

Adjoint-Based Control of a New Eulerian Network Model of Air Traffic Flow

Alexandre M. Bayen, *Member, IEEE*, Robin L. Raffard, and Claire J. Tomlin, *Member, IEEE*

Abstract—An Eulerian network model for air traffic flow in the National Airspace System is developed and used to design flow control schemes which could be used by Air Traffic Controllers to optimize traffic flow. The model relies on a modified version of the Lighthill–Whitham–Richards (LWR) partial differential equation (PDE), which contains a velocity control term inside the divergence operator. This PDE can be related to aircraft count, which is a key metric in air traffic control. An analytical solution to the LWR PDE is constructed for a benchmark problem, to assess the gridsize required to compute a numerical solution at a prescribed accuracy. The Jameson–Schmidt–Turkel (JST) scheme is selected among other numerical schemes to perform simulations, and evidence of numerical convergence is assessed against this analytical solution. Linear numerical schemes are discarded because of their poor performance. The model is validated against actual air traffic data (ETMS data), by showing that the Eulerian description enables good aircraft count predictions, provided a good choice of numerical parameters is made. This model is then embedded as the key constraint in an optimization problem, that of maximizing the throughput at a destination airport while maintaining aircraft density below a legal threshold in a set of sectors of the airspace. The optimization problem is solved by constructing the adjoint problem of the linearized network control problem, which provides an explicit formula for the gradient. Constraints are enforced using a logarithmic barrier. Simulations of actual air traffic data and control scenarios involving several airports between Chicago and the U.S. East Coast demonstrate the feasibility of the method.

Index Terms—Adjoint-based optimization, control of partial differential equations, LWR PDE.

I. INTRODUCTION

THE National Airspace System (NAS) consists of aircraft, control facilities, procedures, navigation and surveillance equipment, analysis equipment, decision support tools, and controller pilots who operate the systems. In this article, the focus is traffic flow management (TFM), which has the goal to maximize throughput while maintaining safety. This entails the design of efficient methods to route aircraft, while preventing the density of aircraft from becoming too large in regions of airspace, and operating efficient reroutes when the weather does not allow traffic to cross a given region of airspace. These tasks are not currently

Manuscript received May 11, 2005. Manuscript received in final form March 27, 2006. Recommended by Associate Editor S. Devasia. This work was supported in part by NASA under Grant NCC 2-5422, by the Office of Naval Research (ONR) under MURI Contract N00014-02-1-0720, by the Defense Advanced Research Projects Agency (DARPA) under the Software Enabled Control Program (AFRL Contract F33615-99-C-3014), and by a Graduate Fellowship of the Délégation Générale pour l'Armement, France.

The authors are with the Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305-4035 USA and are also with the Department of Electrical Engineering and Department of Civil and Environmental Engineering, University of California at Berkeley, Berkeley, CA 94720-1710 USA.

Digital Object Identifier 10.1109/TCST.2006.876904

optimized with respect to throughput or maximal density. Rather, they are prescribed by playbooks, which are procedures that have been established over time, based on controller experience.

The key objective of this article is to design control strategies in the form of “flow patterns,” that is, to come up with ways to route streams of aircraft by generating the corresponding aircraft velocities, rather than optimizing local trajectories of aircraft. Ideally, one would like to automatically generate procedures implementable by air traffic control (ATC), of the following kind: “aircraft on airway 148 at 33 000 ft, fly at 450 kn for the next hour and then accelerate by 10 kn per half hour.” This suggests following an Eulerian approach advocated by Menon *et al.* [24] and dividing the airspace into line elements corresponding to portions of airways, on which the density of aircraft as a function of time and of the coordinate along the line, is modeled. A traditional way to describe the evolution of the density of vehicles in a network is to use a partial differential equation (PDE). This PDE appears naturally in highway traffic and is called the Lighthill–Whitham–Richards (LWR) PDE [22], [30]. In this work, we will derive a modified version of the LWR PDE specifically applicable to the ATC problem of interest.

First, we show that despite the information loss inherent in any Eulerian model, the aircraft count (which is a crucial ATC metric defined in this article) is predicted accurately. Second, we show that fast numerical analysis tools can be applied efficiently to this problem for simulation purposes. The main difference between ours and previous work using LWR models of air traffic [24] or highway traffic [14], [26], [34], is that we generate an optimization problem (with throughput and maximal density as an objective function) using the continuous PDE directly, instead of its discretization. We show that the use of linear numerical schemes to approximate the solution of the PDE perform very poorly, which unfortunately precludes the use of standard linear optimization programs to control the system.

Controlling transportation networks in general is extremely challenging and numerically difficult [15], [26]. In the present case, the control consists of speed assignments and routing policies. We show that we may use flow control techniques [8], which are directly applicable to PDE-driven systems.

Namely, we pose the optimal control of the network as an optimization program, whose variables are solutions to a set of PDEs and satisfy additional inequality constraints. The optimization is performed by updating the control variables in the opposite direction of the gradient of the objective function. The gradient is derived using an adjoint method, specially adapted to the case in which the system is described by a set of PDEs coupled through the boundary conditions, in the presence of constraints. This algorithm does not provide proofs of convergence to a global optimal. However, this method, as well as other flow

control approaches [19], [3], [1], [13], [21], [32], have been shown to work extremely well in practice in fluid mechanics. In addition, though we consider networks of PDEs, the dimension of each PDE is one, enabling online implementations, as solving a set of one dimensional PDEs may be done extremely quickly. As such, we demonstrate the feasibility of generating direct, open-loop control solutions to the air traffic flow control problem using accurate numerical schemes.

There are a few benefits of the above outlined approach over Lagrangian methods, which incorporate all trajectories of all aircraft.

- Most of the Lagrangian methods pose the control problem as an integer optimization program, which is intractable in real-time because it is NP complete. In addition, the solution provided by these methods often takes advantage of actuating single aircraft individually, which precludes the derivation of global policies. The Eulerian framework scales well with the number of aircraft (the larger the number of aircraft is, the more accurate the model becomes, without further computational complexity).
- The method presented in this article is general and can be very easily adapted to specific classes of controllers (smooth, continuous, piecewise affine, etc.): it is possible to use this method to derive a control law in a required format, which is compatible with aircraft capabilities.
- This method can be applied to highway traffic with minor modifications [6], and we believe can be extended to other problems such as networks of irrigation channels [23].

This article is organized as follows. In Section II, we will first rederive the LWR PDE for the case of interest and generalize it to a network. Then, we determine an analytical solution for the case of time-invariant velocity control, which, in Section II-C, will be used for numerical validation purposes. In Section III, we explain how to identify the numerical values of the parameters for the airspace of interest using enhanced traffic management system (ETMS) data. This model is validated against ETMS data in Section IV. In Section V, we derive the adjoint system to our problem, and show how to use it to determine the mean velocity profiles along the links as well as the routing policy. Finally, in Section VI, we show how to apply this to control a very busy portion of airspace: the area enclosed by Chicago, New York, Boston, and the eastern coast of Canada.

II. NEW EULERIAN NETWORK MODEL OF AIRSPACE

A. Modified LWR Model of Air Traffic

In describing the air traffic system, like the road system, one has to first look at aircraft (or cars) present in the system and estimate a density of vehicles. Therefore, given a portion of airspace (airway or sector), one needs to introduce the aircraft count [9] defined as the number of aircraft in that region. Let us consider a portion airway of length L , described by a coordinate $x \in [0, L]$. The number of aircraft in the segment $[0, x]$ at time t is called $n(x, t)$. Thus, $n(L, t)$ represents the aircraft count on the portion of airway $[0, L]$. Assuming a static mean velocity profile v defined on $[0, L]$, $v(x) > 0$ represents the mean velocity of aircraft at location x , and the motion of an aircraft is described by the dynamical system $\dot{x} = v(x)$.

Introducing $K(x) = \int_0^x ds/v(s)$, it is fairly easy to see that if an aircraft were at location x_0 at time t_0 , it would be at x at time $t = t_0 + K(x) - K(x_0)$. Because of the sign of v , K is invertible, and, therefore, x_0 is related to x, t and t_0 by $x_0 = K^{-1}(K(x) - (t - t_0))$.

Consider a point x and a distance $h > 0$ such that $x < x + h < L$. The number of aircraft between x and $x + h$ at t can be related to the number of aircraft at t_0 at locations $x_0 = K^{-1}(K(x) - (t - t_0))$ and $x_h = K^{-1}(K(x+h) - (t - t_0))$ (conservation of aircraft): $n(x + h, t) - n(x, t) = n(K^{-1}(K(x+h) - (t - t_0)), t_0) - n(K^{-1}(K(x) - (t - t_0)), t_0)$. In other words, assuming that there is no inflow at 0

$$n(x, t) = n(K^{-1}(K(x) - (t - t_0)), t_0).$$

Some simple algebra (two successive applications of the chain rule) shows that the space derivative and the time derivative of n are related by

$$\frac{\partial n(x, t)}{\partial t} + v(x) \frac{\partial n(x, t)}{\partial x} = 0.$$

We recognize this as a first-order linear hyperbolic PDE, and can now enunciate the following proposition.

Proposition 1: Let $v(\cdot) : [0, L] \rightarrow \mathbb{R}^+$ be a $PC_1([0, L])$ function with a finite number of discontinuities at $\{x_k\}_{k \in \{0, \dots, K\}}$ on $[0, L]$. Assume $\exists m > 0, m \leq v(x)$ for all $x \in [0, L]$. Let $q^{\text{in}} \in C_0, PC_1([0, T])$ and $n_0 \in C_0, PC_1([0, L])$. Then the following PDE:

$$\begin{cases} \frac{\partial n(x, t)}{\partial t} + v(x) \frac{\partial n(x, t)}{\partial x} = q^{\text{in}}(t), & \text{in } [0, L] \times (0, T] \\ n(x, 0) = n_0(x), & \text{in } [0, L] \times \{0\} \\ n(0, t) = 0, & \text{in } \{0\} \times (0, T] \end{cases} \quad (1)$$

admits a unique continuous (weak) solution, given by

$$\begin{aligned} n(x, t) &= n_0(K^{-1}(K(x) - t)) + \int_0^t q^{\text{in}}(u) du, \\ &\text{if } t \leq \int_0^x \frac{du}{v(u)} \\ n(x, t) &= \int_{t-K(x)+K(0)}^t q^{\text{in}}(u) du, \\ &\text{if } t \geq \int_0^x \frac{du}{v(u)} \end{aligned} \quad (2)$$

where $K(x) = \int_0^x du/v(u)$, and K^{-1} is its inverse.

Proof: See the Appendix.

