

Market-Based Air Traffic Flow Control with Competing Airlines

Steven L. Waslander* and Robin L. Raffard†
Stanford University, Stanford, California 94305

and

Claire J. Tomlin‡
University of California, Berkeley, Berkeley, California 94270

DOI: 10.2514/1.30595

To advance efficient and equitable use of the U.S. National Airspace System during weather disruptions, air traffic management is modeled as a network flow optimization program that explicitly incorporates airline preference information. Market mechanisms are proposed to perform distributed computation of efficient solutions based on novel cost metrics that accurately reflect the cost of delays to the airlines. Two distinct types of network flow models are presented to demonstrate the tradeoff between computational complexity and control input flexibility. The discrete path flow model describes a simplified approach by fixing flow velocity and limiting rerouting options, thereby satisfying the primary assumptions of general equilibrium theory to ensure efficiency of the market outcome. The continuous link flow model allows additional control inputs to better align with current methods used in dealing with adverse weather conditions; namely, ground-delay programs, miles-in-trail restrictions, and flight reroutes. The market-mechanism outcome for both models is shown to be preferred by all airlines over a solution determined without incorporating preference information. Simulation results are presented for feasible problem sizes of both flow models and demonstrate the gains that can be achieved by implementing market mechanisms for air traffic management.

Nomenclature

A	=	resource-path flow matrix
\mathcal{A}	=	allocation rule
C	=	resource capacity
\mathcal{D}	=	divide set
dt	=	discrete time step
\mathcal{F}	=	inflow set
\mathcal{I}	=	link set
J	=	cost function
\mathcal{J}	=	airline set
L_i	=	link length
\mathcal{M}	=	merge set
\mathcal{N}	=	node set
\mathcal{P}	=	path set
q^{in}	=	inflow
\mathcal{R}	=	resource set
\mathcal{S}	=	source set
T	=	time horizon
v	=	flow velocity
Y	=	cumulative flow allocation
y	=	flow allocation
\mathcal{Z}	=	sink set
α	=	price-update step size
β	=	flow split parameter
μ, λ	=	resource price
ν	=	resource payment

ρ	=	flow density
ω	=	path cost weighting

Subscripts

d	=	outflow link index
i	=	link index
j	=	airline index

Superscripts

ce	=	competitive equilibrium
d	=	destination
e	=	en route
m	=	market
o	=	origin
p	=	proportional
s	=	scheduled

I. Introduction

AIR traffic management for the U.S. National Airspace System (NAS) is one of the essential roles played by the Federal Aviation Administration (FAA) in maintaining safe commercial flight operations. The air traffic control system is responsible for ensuring that resource limits such as airport takeoff and landing rates are not exceeded, as well as ensuring that minimum en route aircraft separation is maintained for up to 45,000 flights a day [1]. Coordination of air traffic flow is aggravated by capacity restrictions caused by inclement weather, which are difficult to predict beyond three hours in advance, and complicated by the many stakeholders for whom flight delays result in significant financial burdens. Delays incurred by any individual flight can result in repercussions throughout the network, because aircraft, crews, and passengers may all be scheduled to continue on subsequent flights. The result is a highly connected, dynamic, and unpredictable resource allocation problem with costs worth potentially millions of dollars per day.

The FAA has three main avenues for reconfiguring scheduled operations in the face of airspace disruptions: ground-delay

Received 22 February 2007; revision received 29 May 2007; accepted for publication 5 June 2007. Copyright © 2007 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0731-5090/08 \$10.00 in correspondence with the CCC.

*Ph.D. Candidate, Department of Aeronautics and Astronautics; stevenw@stanfordalumni.org. Student Member AIAA.

†Ph.D. Candidate, Department of Aeronautics and Astronautics; rraffard@stanfordalumni.org. Student Member AIAA.

‡Professor, Department of Electrical Engineering and Computer Science; tomlin@stanford.edu. Member AIAA.

programs (GDPs), miles-in-trail restrictions, and flight reroutes. GDPs can be used to delay network traffic inflow bound for specific constrained resources such as destination airports and en route airspace. Miles-in-trail restrictions fix en route vehicle spacing on crowded routes to smooth flow into intersections and congested areas. Flight reroutes ensure that aircraft are redirected around weather disruptions and are managed through a playbook of fixed procedures developed over years of experience. Currently, each control method is managed separately and adjustments are made primarily through the intervention of experienced traffic control specialists.

