, \ \forall (i,j)\}.$$

A. MaxWeight for Routing and Scheduling

The celebrated MaxWeight policy for the network model (6) (or similarly for (7)) can defined as the policy which, for state q of the system (representing the size of all the queues at all nodes), applies the control

$$\phi^{MW}(q) \in \arg\min_{\zeta \in \mathsf{U}(q)} \langle B(q)\zeta + \alpha, Dq \rangle, \tag{8}$$

for D a fixed positive definite diagonal matrix. It corresponds to maximizing the rate of decrease of the Lyapunov function $V(q) = \frac{1}{2}q^T Dq$, since dV/dt is given by the right hand side of (8). If $D = I$, this is called the back-pressure policy [20].

The only difference with traditional network models is that the matrix B in our model is state-dependent.

Now recall that the matrix $B(q)$ is obtained from the model (4), or its equivalent in the case of no routing. Then (8) leads to

$$\arg\min_{\zeta} \sum_{(i,j)} \sum_m \zeta_{ij}^m \mu_{ij}^m(q_i^m)[D_{jj}^m q_j^m - D_{ii}^m q_i^m].$$

For *each link* (i,j), the maximal back-pressure is defined as $\Theta_{ij}(q) := \max_m \mu_{ij}^m(q_i^m)(D_{ii}^m q_i^m - D_{jj}^m q_j^m)$. Then we see that if $\Theta_{ij}(q) \geq 0$, the MaxWeight policy in state q for link (i,j) gives strict priority to buffers achieving the maximal back pressure:

$$\sum_{m=1}^{l_f} \{\zeta_{ij}^m : \mu_{ij}^m(q_i^m)[D_{ii}^m q_i^m - D_{jj}^m q_j^m] = \Theta_{ij}(x)\} = r_{ij}. \tag{9}$$

This policy is attractive because it has a simple form, and *it can be implemented at each node in the network based solely on buffer levels at the node and its neighbors*. In practice, this means that it can be implemented directly by the controllers who only need to talk to the controllers in the neighboring sectors, similar to the way traffic is managed in the current system. The MaxWeight policy (9) also defines directly a control law for the discrete model (3). This policy indicates to an air traffic controller to give priority in a particular region to aircraft going to a certain destination over other aircraft, based on the information available to the controller and his/her neighbors.

B. Stability of MaxWeight

The MaxWeight policy with state-independent matrix B is also known to be stabilizing for any network that is stabilizable [11]. The case of a state-dependent matrix $B(q)$ is somewhat more complicated and we only provide a sufficient stabilizability condition here. Recalling the discussion leading to Fig. 3, with few aircraft in the system, the flow rates $\mu_{ij}^m(q_i^m)$ are low and the number of aircraft in the system can grow. At high loads, the rates $\mu_{ij}^m(q_i^m)$ are higher, and for the network to be stable these rates should be sufficiently high to keep the queues bounded. Formally, we will use the following definition of stability

Definition 1: The fluid model (6) is said to be stabilizable if there exists constants K and T such that, for any initial condition q_0,

$$\|q(t)\|_\infty \leq K, \ \forall t \geq T\|q_0\|_\infty.$$

Proposition 2: Assume that there exists a positive constants C and ε such that for all q with $\|q\|_\infty > C$, there exists $\zeta \in U$ such that

$$\left\langle B(q)\zeta + \alpha, D\frac{q}{\|q\|_\infty} \right\rangle < -\varepsilon.$$

Then the network (6) is stabilizable and the MaxWeight policy (8) is stabilizing.

Proof: This follows from Lyapunov's direct method using the Lyapunov function $V(q) = \frac{1}{2}q^T Dq$. ∎
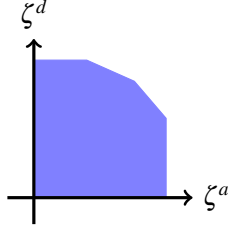
Fig. 4. Airport capacity region: any pair of arrival/departure allocation rates within the filled polytope is achievable. This polytope is determined empirically and depends on the airport runway configuration [12].

Fig. 5. Network Model of an Airport.

Remark 3: Under the same condition as in proposition 2, the MaxWeight policy is stabilizing for the CRW model (7) as well in a precise sense, by adapting [11, chap. 8].

Remark 4: If $B(q) := B$ is a state-independent matrix, the condition of proposition 2 (with $C = 0$) is also necessary for stability.