In (1), q^{in} represents the inflow at the entrance of the link (i.e., at $x = 0$). In highway traffic flow analysis, n is sometimes referred to as cumulative flow. It can be related to the vehicle density through the integral relation

$$n(x, t) = \int_0^x \rho(u, t) du \quad (3)$$

where $\rho(x, t)$ is the vehicle density. It can be checked that the vehicle density satisfies the following PDE:

$$\begin{cases} \frac{\partial \rho(x, t)}{\partial t} + \frac{\partial}{\partial x}(\rho(x, t)v(x)) = 0 \\ \rho(0, t)v(0) = q^{\text{in}}(t) \\ \rho(x, 0) = \rho_0(x). \end{cases} \quad (4)$$

Equation (4) can be related to (1) by a simple integration of ρ along $[0, x]$. Equation (4) is a mass conservation equation. Note that it is very different from the original LWR PDE [22], [30], [2], [15], [12], which is a first order hyperbolic conservation law. In particular, (4) does not have a fundamental diagram, i.e., there is no functional relation between ρ and v , or between ρ and the flux. In the implementation studied in this paper, the function $v(x)$ will represent the control input. It is also possible to rewrite the first equation in (4) as

$$\frac{\partial \rho(x, t)v(x)}{\partial t} + v(x)\frac{\partial}{\partial x}(\rho(x, t)v(x)) = 0 \quad (5)$$

which provides the following corollary.

Corollary 2: The solution of (5) for ρ is given by

$$\rho(x, t) = \begin{cases} \rho(K^{-1}(K(x) - t), 0) \frac{v(K^{-1}(K(x) - t))}{v(x)}, & \text{if } t \leq K(x) \\ \frac{q^{\text{in}}(t - K(x))}{v(x)}, & \text{otherwise.} \end{cases}$$

The interpretation of the corollary follows. The quantity ρv is conserved along the characteristic curves $t - t_0 = K(x) - K(x_0)$. At this stage, ρ is defined by $\rho = \partial n / \partial x$ and satisfies (4). However, unlike for highway traffic, the density ρ might not be the best way to characterize the flow situation at a given time. If the number of aircraft in the system is small, ρ will be a set of spikes, which is intractable numerically. Therefore, a more tractable quantity to work with would be $\delta n / \delta x$, where δn represents the number of aircraft contained in a finite interval of length δx . This quantity does not *a priori* satisfy the PDE (4). It is meaningful to introduce an additional “density-like” quantity called r , which satisfies the PDE and for which we can suggest a physical interpretation

$$r(x, t) = \frac{[n(K^{-1}(K(x) + t_{\text{ref}}), t) - n(K^{-1}(K(x) - t_{\text{ref}}), t)]}{2t_{\text{ref}}v(x)}$$

where t_{ref} is a reference time. $r(x, t)v(x)$ represents the number of aircraft included in a window of $2t_{\text{ref}}$ time units around location x and can be referred to as “time density.” This way of accounting for density is meaningful for air traffic control, since it incorporates a time scale t_{ref} into the density computation and thus, provides access to the time separation between aircraft. It is easy to show that r itself satisfies the same PDE as ρ for any value of t_{ref}

$$\frac{\partial r(x, t)}{\partial t} + \frac{\partial (r(x, t)v(x))}{\partial x} = 0.$$

One can also show that when $t_{\text{ref}} \rightarrow 0$, r , and ρ are the same

$$\begin{aligned} \lim_{t_{\text{ref}} \rightarrow 0} \left[\frac{n(K^{-1}(K(x) + t_{\text{ref}}), t) - n(K^{-1}(K(x) - t_{\text{ref}}), t)}{2t_{\text{ref}}v(x)} \right] \\ = \frac{1}{v(x)}v(x)\frac{\partial n(x, t)}{\partial x} \\ = \rho(x). \end{aligned}$$

At this stage, we have three quantities: ρ , $\delta n / \delta x$, and r . The meaning of ρ as we know it in fluid mechanics assumes a large number of particles (i.e., aircraft) per unit volume (the threshold is defined by the Knüdsen number). In the present case, the number of aircraft we consider will almost certainly be below this number, meaning that the fluid approximation is questionable. This means that instead of using $\rho = \partial n / \partial x$, we will use $\rho \sim \delta n / \delta x$ in the PDE. We will justify this approximation with appropriate validations.

B. Network Model

The model of the previous section describes traffic on a single portion of airway or line element. As was done earlier for highways [15], this model can be generalized to airway networks, i.e., sets of interconnected airways, as shown in Fig. 1. We now derive a framework to describe unidirectional air traffic. We describe the topology of the network by a unidirectional graph (E, V) , in which E is the set of edges or links, and V the set of vertices. We index the links by $i \in \{1, \dots, N\}$, rather than by the indices of the two corresponding vertices. For all $i \in \{1, \dots, N\}$, we call $\mathcal{U}(i)$ the set of upstream links merging into link i , and \mathcal{M} the set of links for which the upstream links are only merging. The number of links merging into a single link is not limited; it is possible to have $|\mathcal{U}(i)| > 2$. If there is a divergence at the end of a link i , we assume for simplicity that there are only two emanating links from the corresponding vertex. We index by i_l and i_r the two emanating links (left and right), and call β_i the portion of the flow going from i to i_l , and $1 - \beta_i$ the proportion of the flow going from i to i_r . We call \mathcal{D} the set of links with a divergence at the end of it. The β_i are not known *a priori* and have to be determined. These coefficients might depend on t as well and, therefore, a dependence $\beta_i(t)$ is included in the model. We call \mathcal{S} the set of sources in the network and \mathcal{T} a sink of the network, at which we might want to perform optimization. We index all variables of the previous section by i : the aircraft density on link i is ρ_i , the coordinate is x_i , the main velocity profile is v_i , etc. Note that we are not using Einstein’s notation; the notation is summarized in Table I. The governing PDE system thus reads

$$\begin{cases} \mathcal{N}_i(\rho_i) := \frac{\partial \rho_i(x_i, t)}{\partial t} + \frac{\partial}{\partial x_i}(\rho_i(x_i, t)v_i(x_i)) = 0 & \forall i \in \{1, \dots, N\} \\ \rho_i(x_i, 0) = \rho_i^0(x_i) & \forall i \in \{1, \dots, N\} \\ \rho_i(0, t)v_i(0, t) = \sum_{j \in \mathcal{U}(i)} \rho_j(L_j, t)v_j(L_j, t) + q_i^{\text{in}}(t) & \forall i \in \mathcal{M} \\ \begin{cases} \rho_{i_l}(0, t)v_{i_l}(0, t) = \beta_i(t)\rho_i(L_i, t)v_i(L_i, t) \\ \rho_{i_r}(0, t)v_{i_r}(0, t) = (1 - \beta_i(t))\rho_i(L_i, t)v_i(L_i, t) \end{cases} & \forall i \in \mathcal{D} \\ \rho_i(0, t)v_i(0, t) = q_i^{\text{in}}(t) & \forall i \in \mathcal{S}. \end{cases} \quad (6)$$

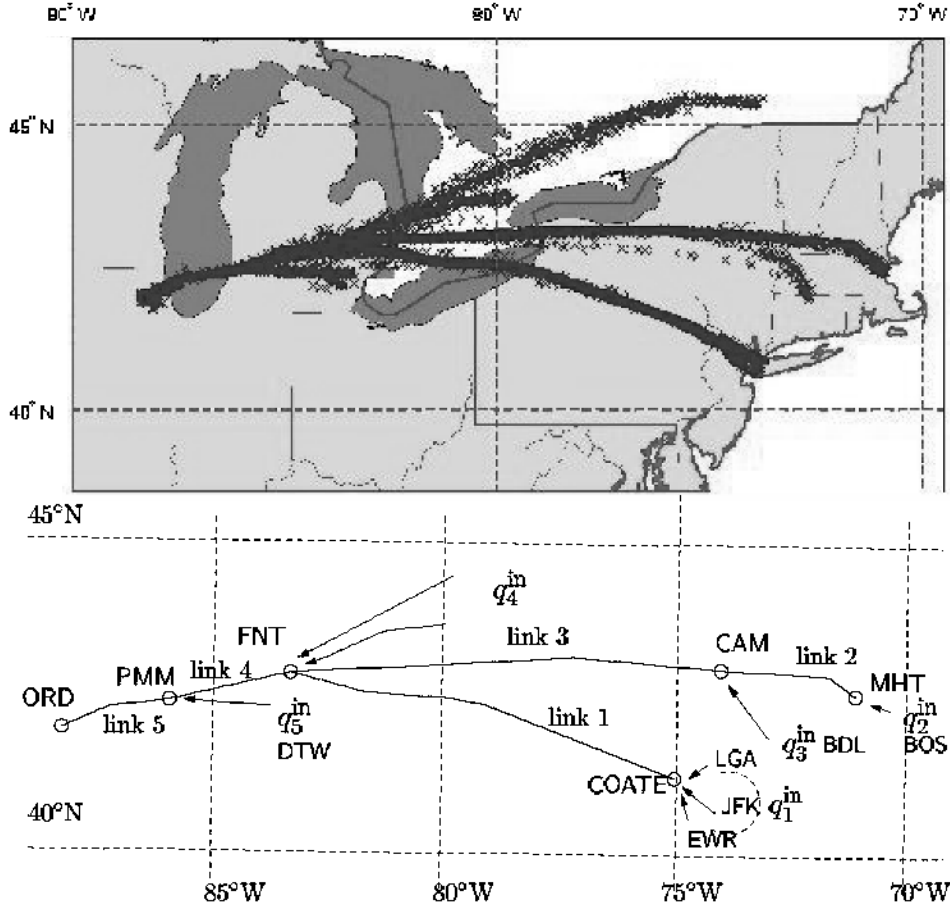


Fig. 1. Top: Tracks of flights incoming into Chicago (ORD). The upper stream comes from Canada, the lower from New York and Boston (BOS). Additional streams merge into the network (Detroit and Hartford Bradley). Bottom: Network model for the tracks shown above, with waypoints labeled. The model includes five links, merging into ORD. The corresponding inflow terms correspond to a single airport as in BOS or Detroit (DTW), or to a set of airports, as in New York (EWR, JFK, LGA).

TABLE I
NOTATION FOR THE NETWORK PROBLEM

N	number of links
S	set of source links
\mathcal{M}	set of links into which other links merge
\mathcal{D}	set of links ending in a fork
$\mathcal{U}(i)$	set of links merging into link i (if $i \in \mathcal{M}$)
i_l, i_r	indices of the two links of a fork if $i \in \mathcal{D}$
L_i	length of link i
x_i	arclength on link i : $x_i \in [0, L_i]$
$\rho_i(x_i, t)$	aircraft density on link i
$\rho_i^0(x_i)$	initial aircraft density on link i
$v_i(x_i)$	velocity profile on link i : $v_i(\cdot) : [0, L_i] \rightarrow \mathbb{R}^+$
$q_i^m(t)$	inflow at $x_i = 0$ for link i (if applicable)
$\beta_i(t)$	portion of ρ_i which flows into link i_l

In the previous system, the PDE (first equation) describes the evolution of ρ_i on each link. The notation $\mathcal{N}_i(\cdot)$ represents the LWR operator. The second equation is the initial condition (i.e., the initial density of aircraft on each link). The third equation expresses the conservation of aircraft at the merging points. The fourth and fifth equations express the conservation of aircraft at the divergence points. The last equation expresses the boundary conditions (inflow at the sources of the network). The sinks of the system are free boundary conditions and, therefore, do not appear in the previous system. The solution of (6) enables the computation of certain metrics useful for ATC. For example, one quantity of interest is the aircraft count per sector.