The control actions undertaken by the FAA are quite effective for situations in which only a simple sequence of such decisions is needed to alleviate moderate capacity restrictions. The superposition of these actions, however, in situations with congested traffic and/or large weather-related disruptions may not lead to efficient or equitable outcomes. With air traffic levels from all types of service continuing to rise steadily and delays affecting over 25% of the flights in 2005 [1], the FAA has consistently called for improvements to existing methodologies for mitigating the effect of weather disruptions on the network [2]. By developing scalable real-time optimization techniques that incorporate preference information and simultaneously consider the effects of multiple control actions, this work presents a common framework for nationwide air traffic management that could result in improved network efficiency and reduced traffic delays.

A key aspect in ensuring efficient use of the NAS is to effectively incorporate airline preferences in air traffic flow optimization methodologies. It is the airlines that are in the best position to accurately determine the relative costs of delay for conflicting flights on constrained routes, and it is the airlines that must bear the cost of inefficient allocation of the available resources. By modeling the airlines as competitive agents and developing mechanisms that allocate resources without requiring disclosure of preferences, it is possible to realize significant gains in efficiency over centralized allocation methods that do not incorporate airline preferences.

Already, the airlines are capable of influencing traffic flow control decisions through the collaborative-decision-making (CDM) process to improve the efficiency of GDPs. Previous failure to reward the timely communication of information resulted in many airlines remaining mute about known delays and mechanical issues, for fear of being assigned a lower priority in landing-slot assignments [3]. This problem was addressed by modifying the landing-slot-assignment process during GDPs from first-come-first-served to one that fixed relative assignment priority based on the existing flight schedule. This modification resulted in much greater airline participation, and over 31 million GDP delay minutes have been avoided in the first decade of operation [4], an approximate 4% annual reduction. There has been much interest in extending the benefits of the CDM system by allowing more flexibility in slot allocation and further incorporating airlines in the decision-making process [3,5].

This work seeks to build on this foundation by developing market mechanisms for the more general air traffic flow control problem incorporating all three FAA control decisions outlined earlier. Similar approaches have been suggested for both electricity power generation [6], which operates daily to price the last 5% of supply in California, and Internet traffic routing [7], which is much harder to implement due to the highly distributed nature of the Internet. In a market mechanism for air traffic control, all airspace resources such as airport and airway capacity are assigned a price-per-unit capacity through an iterative process. At each step in the process, the airlines seek to minimize their delay costs over a finite time horizon by purchasing airspace resources at market prices, after which the FAA updates resource prices to penalize excess demand. The market mechanism can be interpreted as a distributed algorithm for solving the dual problem to the aggregate cost minimization over all airlines subject to the airspace capacity limitations. Convergence of the market mechanism can therefore be guaranteed through standard subgradient descent arguments from convex optimization. Under the

assumptions of general equilibrium theory, the resulting market equilibrium is guaranteed to be efficient [8].

A market mechanism for air traffic flow resource allocation could be implemented as a natural extension of the CDM process, which has opened real-time communication channels between the FAA air traffic management specialists and the airline operation centers responsible for daily flight and crew dispatch decisions. To be successful, however, airline incentive issues must be addressed to ensure both efficient outcomes and participation of the airlines in the process. To this end, a baseline solution can be defined through proportional allocation relative to a predefined flight schedule, and resource payments in the market mechanism are made relative to this default allocation. It is then possible to guarantee that all airlines will prefer the market outcome to the alternative of not participating at all.