IV. ADDITIONAL EXAMPLES

A. Airport with Traffic Control at an Upstream Arrival Fix

It is undesirable to have a large number of aircraft waiting to land in holding patterns in close proximity of an airport. We can however control the rate of aircraft by setting a miles-in-trail (MIT) or minutes-in-trail (MinIT) requirement at certain arrival fixes [16]. The arrival and departure rates at the airport are subject to linear constraints as described by Gilbo [12], [21].

Consider the network model shown in Fig. 5. The equations for its fluid model are

$$\frac{d^+}{dt}q^f = \alpha^s - \mu^f \zeta^f, \quad \frac{d^+}{dt}q^a = \mu^f \zeta^f - \mu^a \zeta^a,$$
$$\frac{d^+}{dt}q^d = \alpha^d - \mu^d \zeta^d, \quad q(t) \geq 0,$$

where q^f, q^a and q^d are the number of aircraft upstream of the arrival metering fix, and in the arrival and departure queues of the airport.

Here μ^a, μ^d are constant rates, inverse of the average required separation time between two aircraft landing or taking-off at the airport respectively [22]. The maximum aircraft flow rate at the arrival fix μ^f is also assumed to be constant. The allocation rates are subject to the constraints $0 \leq \zeta \leq 1$ as well as additional linear constraints

$$C \begin{bmatrix} \zeta^a \\ \zeta^d \end{bmatrix} \leq 1,$$

as shown on Fig. 4. The goal is to minimize a cost function $c(q(t))$ over time (for all t if possible, or a time integral of this cost otherwise), for example

$$c(q) = c^f q^f + c^a q^a + c^d q^d.$$

If c^a is sufficiently large then an optimal control will keep the fluid level in the arrival queue q^a to 0 by reducing ζ^f appropriately. If this queue is maintained to 0, we can write
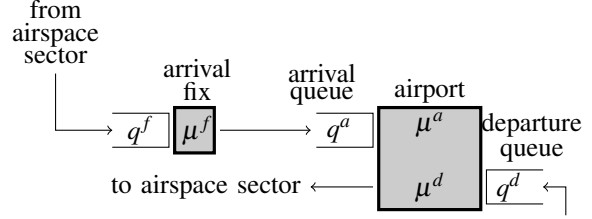
$$\zeta^a = \frac{\mu^f}{\mu^a}\zeta^f,$$

deduce the constraints linking ζ^f and ζ^d and simply optimize over these two controls. [11] presents a number of available control techniques for such a model.

B. Traffic Driven by Demand Instead of Arrival Rates

While the discussion so far has focused on systems in which the traffic flows are driven by the arrival of flights at originating airports that need to serve different destinations, there are potential scenarios in which we would like to schedule and route traffic flows that are driven by demand at the destination airports. This is most frequently seen on days in which severe weather has been forecast at one of the severely congested airports (for example, New York's La Guardia airport) in the system. In such scenarios, traffic flow management will need to coordinate with the airlines in order to best schedule the demand for flights into JFK, before the airport is closed because of weather. In other words, the demand for service into JFK will drive (*pull*) the traffic flows in the system, rather than the *push* model that we have considered in the previous discussions. A more detailed discussion of these two types of network models, and associated control techniques, can be found in [11].

V. SIMULATION

We consider the airspace surrounding San Francisco airport (SFO) in simulating representative scenarios using the approaches described in this paper. As a first step, we consider departures from Los Angeles (LAX) and other airports in southern California into SFO and Seattle (along with Portland and other airports in the Pacific northwest), aircraft flows from SEA to SFO, and flights from Las Vegas (LAS) to SFO. The airspace modeled is primarily within the Oakland Air Route Traffic Control Center (ZOA). The airspace sector boundaries and the traffic routes being considered are shown in Figure 6.

We simulate the behavior of the system under MaxWeight control, with and without routing. The simulations include capacity constraints on the number of aircraft allowed in each sector, which correspond to the maximum number of aircraft that can simultaneously be present in a sector. We model two routes between LAX and SFO, a short one and a long one that merges with other routes coming from LAS and the points East. Traffic on the short route between LAX and SFO competes with the traffic going from LAX to SEA. Fig. 7 shows, for a sample trajectory of the stochastic systems (1) and (3), the number of aircraft on each of the two routes LAX-SFO in the first sectors after leaving LAX. In both
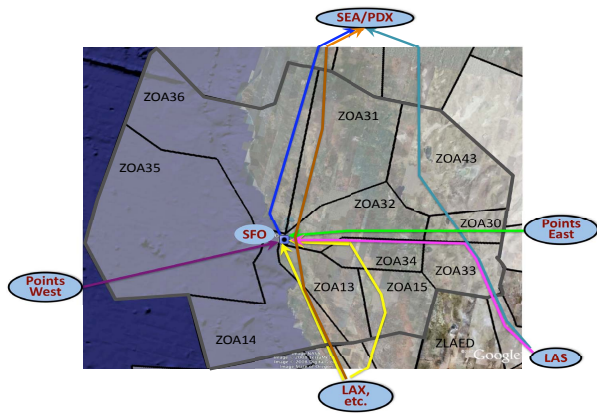
Fig. 6. Airspace region considered in the simulations, with routes. The routes are colored based on the origin-destination pairs.