C. Accuracy of Numerical Solutions

Even for a single link i it is, in general, not possible to solve the system (6) analytically when v depends on time. The solutions ρ of the LWR PDE in the system (6) have very undesirable properties for numerical integration: they are by construction discontinuous;¹ they can develop kinks if the velocity profiles are discontinuous. Several numerical schemes of the original LWR PDE have been the focus of recent research [16] in order to address similar difficulties encountered in the original LWR PDE; they have proved extremely efficient in the case of highway traffic. We have chosen to use three different schemes to compare their respective benefits.

- 1) The well-known Lax–Friedrichs scheme [17].
- 2) A left-centered scheme, inspired by the Daganzo scheme [16] in light traffic

$$\rho_k^{n+1} = \left(\rho_k^n \left(1 - \frac{v(x_k)\Delta T}{\Delta x} \right) + \rho_{k-1}^n \left(\frac{v(x_k)\Delta T}{\Delta x} \right) \right) \times \left(1 - \frac{v(x_k)\Delta T}{\Delta x} + \frac{v(x_{k-1})\Delta T}{\Delta x} \right).$$

- 3) The Jameson–Schmidt–Turkel (JST) scheme. This scheme is nonlinear and has very desirable properties for this work: it captures shocks (which are present in the solutions we

¹Unlike n which is its primitive and is continuous.

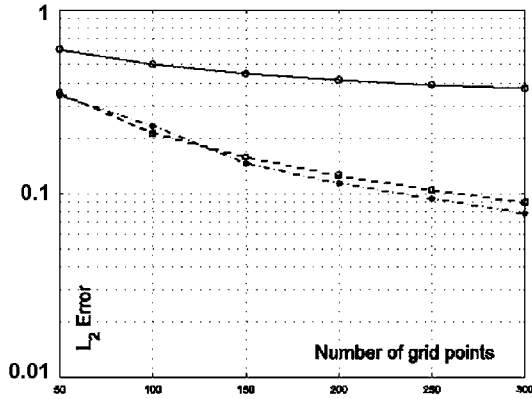


Fig. 2. L_2 error due to the discretization method, as a function of the number of grid points for both schemes. Lax-Friedrichs scheme (solid), Jameson-Schmidt-Turkel scheme (—), left-centered scheme (- -).

compute, as will be seen), and when the PDE has an entropy solution, which is the case for highway traffic in the original LWR setting, it converges to the entropy solution of the problem. Details of this scheme are available in [20].

Even if a numerical scheme is theoretically proven to converge to the analytical solution of a PDE, one usually does not know *a priori* the required gridsize to guarantee that the numerical solution is close to the analytical solution. This type of validation is standard in numerical analysis [17], [16].

We use the method developed earlier to compute the analytical solution of three benchmark problems solvable by hand, involving solutions with shocks and kinks (a detailed description of the benchmark examples is available in [4]). For each of the numerical schemes used, we compute the L_2 error due to the discretization method, as a function of the number of grid points. The result is shown in Fig. 2. This study leads to several conclusions. The Lax-Friedrichs scheme is very diffusive. Its behavior is representative of linear schemes to approximate a hyperbolic PDE. Consequently, we do not think that it is a good idea to use such linear numerical schemes, even if it would have the advantage of making the constraints linear in the resulting optimization program. The left centered scheme is less diffusive, but fails to capture the kinks of the solution. However, it still provides good L_2 convergence. The JST scheme captures shocks accurately because of its anti-diffusive term, and thus, gives the best results overall. It will be used for the rest of this study. Additionally, the JST scheme has the benefit that we can use it both for the direct problem, and for the adjoint. A detailed study of the computational time required to solve this class of problems is out of the scope of this study. For this, we refer the reader to our ongoing work [33], in which we compare the following three models: the original Menon model [24], the present model, and a new cell-based model [31].

III. SELECTION OF MODEL PARAMETERS

A. System Identification: Main Velocity Profiles

In this section, we identify the mean velocity profiles $v_i(x_i)$ on each link. We use enhanced traffic management system (ETMS) data, which we can obtain from NASA Ames (see [9] for a description of ETMS data). From ETMS data, we can obtain useful

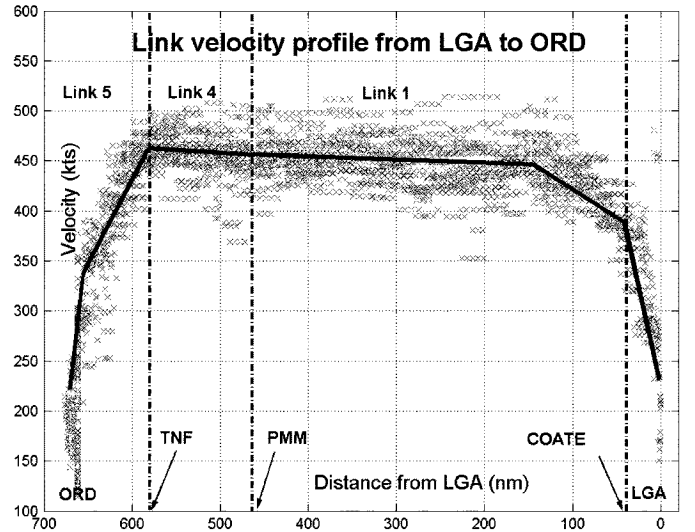


Fig. 3. Example of velocity profile used for the junction LGA-ORD. The horizontal coordinate is the distance from ORD in nm. The corresponding links are shown, as well as the location of the airspace fixes between the links.

flight information at a 3 min rate:² position of each aircraft in the NAS, altitude, velocity, and flight plan (i.e., set of airways and waypoints). From this data, we are able to identify the routes in which traffic is concentrated. Note that in recent work, Menon *et al.* [24] focused on creating an automated tool which performs similar tasks automatically at a NAS-wide level, using FACET [9], a tool developed by NASA Ames.

We analyzed 24 hours of ETMS data and selected all aircraft using the links of the network shown in Fig. 1. We identified all aircraft which used each of the links, and recorded all tracks and corresponding speeds between takeoff and landing. For each of the links shown in Fig. 1, we identified the mean velocity profiles as piecewise affine functions, using a least squares fit. The total number of aircraft used is 220. The result for the flight New York-Chicago is displayed in Fig. 3. The curve is a piecewise affine fit obtained using least squares. As can be seen, once the En Route altitude is reached, the curve fits are almost flat, which means that the aircraft are En Route at a high altitude cruise speed. It can also be seen from Fig. 3 that the data is relatively broadly spread (standard deviation 19.6 kn). This suggests a refinement using multilayer models: dividing the link in sublayers corresponding to altitudes (with different speed profiles) has the benefit of being more precise. In this work, we consider a single layer.

B. Initial and Boundary Conditions

Once the mean velocity profiles are computed, we identify the initial density of aircraft and the inflow (boundary conditions) in the network. The initial position of the aircraft is easy to extract from the ETMS data: at the prescribed time, all airborne aircraft which are on the relevant links are selected.

- 1) For any selected aircraft a , at location x_i^a , on link i , the classical density $\rho_i(x_i, 0) = \rho_i^a(x_i)$ is taken to be a “box” around x_i^a , of length $2L_{\text{ref}}$, where L_{ref} is a reference length

²Current ETMS data can now be obtained at a higher rate, which was not available at the time this work was performed.




Fig. 4. Different predictions obtained by the use of ρ and r for aircraft density. Above: density propagation through the PDE system (6); below: position update from ETMS data and from the PDE.

relevant for the scale of the problem. Calling χ_I the characteristic function of an interval I (equal to 0 outside of I and 1 inside), $\rho_i(x_i, 0)$ is $(1/2L_{\text{ref}})\chi_{[x_i^a - L_{\text{ref}}, x_i^a + L_{\text{ref}}]}(x_i)$. Taking all aircraft initially airborne on link i , the density is

$$\rho_i^{\circ}(x_i) = \sum_{a \text{ in link } i} \frac{\chi_{[x_i^a - L_{\text{ref}}, x_i^a + L_{\text{ref}}]}(x_i)}{2L_{\text{ref}}}.$$

2) Similarly, the density-like function r is computed using the knowledge of the mean velocity profile along link i , called $v_i(x_i)$, and the parameter t_{ref}

$$r_i^{\circ}(x_i) = \sum_{a \text{ in link } i} \frac{\chi_{[x_i^a - t_{\text{ref}}v_i(x_i^a), x_i^a + t_{\text{ref}}v_i(x_i^a)]}(x_i)}{2t_{\text{ref}}v_i(x_i^a)}.$$

These two equations thus, represent the initial conditions for the density and the density-like functions, which we extract from ETMS data.

The inflows (boundary conditions) can also be extracted from ETMS data: each time an aircraft takes off, it will appear in the ETMS data as soon as it is airborne. The ETMS data also shows the filed flight plan, which we select when it intends to use the links of interest to us. $q^{\text{in}}(t)$ is computed the following way. At any instant when the data shows a new aircraft on one of the source links \mathcal{S} , the track is in general passed the entrance point of that link (because of the sampling rate of 3 min). Calling x_i^a the position of this aircraft on link i at the first time it appears, we compute the time t_a at which it crossed the location $x_i = 0$ (using the knowledge of the mean velocity profile on the link). We then use one of the two definitions above to compute $q^{\text{in}}(t)$ corresponding to either ρ or r .

C. Identifying the Numerical Parameters

As explained in Section III-C, we have two ways of describing the density of aircraft in the network, in terms of a density function ρ and a “density-like” function r , which, respectively, account for spatial and temporal distribution of aircraft. The function ρ depends on the numerical parameter L_{ref} , which we need to adjust. The value of this parameter is crucial for predicting aircraft count: Fig. 4 shows how errors can occur in translating density functions into aircraft count. We want to determine the choice of parameters leading to the smallest error in aircraft count prediction.

We first run the following set of experiments. For the link New York–Chicago, we run a set of simulations involving N_{aircraft} aircraft, where N_{aircraft} successively takes all values between 1 and 50. We vary L_{ref} between 0 and 120 nm. For each value

of N_{aircraft} and L_{ref} , we run 400 experiments. Each experiment corresponds to a uniformly distributed random density of N_{aircraft} aircraft along link 1 in $[0, 400]$ (see Fig. 1). The simulation starts at a time t_0 , with the density ρ_i° computed as in the previous section, and computes the solution of the LWR PDE until the time $t_0 + \Delta T$. For the experiments, ΔT was chosen equal to 1 hr (note that the duration of the total flight is on the order of two and a half hours). This solution is compared with the solution obtained by propagating the Lagrangian trajectories of each of the aircraft independently from t_0 to $t_0 + \Delta T$ and computing the resulting density. In mathematical terms, we compare the two following quantities:

- $\rho_i(\cdot, t_0 + \Delta T)$ computed by the LWR PDE [6];
- $\tilde{\rho}_i(\cdot, t_0 + \Delta T) := \sum_a \text{in link } i (\chi_{[x_i^a(t_0 + \Delta T) - L_{\text{ref}}, x_i^a(t_0 + \Delta T) + L_{\text{ref}}]}(\cdot)) / 2L_{\text{ref}}$, where $x_i^a(t_0 + \Delta T)$ is the position of aircraft a at time $t_0 + \Delta T$.