This paper presents two traffic flow models that differ in their approach to contrast the benefits of each for market-based flow control. The *discrete path flow* model assumes that predefined alternate flight plans with fixed velocity profiles have been provided for each scheduled flight over a finite horizon with a finite set of time intervals. Inspired by previous work in Internet routing protocols [9,10], the model results in a convenient convex optimization formulation and meets the assumptions required to ensure efficiency of the market outcome, while still capturing many key aspects of the air traffic control problem. The *continuous link flow* model is defined in the continuous-time domain and does not fix either flow velocity or routing over the network, but instead incorporates inflow, velocity, and routing as control decisions. Although it does not satisfy the assumptions of general equilibrium theory [8], this model more accurately captures the true control possibilities, providing decision-makers with a high degree of control over aircraft flow in the network.

Air traffic modeling at the NAS level has been studied by several researchers and practitioners. Optimal allocation of GDP delays was proposed by Odoni et al. [11,12]. Extended to include en route capacity, an integer program formulation based on sector occupancy was presented by Bertsimas and Stock Patterson [13] for fixed vehicle routing, and for which integer solutions are often found when solving its linear programming (LP) relaxation. However, the dramatic benefit of needing only to solve a single LP quickly disappears if the utility structure is nonlinear or if flight rerouting is included [14]. As a result, combinatorial optimization techniques that directly address individual flight control seem to be impractical for full NAS modeling, and approximate optimization methods have been proposed to address the issue of exponential complexity [15,16].

As an alternative, aggregate models have been proposed that take a control-volume approach and model traffic over the network as the propagation of aircraft density profiles, in effect modeling air traffic as a continuous flow instead of tracking aircraft individually. This network flow method was first introduced by Lighthill and Whitham [17] and Richards [18] in highway-traffic modeling and leads to noninteger solutions to the traffic control problem. The first adaptation to air traffic modeling was a cell flow model introduced by Menon et al. [19], who presented a linear discretized control-volume method that generated much interest in the air traffic management community. A link flow partial-differential-equation (PDE) constrained formulation was proposed by Bayen et al. [20,21], which improves the accuracy of flow propagation and enables control of velocity and rerouting, at the cost of increased computational complexity. The continuous link flow model presented later is an extension of this PDE constrained formulation, with the addition of inflow control and more realistic cost metrics. For a recent survey of aggregate flow modeling techniques, see the work of Sun et al. [22].

The paper proceeds as follows. Section II presents the fundamental theory upon which market mechanisms for air traffic control rest. The discrete path flow model as well as an appropriate airline cost model are then introduced in Sec. III, and a market mechanism is designed for which efficiency and airline participation is guaranteed in Sec. IV. A network simulation is presented in Sec. V, which demonstrates the possibility of solving large network flow problems in the allowable

time frame while incorporating the latest weather information. Section VI introduces the continuous link flow model and a similar market mechanism is defined for this model in Sec. VII. Finally, a comparison of the two models is performed on a smaller network scenario in Sec. VIII.

II. Resource Allocation Equilibrium Properties

Market mechanisms can be thought of as a subset of noncooperative resource allocation methods that explicitly define a price for each resource to be allocated among a set of competitive agents. In the context of air traffic flow control, a resource might be defined as the flow into a destination airport or along an airway with a capacity limit of a certain number of aircraft per 15-min interval. The agents with private preference information are airlines competing to get their flights over the network with minimal delay, and prices could be set centrally by the FAA. Given a specific price for each resource, the agents determine and publish their desired allocation. If the total demand for a resource exceeds the capacity, its price is increased, and if the total demand does not exceed capacity, the price can be decreased. This section proceeds by introducing important concepts from both general equilibrium theory [8] and equity theory [23] defined specifically for air traffic flow control, in which the agents are in fact the airlines and the resources in question are airspace flow capacities.

Consider the resource allocation problem for a single resource of capacity $C \in \mathbb{R}_+$ and a fixed number $N_{\mathcal{J}}$ of airlines ($j \in \mathcal{J}$). The airlines are allocated an amount $y_j \in \mathbb{R}_+$ of the resource such that

$$\sum_{j \in \mathcal{J}} y_j \leq C$$

Suppose further that a predefined publicly available flight schedule exists, $y_j^s \in \mathbb{R}_+$, such that

$$\sum_{j \in \mathcal{J}} y_j^s > C$$

and that convex airline cost functions $J_j: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ can be expressed in terms of the deviation from the flight schedule. The single-resource allocation problem is summarized as follows:

Minimize

$$\sum_{j \in \mathcal{J}} J_j(y_j) \quad (1)$$

subject to the constraints

$$\begin{aligned} \sum_{j \in \mathcal{J}} y_j &\leq C \\ y &\geq 0 \end{aligned}$$

First, without incorporating airline cost information, it is possible to define a schedule-based allocation rule $\mathcal{A}_j(C, y^s)$ that converts the scheduled flow allocation to a feasible flow allocation for each airline proportional to the predefined schedule.

$$y_j^p = \mathcal{A}_j(C, y^s) = \frac{y_j^s}{\sum_{j \in \mathcal{J}} y_j^s} C \quad (2)$$

where y_j^p is the resulting schedule-based flow allocation for airline j , and $\mathcal{A}(C, y^s)$ is the allocation rule for all airlines. The schedule-based allocation rule satisfies two interesting properties: namely, impartiality and collusion-proofness.

An allocation rule $\mathcal{A}(C, y^s)$ is *impartial* if for any permutation π on \mathcal{J} , the resulting allocation is permuted identically; that is, $\mathcal{A}(C, \pi(y^s)) = \pi(\mathcal{A}(C, y^s))$.

An impartial allocation rule is *collusion-proof* if for any schedule, no subset of airlines can increase their allocation by combining their claims.

In fact, proportional allocation is the unique allocation rule that satisfies these two properties [23]. Both the strength and weakness of

the proportional allocation rule is that it is independent of airline cost functions and private information. Because the schedule is public, impartiality and collusion-proofness ensure that the allocation is indeed fair and seems so to all airlines. It is, however, highly probable that scheduled flights can vary dramatically in the costs that are incurred by delays, and therefore, any allocation mechanism that relies on schedule information alone cannot ensure that another solution does not exist that all airlines prefer. By incorporating airline preferences, it is possible to define a competitive equilibrium as a flow allocation $y^{ce} \in \mathbb{R}_+^{N_{\mathcal{J}}}$ and a price $\mu \in \mathbb{R}_+$, for which the allocation is preferred by all airlines over any other allocation they could purchase at the current market price.

The pair (y^{ce}, μ) is a *competitive equilibrium* if y_j^{ce} minimizes each airline's cost function

$$y_j^{ce} = \inf_{y_j} \{J_j(y_j) + \mu y_j\}$$

and the price μ satisfies

$$\mu \left(\sum_{j \in \mathcal{J}} y_j - C \right) = 0$$

By enforcing that all airline cost functions are to be defined in a common unit of measure (USD, for example), an efficient solution can be defined.

An *efficient* solution y^* to the resource allocation problem is one that minimizes aggregate cost

$$y^* = \inf_y \sum_{j \in \mathcal{J}} \{J_j(y_j)\}$$

such that

$$\begin{aligned} \sum_{j \in \mathcal{J}} y_j &\leq C \\ y &\geq 0 \end{aligned}$$

It is important to note that because the schedule-based allocation does not depend on preference information, the allocation cannot be guaranteed to be efficient and is therefore inherently problematic. The potential exists for discussion among airlines to result in dissatisfaction with an inefficient process, as was observed before modifications to the CDM process [3]. The utility of the concept of proportional allocation arises from its use as a normative standard that defines a baseline from which to investigate cost-based allocations. Given the proportional flow allocation y^p , acceptable flow allocations are those that are no worse than y^p for any airline.

An allocation rule is *acceptable* if no airline's cost is increased above its cost resulting from a schedule-based proportional allocation.

The fundamental theorems of welfare economics assert that the competitive equilibrium of the resource allocation problem results in an efficient solution and that such an equilibrium exists and is unique under certain assumptions [8,24].

Theorem 1 (first theorem of welfare economics): Assume that each airline's cost function, $J_j(y_j)$ is convex, nonincreasing, and continuously differentiable. Then any competitive equilibrium defined by a price μ and a resource allocation vector y is an efficient solution to the resource allocation problem.