(a) routing predetermined (see II-B)

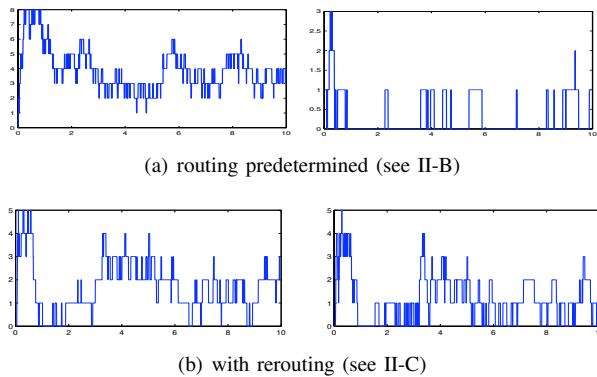

(b) with rerouting (see II-C)

Fig. 7. Aircraft count (y-axis) over a 10 hour interval (x-axis is time) in the south sectors of the LAX-SFO routes. Left: short route. Right: long route.

cases we use the back-pressure policy, i.e., $D = I$ in (9). We note in Fig. 7(b) that the the back-pressure policy with routing tends to balance the load on both routes evenly, which may not be optimal from the point of view of operating costs or passenger delay, if one route is significantly faster than the other. Fig. 7(a) on the other hand assumes a fixed predetermined fraction of aircraft on each route. The knowledge that one route is shorter and therefore preferable can be incorporated in the MaxWeight policy with routing by increasing the weight D_{ii} for the first buffer i on the longer route, thereby prioritizing routing on the shorter route (which has lower buffer weight).

VI. CONCLUSION

The increasing levels of congestion in the NAS motivate the development of decision support tools for air traffic controllers to aid them in managing traffic flows into and out of their facilities. The main contribution of this paper is the development of an Eulerian model of traffic flows in the NAS that reflects the current manner in which air traffic controllers regulate traffic flows and coordinate with their neighbors, and the subsequent development of a distributed feedback control law for the system. The proposed model can also provide insight into the behavior of the system as it responds to capacity and demand disruptions, and helps

in developing appropriate responses in the event of severe weather or airport closures.

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 and S. Stock Patterson, "The air traffic flow management problem with enroute capacities," *Operations Research*, vol. 46, no. 3, pp. 406–422, May-June 1998.

[3] R. W. Schwab, A. Haraldsdottir, and A. W. Warren, "A requirements-based CNS/ATM architecture," in *Proceedings of the SAE/AIAA World Aviation Congress*, 1998.

[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] 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.

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

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

[13] K. M. Malone, "Dynamic queueing systems: Behavior and approximations for individual queues and for networks," Ph.D. dissertation, Massachusetts Institute of Technology, 1995.

[14] D. Long, D. Lee, J. Johnson, E. Gaier, and P. Kostiuk, "Modeling air traffic management technologies with a queueing network model of the national airspace system," NASA, Tech. Rep. CR-1999-208988, 1999.

[15] B. G. Chandran, "Predicting airspace congestion using approximate queueing models," Master's thesis, University of Maryland at College Park, 2002.

[16] M. Nolan, *Fundamentals of Air Traffic Control*, 4th ed. Brooks Cole, 2003.

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

[18] A. Mandelbaum and G. Pats, "State-dependent stochastic networks. Part I: Approximations and applications with continuous diffusion limits," *The Annals of Applied Probability*, vol. 8, no. 2, pp. 569–646, 1998.

[19] Federal Aviation Administration, "New York Area Standard Operating Procedures and Letters of Agreement," Order 7110.1A, 2005.

[20] L. Tassiulas, "Adaptive back-pressure congestion control based on local information," *IEEE Transactions on Automatic Control*, vol. 40, no. 2, pp. 236–250, February 1995.

[21] E. P. Gilbo, "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.

[22] R. de Neufville and A. Odoni, *Airport Systems: Planning, Design and Management*. McGraw-Hill, 2003.

[23] 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.