In order to characterize the best choice of numerical parameters, we compute the following two quantities (notations refer to Fig. 1):

- relative density error, defined by

$$\frac{\sum_{i=1,4,5} \int_0^{L_i} |\rho_i(x_i, t_0 + \Delta T) - \tilde{\rho}_i(x_i, t_0 + \Delta T)| dx_i}{\sum_{i=1,4,5} \int_0^{L_i} \rho_i(x_i, t_0 + \Delta T) dx_i}.$$

This quantity represents the error in density prediction due to the propagation of ρ by the PDE;

- absolute aircraft count error, defined by

$$\sum_{i=1,4,5} \sum_{\text{sublinks of } i} \left| \int_0^{L_i} \rho_i(x_i, t_0 + \Delta T) dx_i - \#(a | a \in \text{link } i) \right|$$

where $\#$ means number. This quantity is the sum of count error for all sublinks of links 1, 4, and 5. Typically, a link is divided into sublinks which correspond to different airspace sectors. For example, if link 1 goes through 8 sectors, we divide it in 8 sublinks and are interested in the aircraft counts on these sublinks. This error thus estimates the difference between the number of aircraft predicted by the PDE and the number obtained by a Lagrangian propagation of aircraft, where the error is the sum of all errors on the sublinks.

The computation of both quantities is illustrated in Fig. 5. In this figure, for each of these sublinks, we compare the number of aircraft predicted by the method (depicted by arrows, which are computed from the density) with the number of aircraft obtained by a Lagrangian propagation of the trajectories. The error is the sum of errors for all sublinks, i.e., the sum of the errors in sector counts. The relative density error and absolute aircraft count error are averaged (over the 400 runs) and plotted for the range of n and L_{ref} considered. The result is shown in Fig. 6. The left plot shows the relative density error. As expected, the error decreases when the number of aircraft increases and L_{ref} increases. The right plot shows the absolute aircraft count error, averaged over 400 simulations. For this plot, each of the links 1, 4, and 5 have been divided in sublinks (20 total), of about 50 nm length. This is a worst case scenario, i.e., the number of relevant sectors for a flight of this length is never higher. One can see

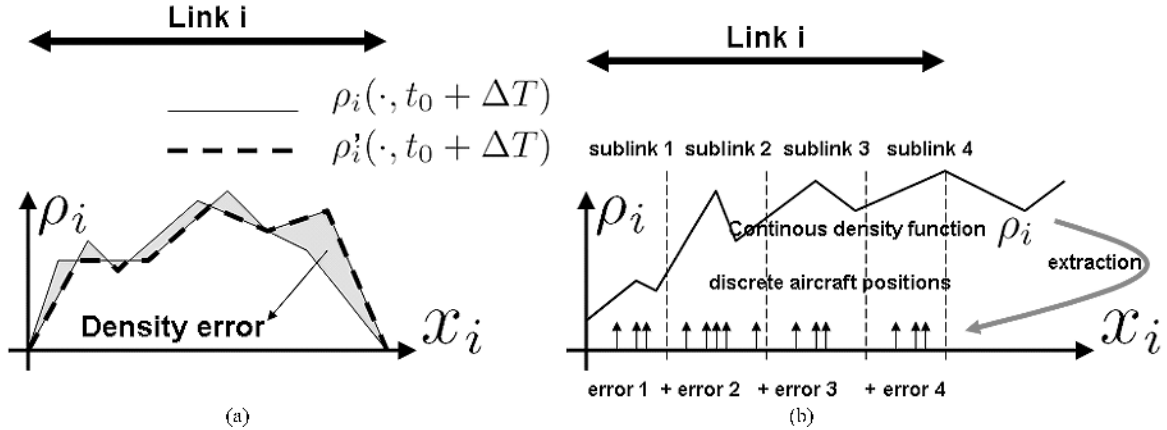


Fig. 5. (a) Illustration of the computation of the relative density error depicted in Fig. 6. The difference between the two density curves (shaded area) is divided by the area below the ρ_i curve. (b) Illustration of the computation of the error in aircraft count. The link is divided into sublinks (which can correspond to sectors).




Fig. 6. (a) Relative density error between the density predicted by the Eulerian PDE propagation of the density. (b) Absolute aircraft count error for the junction New York-Chicago.

that for $L_{\text{ref}} \leq 60$ and $N_{\text{aircraft}} \geq 25$, the average aircraft count error is always extremely small.

The best choice for L_{ref} is thus obtained at the intersection of the lowest level sets of both plots of Fig. 6, i.e., for a range of $L_{\text{ref}} \in [20, 60]$ and $N_{\text{aircraft}} \geq 20$. Fig. 6 can also be interpreted as follows. The region in the top-right corner and the bottom-left corner are both regions in which the model might not be applicable. As can be seen, the relative error or absolute count error exceeds values that might be realistic for practical purposes (15% error and absolute aircraft count error of 7). These regions are to be avoided.

IV. VALIDATION OF THE MODEL

In the previous section, we have shown that the use of the modified LWR PDE either with r (with any t_{ref}) or $\rho \sim \delta n / \delta x$ (with an appropriate choice of L_{ref}) enables accurate aircraft count predictions. In this section, we validate the model against real data.

A. Static Validation

In the first experiment, we use the static velocity profiles $v_i(\cdot)$ determined in the previous sections for the validation of

the method. We use a 6-hr ETMS data set. From this data set, we extract the position of the aircraft, at the initial time, construct the corresponding initial aircraft density, and propagate it through the PDE system. At any given time, we compare the aircraft count predicted by our method and the aircraft count provided by the ETMS data (which is exact, since it provides the position of each aircraft). We compute the error in aircraft count for a set of ten sublinks for the network shown in Fig. 1. The result is shown in Fig. 7(a). The window width L_{ref} was taken equal to 15 nm. One can see on the left plot that the total error (for all airborne aircraft in this airspace) is relatively low (the maximal error is 7 aircraft). In fact, the results are much better than they seem: most of the errors come from the fact that the aircraft distribution is such that there is always at least one or two aircraft close to a sublink boundary, which will thus be counted in the wrong sublink. In fact, this is not really a problem, as it is more an artifact of the computation rather than a true error (Fig. 11 illustrates that the density unambiguously shows where the aircraft is). Furthermore, some of the errors in aircraft count are due to errors present in the ETMS data (some have clearly erroneous data; this fact has also been reported in [11]).




Fig. 7. (a) Error in aircraft count for the static validation over a five hour period. (b) Error in aircraft count for the dynamic validation over a 5-hr period.

B. Dynamic Validation

We extend the validation to a case in which the velocity profiles are time dependent, i.e., $v_i(x_i, t)$. The details of the identification of these profiles are technical and are not explained here. The comparison is the same as in the static case. The results are shown in Fig. 7 and are more accurate than the static results, as expected. The same remarks apply, and the results are again affected by the quality of ETMS data and the inclusion of the computation artifact. The only weakness of this validation is that the simulation is run using data from the same day as the data used in identification. A way to improve this would be to perform the velocity identification with data of a given day over a 24-hr period, and validate it over the next 24-hr period, using the fact that there is periodicity in the traffic for normal days. This was not done here due to lack of available data. An animation (.avi movie file) corresponding to the snapshots of Fig. 7 is available.³

In both cases, the validation is very encouraging and shows strong predictive capability for our model. The model was also tested successfully using data from the western states (Oakland Center with traffic incoming into Bay Area airports), though for brevity these results are not included here.

V. NETWORK CONTROL VIA ADJOINT METHOD

Consider solving the following problem: maximize the throughput (i.e., flux of landing aircraft) at a destination airport, while maintaining the density of aircraft everywhere lower than a given threshold. Let us call $\rho_{\max,i}$ the maximal allowed density on link i , $v_{\max,i}(\cdot)$ and $v_{\min,i}(\cdot)$ the maximal and minimal achievable speeds on link i (which can depend on location). Using the notations of Section II-B, the optimization problem thus reads

$$\begin{aligned}
 \min : & - \sum_{i:i \in \mathcal{F}} \int_0^T \rho_i(L_i, t) v_i(L_i, t) dt \\
 \text{s.t.} : & (6) \\
 & \rho_i(x_i, t) \leq \rho_{\max,i} \\
 & \forall i \in \{1, \dots, N\}, \forall x_i \in [0, L_i], \forall t \in [0, T] \\
 & v_{\min,i}(x_i) \leq v_i(x_i) \leq v_{\max,i}(x_i) \\
 & \forall i \in \{1, \dots, N\}, \forall x_i \in [0, L_i], \forall t \in [0, T] \\
 & 0 \leq \beta_i(t) \leq 1, \quad \forall i \in \mathcal{D}, \forall t \in [0, T]. \quad (7)
 \end{aligned}$$

The difficulty posed by the constraints can be avoided in practice by using a barrier function as commonly done in optimiza-

tion [10], in which the cost is augmented by a logarithmic term, which prohibits violation of the constraints.

$$\begin{aligned}
 \min : & H(v, \beta, \rho) := - \sum_{i:i \in \mathcal{F}} \int_0^T \rho_i(L_i, t) v_i(L_i, t) dt \\
 & - \frac{1}{M} \sum_{i=1}^N \int_0^T \int_0^{L_i} \log((\rho_{\max} - \rho_i(x_i, t)) \\
 & (v_{\max,i} - v_i(x_i, t))(v_i(x_i, t) - v_{\min,i})) dx_i dt \\
 & - \frac{1}{M} \sum_{i:i \in \mathcal{D}} \int_0^T \log(\beta_i(t)(1 - \beta_i(t))) dt \\
 \text{s.t.} : & (6). \quad (8)
 \end{aligned}$$

We call $H(v, \beta, \rho)$ the augmented cost function. When ρ, β , and v are used without indices, it means that they are vectors, i.e., $v = [v_1, \dots, v_N]$. Note that the two last constraints in the optimization program (7) have disappeared into the cost function. This constrained optimization problem is easier to solve in practice. It is asymptotically equivalent to the problem of interest when $M \rightarrow +\infty$. We use an adjoint method to algebraically compute the gradient of the cost function. This method was extensively used [8] in flow control. We now adapt the adjoint method to the case in which we have a set of PDEs coupled through the boundary conditions, and subject to constraints. The adjoint method computes the gradient of the cost function $H(v, \beta, \rho)$ when ρ is an implicit function of v and β via the dynamics (6). Let us denote \mathcal{J} the cost function of the two variables v and β : $\mathcal{J} : (v, \beta) \rightarrow \mathcal{J}(v, \beta) = H(v, \beta, \rho)$, where ρ is the solution of the PDE system (6). We compute the linearized (6), which we will use to compute the gradient of the cost function in the optimization program (8). We denote by $(\cdot)'$ the linearized quantities around a nominal value denoted by (\cdot) : $\rho_i = \bar{\rho}_i + \rho'_i$. We call $\mathcal{N}'_i(\cdot)$ the linearized LWR operator and $q_i = \rho_i v_i$. In order to abbreviate the notation, we will write $q'_i = \rho'_i \bar{v}_i + \bar{\rho}_i v'_i$ and $\bar{q}_i = \bar{\rho}_i \bar{v}_i$. We omit the time and space dependence when they are obvious. The linearized (6) reads

$$\begin{cases}
 \mathcal{N}'_i(\rho'_i) := \mathcal{N}'_i \rho'_i = \frac{\partial \rho'_i}{\partial t} + \frac{\partial \rho'_i \bar{v}_i}{\partial x} = - \frac{\partial \bar{\rho}_i v'_i}{\partial x} & \forall i \in \{1, \dots, N\} \\
 \rho'_i(x, 0) = 0 & \forall i \in \{1, \dots, N\} \\
 q'_i(0, t) = \sum_{j \in \mathcal{U}(i)} q'_j(L_j, t) & \forall i \in \mathcal{M} \\
 q'_{ii}(0, t) = \beta'_i(t) \bar{q}_i(L_i, t) + \bar{\beta}_i(t) q'_i(L_i, t) & \forall i \in \mathcal{D} \\
 q'_{ir}(0, t) = -\beta'_i(t) \bar{q}_i(L_i, t) + (1 - \bar{\beta}_i(t)) q'_i(L_i, t) & \forall i \in \mathcal{D} \\
 \rho'_i(0, t) \bar{v}_i(0) + \bar{\rho}_i(0, t) v'_i(0, t) = 0 & \forall i \in \mathcal{S}.
 \end{cases} \quad (9)$$

³<http://cherokee.stanford.edu/~bayen/TCST06.html>.