Theorem 2 (second theorem of welfare economics): Assume that each airline's cost function $J_j(y_j)$ is convex, nonincreasing, and continuously differentiable. Then for any allocation y , there exists a price μ such that (y, μ) defines a competitive equilibrium. If the cost functions are also strictly decreasing, then the equilibrium price vector and efficient allocation are unique.

It is possible to extend these results to the resource allocation problem with multiple resources, as is required for air traffic flow control by extending the definitions of preference and monotonicity to the multivariate domain. *Weak preference* between two vectors $x, \tilde{x} \in \mathbb{R}^N$, is defined as

$$x \geq \tilde{x} \Leftrightarrow x_n \geq \tilde{x}_n \quad \forall n \in N$$

and *strong preference* is defined as

$$x > \tilde{x} \Leftrightarrow x_n \geq \tilde{x}_n \quad \forall n \in N$$

and $\exists i \in N$ such that $x_i > \tilde{x}_i$. *Weak nonincreasing monotonicity* is then defined for a function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ as $x \geq \tilde{x} \Rightarrow f(x) \leq f(\tilde{x})$, and the function f is referred to simply as a nonincreasing function.

It is important to note that the key assumptions used in guaranteeing efficiency and existence of a competitive equilibrium are the convexity of the allocation domain as well as the convexity and nonincreasing nature of the airline cost functions in all resources over the allocation domain. For air traffic flow modeling, however, these assumptions represent a restriction on the possible models that can be used. Therefore, this paper presents two models: the discrete path flow model, which satisfies the assumptions of the first and second theorems of welfare economics by design, and the continuous link flow model, which includes nonlinear PDE flow constraints that violate the assumption of convexity of the allocation domain but offer a more precise representation of traffic flow over the network, as well as additional control inputs for velocity regulation.

III. Discrete Path Flow Model

This section proceeds by defining an air traffic flow model and airline models that satisfy the assumptions of general equilibrium theory and for which a distributed allocation mechanism is derived that is guaranteed to converge to the competitive equilibrium. By defining payments relative to an equitable but inefficient proportional allocation, a guarantee is provided that each airline will always find the market outcome acceptable; that is, they will weakly prefer the market-mechanism outcome to the proportional allocation. This theoretical foundation represents a formal justification for the application of market mechanisms to the air traffic control problem.

A. Network Definition

Consider an air traffic network defined by fixed sets of nodes $n \in \{1, \dots, N\} = \mathcal{N}$ and links $i \in \{1, \dots, I\} = \mathcal{I}$. The node incidence matrix $A^N \in \{0, 1\}^{N \times N}$ captures the connectivity of the network, with $A^N(a, b) = 1$ if node a flows to node b and 0 otherwise.[§] Similarly, a link connectivity matrix $A^I \in \{0, 1\}^{I \times I}$ is defined with $A^I(a, b) = 1$ if link a flows into link b and 0 otherwise.

A subset of nodes $S \subseteq \mathcal{N}$ is defined as the set of sources from which flow can originate, and $Z \subseteq \mathcal{N}$ is defined as the set of sinks from which flow exits the network. Each link $i \in \mathcal{I}$ has a set of links \mathcal{M}_i , for which the entire flow merges into link i , and a set of links \mathcal{D}_i , for which the members each receive a portion of flow from link i . Let $\mathcal{M} = \bigcup_{i \in \mathcal{I}} \mathcal{M}_i$ represent the set of all merging links and define \mathcal{D} similarly. Then, for each diverging link $d \in \mathcal{D}$, there exists a nonempty set \mathcal{F}_d of inflow links for which a portion of the flow is able to flow into link d .

A simple air traffic network with five links is presented in Fig. 1 to illustrate the various components of the definition, in which allowable flows are indicated by dashed and dotted lines. The source set is $S = \{a, b, c\}$ and the sink set is $Z = \{e, f\}$. For this network, there are no merge sets, because all flow is split between multiple links. The divide sets $\mathcal{D}_1 = \mathcal{D}_2 = \mathcal{D}_3 = \{4, 5\}$ and inflow sets $\mathcal{F}_4 = \mathcal{F}_5 = \{1, 2, 3\}$ capture the properties that flow from links 1 and 2, and link 3 can be split into links 4 and 5.