The first variation of \mathcal{J} is obtained from (8)

$$\begin{aligned} \mathcal{J}' = & - \sum_{i:i \in \mathcal{F}} \int_0^T \rho'_i(L_i, t) \bar{v}_i(L_i, t) + \bar{\rho}_i(L_i, t) v'_i(L_i, t) \\ & + \frac{1}{M} \sum_{i:i \in \mathcal{D}} \int_0^T \left(\frac{\beta'_i(t)}{1 - \bar{\beta}_i(t)} - \frac{\beta'_i(t)}{\bar{\beta}_i(t)} \right) dt \\ & + \frac{1}{M} \sum_{i=1}^N \int_0^T \int_0^{L_i} \left(\frac{\rho'_i(x_i, t)}{\rho_{\max, i} - \bar{\rho}_i(x_i, t)} \right. \\ & \left. + \frac{v'_i(x_i, t)}{v_{\max, i} - \bar{v}_i(x_i, t)} - \frac{v'_i(x_i, t)}{\bar{v}_i(x_i, t) - v_{\min, i}} \right) dx_i dt. \end{aligned}$$

An integration by parts leads to the following identity for any two functions ρ'_i and ρ_i^*

$$\begin{aligned} \int_0^T \int_0^{L_i} \rho_i^* \mathcal{N}'_i \rho'_i dx_i dt = & \int_0^T [\rho_i^* \rho'_i \bar{v}_i]_0^{L_i} dt \\ & - \int_0^T \int_0^{L_i} \rho'_i \bar{v}_i \frac{\partial \rho_i^*}{\partial x_i} dt dx_i \\ & + \int_0^{L_i} [\rho_i^* \rho'_i]_0^T dx_i \\ & - \int_0^T \int_0^{L_i} \rho'_i \frac{\partial \rho_i^*}{\partial t} dt dx_i \end{aligned}$$

which can be rewritten using the standard inner product denoted $\langle \cdot, \cdot \rangle_i$ for the domain $[0, L_i] \times [0, T]$

$$\langle \rho_i^* | \mathcal{N}'_i \rho'_i \rangle_i = \langle \mathcal{N}_i^* \rho_i^* | \rho'_i \rangle_i + b_i \quad (10)$$

where

$$\begin{cases} \mathcal{N}_i^* = -\frac{\partial(\cdot)}{\partial t} - \bar{v}_i \frac{\partial(\cdot)}{\partial x_i} \\ b_i = \int_0^{L_i} [\rho_i^* \rho'_i]_0^T dx_i + \int_0^T [\rho_i^* \rho'_i \bar{v}_i]_0^{L_i} dt. \end{cases}$$

We will denote by $\langle \cdot, \cdot \rangle_{[0, T]}$ the standard inner product in $[0, T]$. \mathcal{N}_i^* is called the adjoint operator of \mathcal{N}'_i . In order to express the first variation of \mathcal{J} as a function of the v'_i and β'_i only, we choose an adjoint density field ρ_i^* that cancels all the terms containing ρ'_i in the cost function. First, in order to eliminate the term $(1/M) \sum_{i=1}^N \int_0^T \int_0^{L_i} \rho'_i(x_i, t) / (\rho_{\max, i} - \bar{\rho}_i(x_i, t)) dx_i dt$, we choose ρ_i^* such that

$$\mathcal{N}_i^* \rho_i^* = \frac{1}{M(\rho_{\max, i} - \bar{\rho}_i)}. \quad (11)$$

This is a first-order linear hyperbolic PDE, which is well posed if $\bar{\rho}_i$ is known and both the boundary conditions at one location and the initial conditions at one time are specified. This allows




Fig. 8. Network model.

us to enforce two other conditions for ρ_i^* in order to cancel all the terms containing ρ'_i . We can choose

$$\begin{cases} \rho_i^*(x_i, T) = 0 & \forall i \in \{1, \dots, N\}, \forall x_i \in [0, L_i] \\ \rho_i^*(L_i, t) = -1 & \forall i \in \mathcal{F}, \forall t \in [0, T] \\ \rho_i^*(0, t) = \rho_j^*(L_j, t), \forall i \in \mathcal{M}, \forall j \in \mathcal{U}(i), \forall t \in [0, T] \\ \rho_i^*(L_i, t) = \bar{\beta}_i(t) \rho_{i_r}^*(0) + (1 - \bar{\beta}_i(t)) \rho_{i_r}^*(0) & \forall i \in \mathcal{D}, \forall t \in [0, T]. \end{cases} \quad (12)$$

These conditions have been chosen by necessity of the algebraic derivation, in order to cancel appropriate terms in the perturbation of the cost function. After some algebra, using (10)–(12), we are able to express the first variation \mathcal{J}' of \mathcal{J} as a function of the first variations control variables only (v'_i and β'_i), as well as nominal and adjoint quantities, which we can evaluate. The result reads

$$\begin{aligned} \mathcal{J}' = & \sum_{i=1}^N \left\langle \bar{\rho}_i \frac{\partial \rho_i^*}{\partial x_i} + \frac{1}{M} \left(\frac{1}{v_{\max, i} - \bar{v}_i} - \frac{1}{\bar{v}_i - v_{\min, i}} \right) v'_i \right\rangle_i \\ & + \sum_{i \in \mathcal{D}} \left\langle \bar{\rho}_i(L_i) \bar{v}_i(L_i) (\rho_{i_r}^*(0) - \rho_{i_r}^*(0)) \right. \\ & \left. + \frac{1}{M} \left(\frac{1}{1 - \bar{\beta}_i} - \frac{1}{\bar{\beta}_i} \right) \beta'_i \right\rangle_{[0, T]} \end{aligned}$$

where again $\langle \cdot, \cdot \rangle_i$ denotes the inner product for the domain $[0, L_i] \times [0, T]$ and $\langle \cdot, \cdot \rangle_{[0, T]}$ for $[0, T]$. The functions $v_i(\cdot, \cdot)$ and $\beta_i(\cdot)$ generated by this method might be ill-behaved and, thus be inappropriate for practical air traffic control applications. We can alleviate this difficulty by projecting the descent direction $\bar{\rho}_i(\partial \rho_i^* / \partial x_i) + 1/M(1/(v_{\max, i} - \bar{v}_i) - 1/(\bar{v}_i - v_{\min, i}))$ into a vector space \mathcal{E} of appropriate functions, for example the set of continuous functions with bounded derivative, or the set of continuous piecewise affine functions.

VI. CONTROLLER DESIGN

In this section, we demonstrate the effectiveness of the adjoint method by applying it to the air traffic model. Fig. 8 shows the area which we will control (enclosed by a box). The inflows into the box are thus now q_1^{in} and q_2^{in} as shown in Fig. 8. We want to impose the following constraint: for all links, the density should be below a threshold ρ_{\max} which we impose. We allow the flow

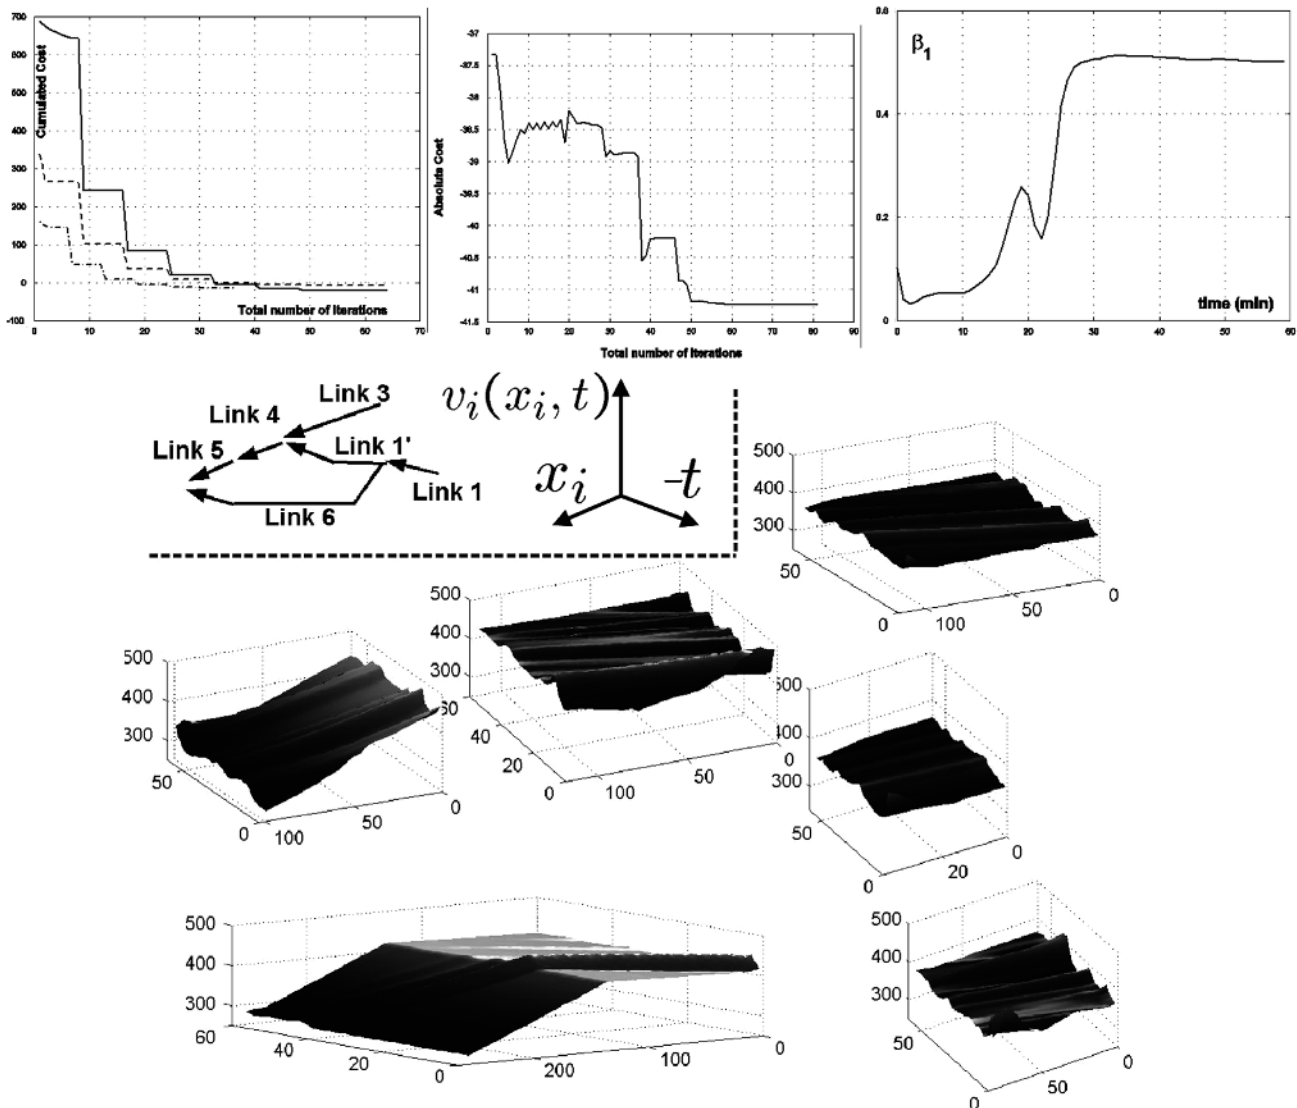


Fig. 9. Top left: Decrease of the cost as a function of the iterations for the three scenarios. The increases in M are clearly visible (steps), while the gradient descent is more subtle. Congested traffic (solid); heavy traffic (- -); normal traffic (—). Top middle: Decrease of the true cost as a function of the iterations. The true cost is the cost \mathcal{J} without the barrier terms. The method does not guarantee the monotonicity of the decrease but only the convergence. Top right: Evolution of the β_1 parameter as a function of time. Bottom: Evolution of the velocity fields as a function of time for the different links. Each of the plots corresponds to a link, (see top-left corner). The axis of each subplot are: x_i (arclength along the link), t (time) and $v_i(x_i, t)$, the velocity distribution.

to be split into a new link (link 6), in order to aid satisfaction of the maximal density constraints. We call β_1 the corresponding split factor: β_1 is the fraction of the flow which stays on link 1 (called 1 bis); $1 - \beta_1$ is the fraction which is routed through link 6. This new link might use another arrival into the airport (it enters the arrival airspace from another direction).⁴ We simulate the following three scenarios.

Scenario 1: Normal Traffic. (Real data) We take ETMS data, from which we extract initial conditions and inflows. We impose a restriction on the density and control the flow.

Scenario 2: Heavy Traffic. (Modified real data) We take the same data as for the previous case, and add additional aircraft in order to more heavily overload the network.

⁴Note that using β_1 is equivalent to using turning proportions in road traffic and might not be the best way to represent network traffic. It could be better to define an assignment proportion, i.e., a coefficient indexed by destination. This might be implemented in the future (as a part of the control strategy), using a framework such as the one developed by Papageorgiou [25].

Scenario 3: Congested Network. We generate data with very high densities of aircraft. This situation does not use ETMS data; it is generated randomly.

Fig. 9 shows the decrease in cost for the three scenarios as a function of the total number of iterations (i.e., iterations on M and gradient advances). As can be seen in this figure, the more congested the situation is, the higher the cost. The evolution of the cost with iterations exhibits two distinct behaviors, as often with barrier methods [10]: large jumps corresponding to the increases in M , and shallower decreases corresponding to the gradient advances. Convergence is clearly observed for the three scenarios. We display some of the results for the third case. An animation (in form of an .avi movie file) corresponding to each of the three scenarios is available.³ We now describe in detail the scenario corresponding to Case 3. We run a one-hour simulation. Fig. 9 shows the aircraft density on all links at various instants, in the absence of control: the velocity is the mean velocity profile determined for each link, and no aircraft is al-

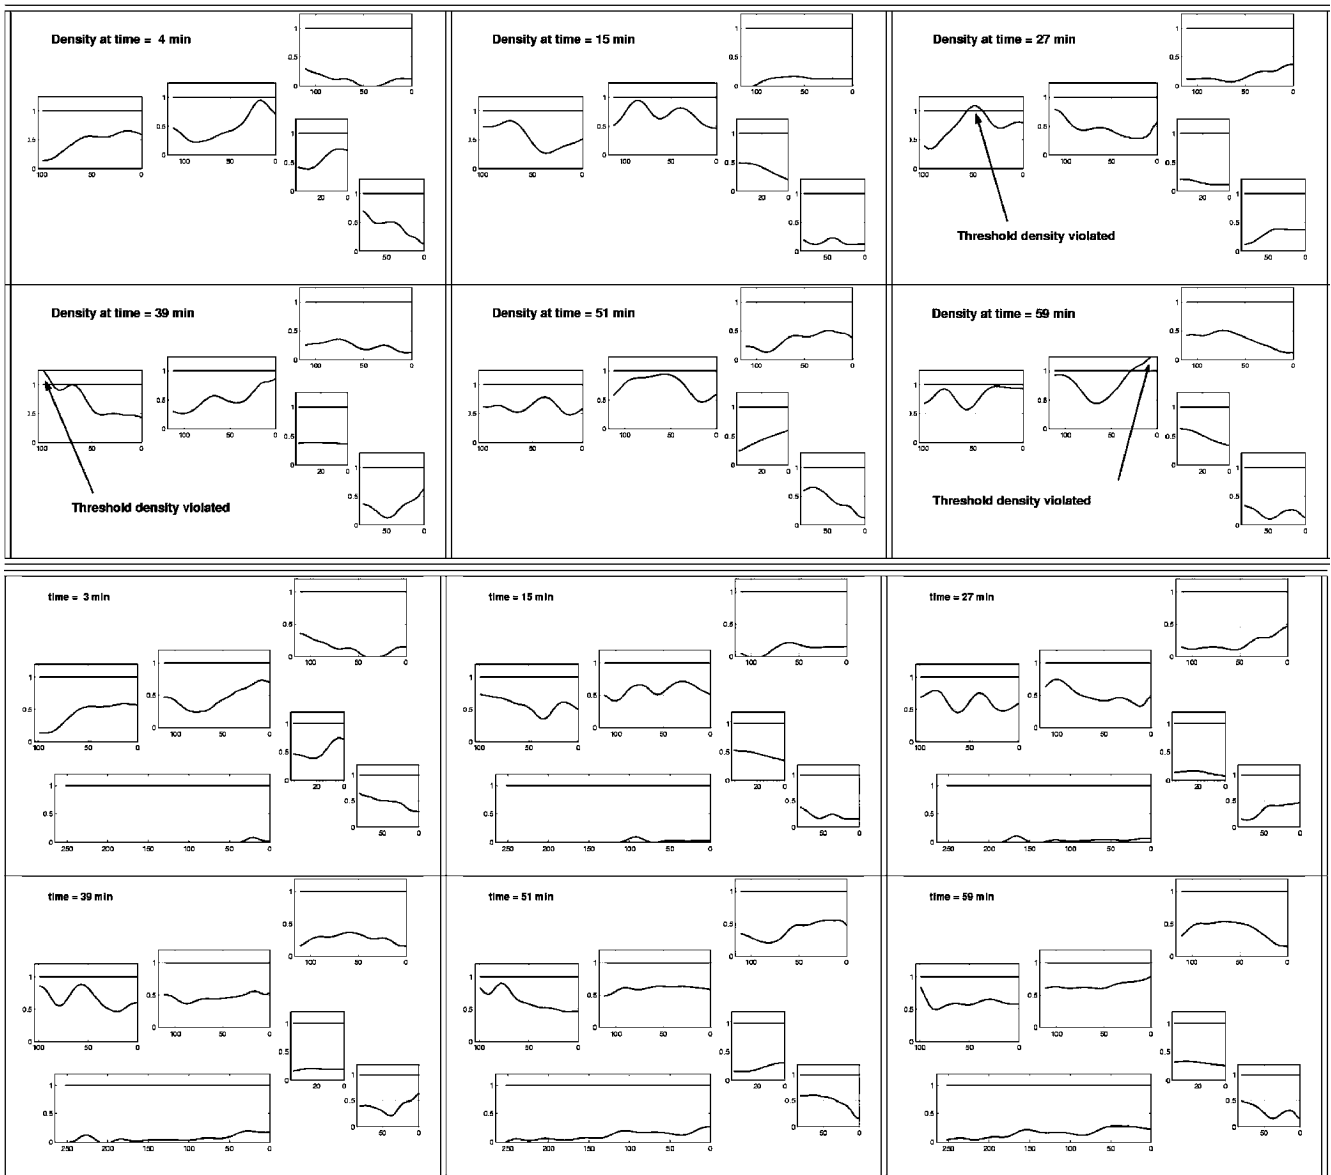


Fig. 10. Top 6 subfigures: Evolution of the aircraft density on the different links in the absence of control. Each of the subplot shows the density distribution at a given time on the corresponding link as in Fig. 9 (the horizontal coordinate represents location, the vertical represents density). The horizontal line represents the density threshold (all quantities are nondimensionalized by ρ_{\max} , so that the threshold density is 1). As can be seen, the density threshold is violated in link 5 at $t = 27$, $t = 39$ and $t = 59$. Bottom 6 subfigures: Evolution of the aircraft density with control applied. Note that link 6 is now open and used. This prevents the second violation of density threshold observed in Fig. 10 ($t = 59$): some of the flow is directly routed from link 1 to link 6. The first violation seen in the top 6 subfigures is avoided by speed changes. This figure is also available in form of a .avi file.³

lowed into link 6 (i.e., $\beta_1 = 1$). The initial density is shown in the top-left corner. The inflow into links 1 and 3 is such that at time $t = 27$, the density threshold (represented by the horizontal line on each subplot) is violated until time $t = 45$. At time $t = 55$, it is violated again, until the end of the experiments. Fig. 9 shows the same experiment when link 6 is now opened to traffic, and velocity control is enabled. As can be seen, about half of the traffic incoming into link 1 is rerouted into link 6, and the other half into link 1 bis. Fig. 9 shows the variation of β_1 with time. As can be seen, around $t = 20$ min, there is a peak of about 25% of aircraft routed into link 6, which settles to 50% at $t = 30$. The routing control enables avoidance of violation of maximal density shown in Fig. 10. The first violation is avoided by velocity changes.

The velocity profiles $v_i(x_i, t)$ are shown in Fig. 9. Each of the subplots corresponds to one of the links. For links 5 and 6, one can clearly see the descent velocity profiles. Also, for link 6 (subfigure below), one can see a ridge. It corresponds to a set of aircraft which have to fly at high speed into the airport. One can also see similar ridges on the other subplots, which have the same interpretation. For any ridge, the Controller command could be to the corresponding set of aircraft: “fly direct at 420 kn direct into [the next waypoint].” Note that in the absence of control, the first violation of the aircraft density threshold occurs 33 min after the beginning of the experiment, almost at the end of the network, which is not intuitive. This shows the efficiency of the method, which is capable of generating the right routing and speed assignments to prevent undesirable events from hap-




Fig. 11. Display of the traffic situation for the static validation. The density of the links is depicted by the color. The colored rectangles shown in this plot represent the density. The color scale is: white for zero density; black for highest density. The actual aircraft positions are superimposed (triangles). Traffic is shown at t_0 (top), $t_0 + 8$ min, $t_0 + 16$ min, etc. As can be seen, the peaks of density corresponds to the actual positions of the aircraft.

pening much later. Finally, the simulations are also depicted on a U.S. map in Fig. 12. One can see that before $t = 27$, all aircraft choose the direct route through link 5 to Chicago (it is shorter). After $t = 27$, the excessive amount of flow incoming into links 1 and 3 forces the flow to be split through links 1 bis and 6.