B. Discrete Path Flow Model

Instead of allowing complete authority over flow velocity and route selection, the path flow model restricts the flow to predefined paths with fixed velocity profiles. The result is that continuity is conserved by the path definition and that capacity constraints need only be ensured upon entry to a link. The formulation is further




Fig. 1 A simple air traffic network with five links.

simplified by discretizing time with time step δt , as $t \in \{0, \delta t, \dots, T\} = \mathcal{T}$, where $T \in \mathbb{R}_+$ represents the finite time horizon of interest and $N_T = |\mathcal{T}|$. This allows for the definition of a finite set of available link/time pairs, referred to as en route resources, $r = (i, t) \in \mathcal{R} = \mathcal{I} \times \mathcal{T}$, with $R = |\mathcal{R}|$ being the total number of en route resources. Each en route resource has associated with it a capacity limit, summarized by an en route resource capacity limit vector $C^e \in \mathbb{R}_+^R$.

In the discrete path flow model, each flight must consume a sequence of en route resources as it travels over the network. Flow over the network must travel along a path p , which is defined by the sequence of resources it consumes:

$$p = (r^1, \dots, r^{N_p}) \quad (3)$$

where $N_p \in \mathbb{N}$ is the length of the path. Implicit in the definition of the paths is the speed that flights travel along each link, which is fixed for all flights. Flow can either be delayed by shifting it to a path over the same links with a later departure time or be rerouted along a different set of links. Furthermore, each path is defined for a specific airline,[¶] so that $p \in \mathcal{P}_j$, $\mathcal{P} = \bigcup_{j \in \mathcal{J}} \mathcal{P}_j$, and $P = |\mathcal{P}|$. Each en route resource is assigned to the path at the time of entry onto a link and therefore encapsulates the dynamics of the network flow in a matrix representation. The path-resource matrix $A^e \in \{0, 1\}^{R \times P}$ is defined as

$$A^e(r, p) = \begin{cases} 1 & \text{path } p \text{ consumes resource } r \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

Traffic flow along each path can now be allocated such that capacity limitations are observed. Define $y \in \mathbb{R}_+^P$ as the flow allocation along all paths, then en route capacity limit constraints are simply

$$A^e y \leq C^e \quad (5)$$

It should be noted that this model restricts control actions to the flow equivalent of ground delays and flight reroutes, although it does not include en route delay due to the fixed velocity requirement along each link. This restriction is relaxed in Sec. VI as part of the continuous link flow model.

Similarly, airport departure and arrival flow can be restricted to satisfy runway capacity limitations. Because multiple paths can originate or be destined for each airport node, arrival and departure flow can be determined by summing over the corresponding path flows. Define the matrices $A^o \in \{0, 1\}^{(|S|N_T) \times P}$ and $A^d \in \{0, 1\}^{(|Z|N_T) \times P}$ to map any path flow allocation y to an origin outflow $y^o \in \mathbb{R}^{|S|N_T}$ and destination inflow $y^d \in \mathbb{R}^{|Z|N_T}$. Let

$$A^o(s, p) = \begin{cases} 1 & \text{path } p \text{ starts at origin/time pair } s \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

and A^d is defined similarly. Then airport capacities are constrained by

$$A^o y \leq C^o, \quad A^d y \leq C^d \quad (7)$$

where $C^o \in \mathbb{R}^{|S|N_T}$ and $C^d \in \mathbb{R}^{|Z|N_T}$ are the limits on safe airport operation.

C. Discrete Airline Model

To react strategically to unpredictable phenomenon such as weather, it is necessary to define a preference relation over possible outcomes for each airline. The airline preferences are modeled by minimizing the quadratic deviation from the cumulative scheduled

[§]The notation $A^N(a, b)$ refers to the (a, b) element of matrix A^N .

[¶]For the path flow model, the subscript j shall be used to denote the components of a set, vector, or matrix that corresponds to the j th airline.

destination inflow. This cost metric was selected to penalize the largest contributor to airline delay costs: namely, aircraft arrival delay, which is captured in the discrete path flow model through the cumulative deviation from the scheduled arrival flow. The quadratic nature of the cost function penalizes large deviations from the flight schedule more heavily than small deviations, because longer delays tend to impact subsequent scheduled flights more severely, driving a domino effect of cost escalation. Although deviations from the planned routes, departure delays, and en route velocity changes can impact airline costs, it is reasonable to assume that arrival delays are the most prominent concern. However, concave nondecreasing cost functions can easily be designed that also penalize additional quantities such as cumulative deviation from the scheduled departure flow or deviation from scheduled path flows.

To simplify notation, let $Y^d \in \mathbb{R}_+^{V \times V}$ be the cumulative destination outflow for each airport. Cumulative arrival inflow can be expressed as $Y^d = \Gamma A^d y$, where $\Gamma \in \{0, 1\}^{V \times V}$ is a block diagonal matrix with lower triangular blocks of dimension N_T . Let y^{sd} be the scheduled arrival flow and let Y^{sd} be defined equivalently to Y^d . A representative airline cost function can be defined as

$$J_j(y_j) = \left\| \Omega_j \max(0, Y_j^{\text{sd}} - Y_j^d) \right\|_2^2 \quad (8)$$

where $\Omega_j \in \mathbb{R}_+^{V \times V}$ is a diagonal weighting matrix that captures each airline's preferences over arrival flows at each time epoch. This cost formulation is convex and nonincreasing, because it is a positive linear combination of convex nonincreasing functions in the path flow variables y_j .

D. Central Optimization

It is now possible to define a centralized optimization program with complete information of airline preferences. Let A and C be defined as

$$A = \begin{bmatrix} A^e \\ A^o \\ A^d \end{bmatrix}, \quad C = \begin{bmatrix} C^e \\ C^o \\ C^d \end{bmatrix} \quad (9)$$

Note that the matrix A can be divided into $N_{\mathcal{J}}$ submatrices $A = [A_1 | \dots | A_{N_{\mathcal{J}}}]$, which map airline path flow allocations y_j to the resources consumed. The central path flow allocation program is therefore analogous to the single resource allocation program, problem (1):

Minimize

$$\sum_{j \in \mathcal{J}} J_j(y_j) \quad (10)$$

subject to the constraints

$$\begin{aligned} Ay &\leq C \\ y &\geq 0 \end{aligned}$$

for which the solution y^c represents the ideal solution to the multi-airline resource allocation problem if airline preferences were public information. The convex formulation ensures that for problems of reasonable size, the optimization can be performed conveniently using standard convex optimization tools such as CVX [25] and SeDuMi [26] in the Matlab [27] environment.

IV. Market Mechanism for the Discrete Path Flow Model

The problem of managing air traffic flow is essentially one of allocating scarce resources to the airlines, each of which has both prior claims to the resources based on the preapproved flight schedule as well as private information concerning their preferences over allocations. This section develops a market mechanism for the central path flow allocation program by applying dual decomposition and defining both local airline path flow optimizations and a central pricing-update technique.

A. Schedule-Based Allocation

Before developing a market for air traffic resources, a solution is developed that does not require airline preference information but does satisfy capacity constraints on the network. Scheduled path flows are scaled back proportionally relative to the worst capacity violation on any of the resources they consume, resulting in a feasible schedule-based proportional path flow allocation y^p , as defined in Eq. (2), that does not allow rerouting on unconstrained routes. This serves as the baseline solution from which improvement can be achieved by applying a market mechanism.

B. Discrete Market Mechanism

The market mechanism is developed in three steps. First, the dual function for the central path flow allocation program is formed with Lagrange multipliers $\mu \in \mathbb{R}_+^{N_c}$ that are referred to as resource prices, one for each path or airport constraint:

$$g(\mu) = \inf_{y \geq 0} \left\{ \sum_{j \in \mathcal{J}} J_j(y_j) + \mu^\top (Ay - C) \right\} \quad (11)$$

where N_c is the total number of central network constraints defined in Eq. (9). The dual problem to the central path flow allocation problem is

$$\sup_{\mu \geq 0} \{g(\mu)\} \quad (12)$$

Next, the dual problem is subdivided into dual-function evaluations given fixed resource prices and a resource price optimization. The dual problem can be rewritten as

$$\sup_{\mu \geq 0} \{g(\mu)\} = \sup_{\mu \geq 0} \left\{ \sum_{j \in \mathcal{J}} \inf_{y_j \geq 0} [J_j(y_j) + \mu^\top A_j y_j] - \mu^\top C \right\} \quad (13)$$