The expression of the cost function \mathcal{J} can be replaced by any arbitrary user-defined cost as long as the integrand is smooth. The goal of this paper was to prove the feasibility of the technique (with a particular cost function), but extending this to any cost function is a straightforward process (the only thing which is needed is to recompute the expression of the gradient based on the new cost, following the steps outlined here). In particular, in the work of [27], the authors use an integral form with quadratic penalty. This can be interpreted as penalizing the cumulative delay minutes at each point in time, and penalizing more severely large deviations from the scheduled flow than small ones.

VII. CONCLUSION

We have derived an Eulerian model of the airspace based on a modified LWR partial-differential equation. The network structure of the airspace was modeled as a set of coupled LWR PDEs. Given initial positions of aircraft and airport inflows, this system of PDEs enables the prediction of the aircraft density. An analytical solution was derived for a single link in the case, in which the mean velocity profiles of aircraft along airways do not vary with time (just with space). It can be used for multiple links as well. ETMS data was used to identify the numerical parameters associated with this model. The data is also used to validate the model, i.e., to demonstrate good predictive capability of this method.

We first have shown how to use efficient numerical schemes to simulate the network. We have discarded linear numerical schemes because of their poor performance. We have used the Jameson–Schmidt–Turkel scheme as our main numerical scheme to perform numerical simulations of the network. We have posed the problem of maximizing throughput at a destination airport while maintaining the aircraft density below a certain threshold as an optimization program. The inequality constraints of this program have been handled using a log-barrier method. The adjoint problem was derived and used to compute the gradient of the augmented cost function. The resulting optimization and control schemes were applied to a real air traffic case. Simulations show that this method enables automated control of realistic scenarios as well as highly congested situations. The output of the code is a set of time dependent velocity profiles to apply to the network, and a policy telling how to split the flow in areas of diverging traffic. These outputs could, thus, be used by the Traffic Management Unit in charge of managing the flow: they provide high-level policies to apply to the aircraft streams, which are directly understandable by human controllers.

Finally, this formulation of the air traffic flow control problem as an optimization program of PDEs allows for many refinements in the control procedure. For instance, gradient descent may be replaced by more sophisticated optimization methods such as approximate Newton method [29] in order to ensure real time convergence of the algorithm. Furthermore, using this model, a decentralized control policy can also be derived using

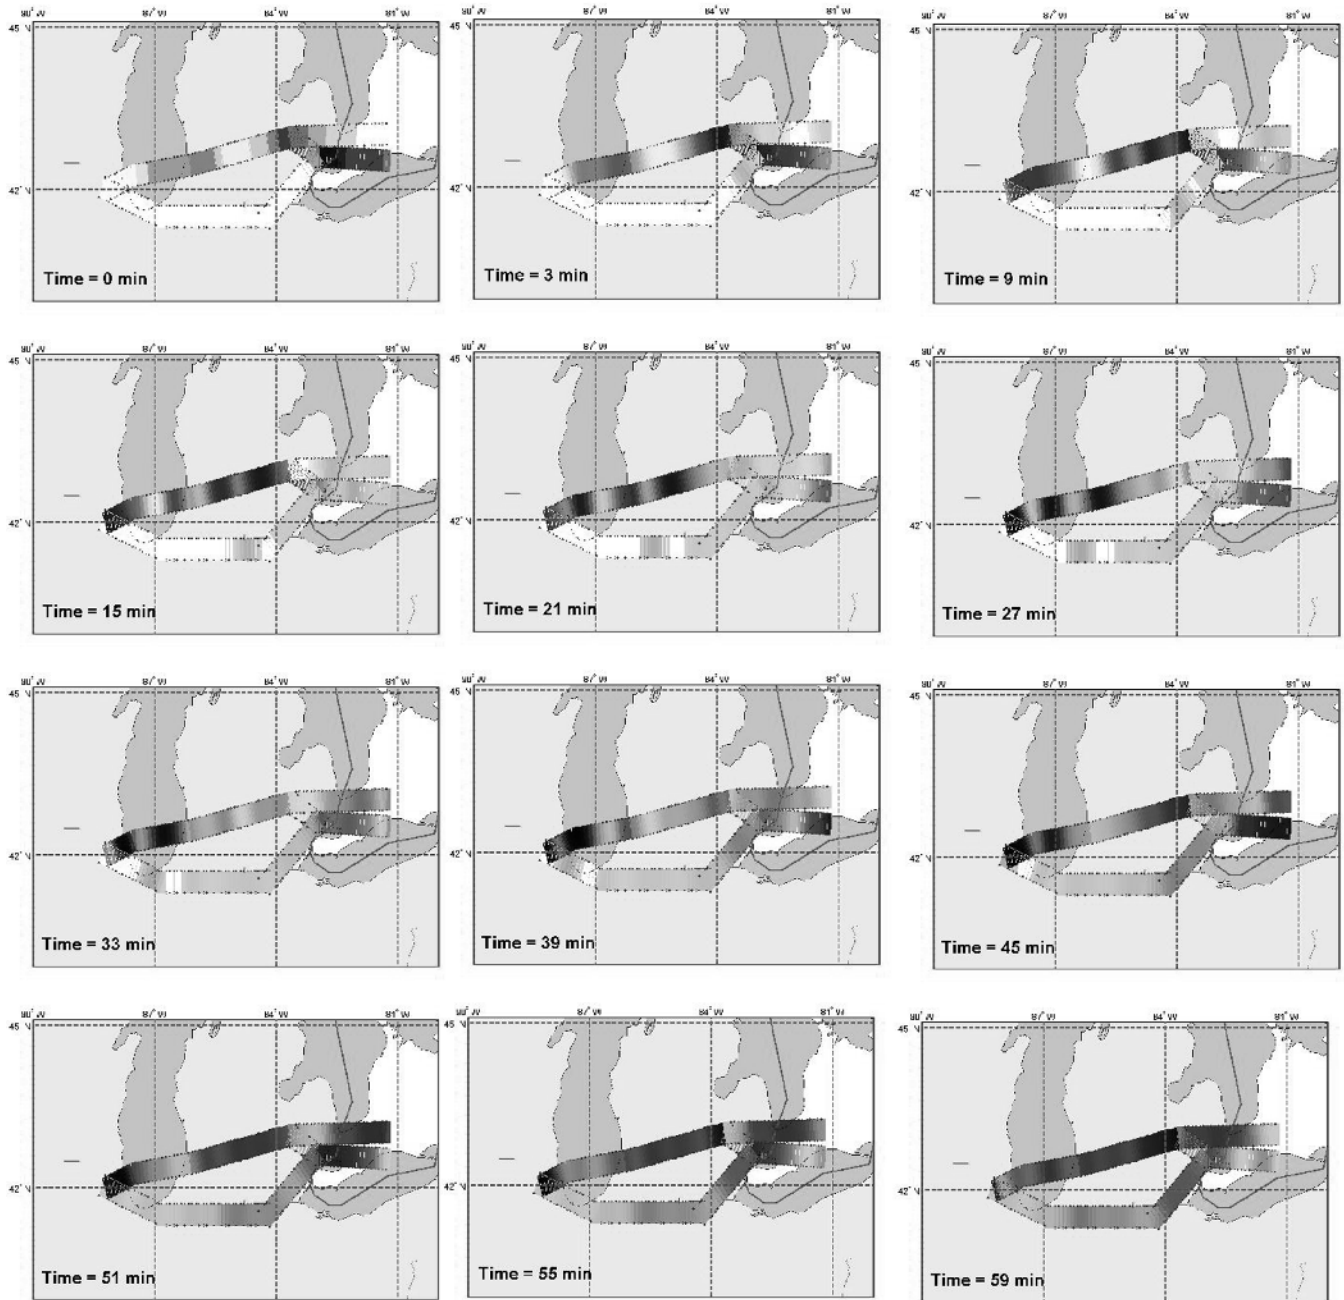


Fig. 12. Aircraft density in the network around Chicago in presence of control and velocity assignment. The density of the links is depicted by the color. The colored rectangles shown in this plot represent the density. The color scale is: white for zero density; black for highest density. As can be seen and was shown in Fig. 10, a good portion of the flow is routed into link 6 starting at time $t = 51$.

decomposition techniques in order to allow different airlines to separately optimize the flow of their aircraft, while maintaining safety criteria [28]. This method has since been applied to highway traffic as well [5], [18] and looks promising for other applications involving networks of partial differential equations.

VIII. PROOF OF PROPOSITION 1

Existence: K is well defined because $v(x) \geq m$ for all $x \in [0, L]$. Its inverse exists because K is (strictly) increasing. It is easy to check that (2) satisfies (1) almost everywhere, and

that it is continuous. This solution has been constructed using a technique analogous to the algorithm of Bayen and Tomlin [7] based on the method of characteristics.

Uniqueness: Let us call n_1 and n_2 two continuous weak solutions of (1). Call $\delta := n_1 - n_2$. δ satisfies: $\partial\delta/\partial t + v(x)\partial\delta/\partial x = 0$ a.e. in $[0, L] \times (0, T]$, $\delta(x, 0) = 0$ in $[0, L] \times \{0\}$ and $\delta(0, t) = 0$ in $\{0\} \times (0, T]$. Multiplying this PDE by δ and integrating from $x_0 = 0$ to the first discontinuity x_1 of $v(\cdot)$ gives

$$\int_{x_0}^{x_1} \delta(u, t) \frac{\partial\delta}{\partial t}(u, t) du + \int_{x_0}^{x_1} v(u) \delta(u, t) \frac{\partial\delta}{\partial x}(u, t) du = 0$$

from which we deduce

$$\frac{1}{2} \frac{d}{dt} \int_{x_0}^{x_1} \delta(u, t)^2 du + \int_{x_0}^{x_1} v(u) \delta(u, t) \frac{\partial \delta}{\partial x}(u, t) du = 0.$$

Integrating by parts gives

$$\begin{aligned} \frac{1}{2} \frac{d}{dt} \int_{x_0}^{x_1} \delta(u, t)^2 du &\leq \int_{x_0}^{x_1} v'(u) \frac{1}{2} \delta(u, t)^2 du \\ &\quad - \left[v(u) \frac{1}{2} \delta(u, t)^2 \right]_{x_0}^{x_1} \\ &\leq \int_{x_0}^{x_1} v'(u) \frac{1}{2} \delta(u, t)^2 du \end{aligned}$$

since $\delta(x_0, t) = 0$ and $v(x_1) > 0$. Using the fact that $v(\cdot) \in C_1([x_0, x_1])$, $\exists M > 0$, $|v'(x)| \leq M$ for all $x \in [x_0, x_1]$. Then, we use the fact that

$$\int_{x_0}^{x_1} v'(u) \frac{1}{2} \delta(u, t)^2 du \leq M \int_{x_0}^{x_1} \frac{1}{2} \delta(u, t)^2 du$$

so that we can rewrite the inequality as

$$\frac{1}{2} \frac{d}{dt} \int_{x_0}^{x_1} \delta(u, t)^2 du \leq M \int_{x_0}^{x_1} \frac{1}{2} \delta(u, t)^2 du.$$

Using Gronwall's lemma, this last inequality implies $\delta(x, t) = 0$ almost everywhere in $[x_0, x_1]$. By continuity, $n_1(x, t) = n_2(x, t)$ everywhere in $[x_0, x_1]$, and, therefore, at x_1 . The same proof applies to $[x_1, x_2]$ since $n_1(x_1, t) = n_2(x_1, t)$ for all t . By induction on x_k , n_1 and n_2 are equal everywhere in $[x_0, x_k]$ and, therefore, in $[0, L]$.


ACKNOWLEDGMENT

The authors would like to thank Dr. P. K. Menon and Dr. K. Bilimoria for conversations which inspired this work, Dr. G. Chatterji for his ongoing support and suggestions which went into modelling this work, and Dr. G. Meyer for his support in this project. They also thank Prof. T. Bewley for useful conversations regarding the application of the adjoint method to flow control, and his help in the original formulation of the control problem. Prof. T.-P. Liu helped define the PDE used for this model.

REFERENCES

- [1] O. M. Aamo and M. Krstic, *Flow Control by Feedback*. New York: Springer-Verlag, 2002.
- [2] R. Ansorge, "What does the entropy condition mean in traffic flow theory?," *Transp. Res.*, vol. 24B, no. 2, pp. 133–143, Apr. 1990.
- [3] B. Bamieh, F. Paganini, and M. A. Dahleh, "Distributed control of spatially-invariant systems," *IEEE Trans. Autom. Control*, vol. 47, no. 7, pp. 1091–1107, Jul. 2002.
- [4] A. M. Bayen, "Computational control of networks of dynamical systems: Application to the National Airspace System," Ph.D. dissertation, Dept. Aeronautics and Astronautics, Stanford Univ., Stanford, CA, 2004.
- [5] A. M. Bayen, R. Raffard, and C. J. Tomlin, "Eulerian network model of air traffic flow in congested areas," in *Proc. Amer. Contr. Conf.*, 2004, pp. 5520–5526.
- [6] A. M. Bayen, R. Raffard, and C. J. Tomlin, "Network congestion alleviation using adjoint hybrid control: Application to highways," in *Number 2993 in Lecture Notes in Computer Science*. New York: Springer-Verlag, 2004, pp. 95–110.
- [7] A. M. Bayen and C. J. Tomlin, "A construction procedure using characteristics for viscosity solutions of the Hamilton-Jacobi equation," in *Proc. 40th IEEE Conf. Dec. Contr.*, 2001, pp. 1657–1662.
- [8] T. R. Bewley, "Flow control: New challenges for a new renaissance," *Prog. Aerosp. Sci.*, vol. 37, no. 1, pp. 1–119, Jan. 2001.
- [9] K. Bilimoria, B. Sridhar, G. Chatterji, K. Seth, and S. Graabe, "FACET: Future ATM concepts evaluation tool," in *Proc. 3rd USA/Eur. Air Traffic Manage. R&D Seminar*, 2001, on CDROM.
- [10] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [11] G. Chatterji, B. Sridhar, and D. Kim, "Analysis of ETMS data quality for traffic flow management decisions," in *Proc. AIAA Conf. Guid., Nav. Contr.*, 2003, AIAA-2003-5626.
- [12] C. Chen, Z. Jia, and P. Varaiya, "Causes and cures of highway congestion," *IEEE Control Syst. Mag.*, vol. 21, no. 4, pp. 26–33, Dec. 2001.
- [13] P. D. Christofides, *Nonlinear and Robust Control of Partial Differential Equation Systems: Methods and Applications to Transport-Reaction Processes*. Cambridge, MA: Birkhäuser, 2001.
- [14] C. Daganzo, "The cell transmission model: A dynamic representation of highway traffic consistent with the hydrodynamic theory," *Transport. Res.*, vol. 28B, no. 4, pp. 269–287, Aug. 1994.
- [15] C. Daganzo, "The cell transmission model, Part II: Network traffic," *Transport. Res.*, vol. 29B, no. 2, pp. 79–93, Apr. 1995.
- [16] C. Daganzo, "A finite difference approximation of the kinematic wave model of traffic flow," *Transport. Res.*, vol. 29B, no. 4, pp. 261–276, Aug. 1995.
- [17] C. Hirsch, *Numerical Computation of Internal and External Flows*. New York: Wiley, 1988.
- [18] D. Jacquet, C. Canudas de Wit, and D. Koenig, "Optimal ramp metering strategy with an extended LWR model: Analysis and computational methods," in *Proc. 16th IFAC World Congr.*, 2005, to be published.
- [19] A. Jameson, "Aerodynamic design via control theory," *J. Scientific Comput.*, vol. 3, no. 3, pp. 233–260, Sep. 1988.
- [20] A. Jameson, "Analysis and design of numerical schemes for gas dynamics I: Artificial diffusion, upwind biasing, limiters and their effect on accuracy and multigrid convergence," *Int. J. Computational Fluid Dyn.*, vol. 4, pp. 171–218, Sep. 1995.
- [21] M. Krstic, "On global stabilization of Burgers' equation by boundary control," *Syst. Control Lett.*, vol. 37, pp. 123–142, Jul. 1999.
- [22] M. J. Lighthill and G. B. Whitham, "On kinematic waves. II. A theory of traffic flow on long crowded roads," in *Proc. Royal Soc. London*, 1956, pp. 317–345.
- [23] X. Litrico, "Robust IMC flow control of SIMO dam-river open-channel systems," *IEEE Trans. Control Syst. Technol.*, vol. 10, no. 5, pp. 432–437, May 2002.
- [24] P. K. Menon, G. D. Sweriduk, and K. Bilimoria, "New approach for modeling, analysis, and control of air traffic flow," *AIAA J. Guid., Contr., Dyn.*, vol. 27, no. 5, pp. 731–5090, Sep/Oct. 2004.
- [25] A. Messmer and M. Papageorgiou, "Route diversion control in motorway networks via nonlinear optimization," *IEEE Trans. Control Syst. Technol.*, vol. 3, no. 1, pp. 144–154, Mar. 1995.
- [26] L. Munoz, X. Sun, R. Horowitz, and L. Alvarez, "Traffic density estimation with the cell transmission model," in *Proc. Amer. Contr. Conf.*, 2003, pp. 3750–3755.
- [27] R. Raffard, S. L. Waslander, A. M. Bayen, and C. J. Tomlin, "Toward efficient and equitable distributed air traffic flow control," presented at the Proc. Amer. Contr. Conf., 2006, Minneapolis, MN.
- [28] R. Raffard, S. L. Waslander, A. M. Bayen, and C. J. Tomlin, "A cooperative, distributed approach to multi-agent Eulerian network control: Application to air traffic management," in *Proc. AIAA Guid., Nav. Contr. Conf. Exhibit*, 2005, AIAA-2005-6050.
- [29] R. L. Raffard and C. J. Tomlin, "Second order optimization of ordinary and partial differential equations with applications to air traffic flow," in *Proc. Amer. Contr. Conf.*, 2005, pp. 798–803.
- [30] P. I. Richards, "Shock waves on the highway," *Oper. Res.*, vol. 4, no. 1, pp. 42–51, Feb. 1956.
- [31] C. Robelin, D. Sun, G. Wu, and A. M. Bayen, "En-route air traffic modeling and strategic flow management using mixed integer linear programming," in *INFORMS Annu. Meeting*, 2005.
- [32] R. C. Smith and M. A. Demetriou, *Research Directions in Distributed Parameter Systems*. Philadelphia, PA: SIAM, 2000.

- [33] D. Sun, S. Yang, I. S. Strub, B. Sridhar, K. Sheth, and A. M. Bayen, "Assessment of the respective performance of three Eulerian air traffic flow models," presented at the Proc. AIAA Conf. Guid., Nav. Contr., Keystone, CO, 2006.
- [34] Y. Wang, M. Papageorgiou, and A. Messmer, "Motorway traffic state estimation based on extended Kalman filter," in *Proc. Euro. Contr. Conf.*, 2003.



Alexandre M. Bayen (S'02–M'04) received the B.S. degree in applied mathematics from the Ecole Polytechnique, Palaiseau, France, in 1998, the M.S. and Ph.D. degrees in aeronautics and astronautics from Stanford University, Stanford, CA, in 1999 and 2004, respectively.

He was a Visiting Researcher at NASA Ames, Moffett Field, CA, from 2001 to 2003. From 2004 to 2005, he worked for the Department of Defense, France, where he held the rank of Major. During that time, he was the Research Director of the Laboratoire de Navigation Autonome at the Laboratoire de Recherches Balistiques et Aérodynamiques, Vernon, France. Since March 2005, he has been an Assistant Professor in the Department of Civil and Environmental Engineering at the University of California at Berkeley, Berkeley. His research interests include control of distributed parameter systems, combinatorial optimization, hybrid systems, and air traffic automation.

Dr. Bayen is a recipient of the Graduate Fellowship of the Délégation Générale pour l'Armement (1998–2002) from France, and the Ballhaus Prize for best doctoral thesis from the Department of Aeronautics and Astronautics at Stanford University (2004).



Robin L. Raffard received the M.S. degree in mechanical engineering from the Ecole Centrale Paris, France, in 2002, and the M.S. degree in aeronautics and astronautics from Stanford University, Stanford, CA, in 2003. He is currently working towards the Ph.D. degree in aeronautics and astronautics at the same school.

His research interests include distributed optimization and optimal control of systems governed by differential equations, with applications in air traffic flow, systems biology, and stochastic systems.



Claire J. Tomlin (S'93–M'99) received the Ph.D. degree in electrical engineering from the University of California at Berkeley, Berkeley, in 1998. She also received the M.Sc. degree from Imperial College, London, in 1993, and the B.A.Sc. degree from the University of Waterloo, Canada, in 1992, both in electrical engineering.

She is an Associate Professor in the Department of Electrical Engineering and Computer Sciences at the University of California at Berkeley, and is an Associate Professor in the Department of Aeronautics and Astronautics at Stanford University, Stanford, CA, where she also holds the Vance D. and Arlene C. Coffman Faculty Scholarship in the School of Engineering. She joined Stanford in September 1998, as a Terman Assistant Professor, and received tenure at Stanford in November 2004. In July 2005, she joined Berkeley as an Associate Professor. She has held visiting research positions at NASA Ames, Honeywell Labs, and the University of British Columbia. Her research interests include control systems, specifically hybrid control theory, and she works on air traffic control automation, flight management system analysis and design, and modeling and analysis of biological cell networks.

Dr. Tomlin is a recipient of the Eckman Award of the American Automatic Control Council (2003), MIT Technology Review's Top 100 Young Innovators Award (2003), the AIAA Outstanding Teacher Award (2001), a National Science Foundation (NSF) Career Award (1999), and the Bernard Friedman Memorial Prize in Applied Mathematics (1998).