Dual-function evaluation can now be decomposed into $N_{\mathcal{J}}$ path flow subproblems,

$$\inf_{y_j \geq 0} [J_j(y_j) + \mu^\top A_j y_j] \quad (14)$$

and the dual problem is solved by maximizing over resource prices μ . Because the dual function is the pointwise infimum of a family of convex functions, it cannot be assumed to be continuously differentiable. Hence, nonsmooth optimization techniques such as the subgradient method [28] must be applied to solve the dual problem. Let k index the algorithm iterations. Given the current flow requests y^k from all airlines, a subgradient $\partial g / \partial \mu^k \in \mathbb{R}^{N_c}$ provides an ascent direction for the dual function:

$$\frac{\partial g}{\partial \mu^k} = Ay^k - C \quad (15)$$

Maximization of the dual function is performed iteratively using a predetermined square-summable but not summable step size, because convergence results have been established for such an approach [28]. The pricing update assumes the form

$$\alpha^k = \frac{\alpha_0}{k} \quad (16)$$

$$\mu^{k+1} = \max \left(0, \mu^k + \alpha^k \frac{\partial g}{\partial \mu^k} \right) \quad (17)$$

where $\alpha^k \in \mathbb{R}_+$ is the k th iteration step size.

C. Airline Market Optimization

Each airline is interested in minimizing the total cost it incurs, trading off the cost of purchasing additional resources with its preferences over flow allocations. A resource payment $v_j \in \mathbb{R}$ is defined for each airline relative to the baseline proportional allocation defined in Eq. (2),

Algorithm 1 Discrete path flow model market mechanism

-
- 1) Set prices to zero, $\mu = 0$.
 - 2) Repeat.
 - 3) Central authority publishes prices μ^k .
 - 4) Airlines solve airline path flow allocation program [problem (19)].
 - 5) Airlines submit flow requests y_j^k .
 - 6) Central authority updates prices $\mu^{k+1} = \max[0, \mu^k + \alpha^k(\partial g / \partial \mu^k)]$.
 - 7) Until $\|\mu^{k+1} - \mu^k\| < \epsilon$ or $k = k^{\max}$.
-

$$v_j = \mu^\top A_j (y_j - y_j^p) \quad (18)$$

which charges each airline for any additional resources consumed and rewards the airline for any resources relinquished.

The airline path flow allocation program can now be defined, which precisely solves the dual-function-evaluation subproblems defined in Eq. (14):

Minimize

$$J_j(y_j) + v_j \quad (19)$$

subject to the constraint

$$y_j \geq 0$$

The airline path flow allocation represents the market-based tradeoff that each airline is faced with and can be interpreted as follows: Given the current resource prices, minimize the total cost of deviating from the scheduled flow and purchasing the resources consumed by the flow. With linear constraints and a convex cost, the local optimization problem is convex.

If desired, airlines can add additional constraints to their airline path flow allocation program not captured by the central path flow allocation program constraints. For example, airlines may wish to ensure that aircraft are not required to take off ahead of their scheduled departure times. This can be encoded by adding the constraints that for each airline, the cumulative flow leaving an airport node does not exceed the cumulative scheduled flow at any point in time. Similarly, additional constraints can be included in the problem formulation to limit early arrival traffic or to limit the difference between arrival and departure flow at an airport to satisfy gate capacity constraints.

The complete market-based distributed algorithm is defined in Algorithm 1, and the resulting flow allocation shall be referred to as y^m .

D. Acceptability

The advantage of defining resource payments relative to the equitable public baseline defined by proportional allocation is that airlines can only improve their outcome by participating in the market. The following proposition provides a guarantee that inclusion in the market is always in the airlines' best interest.

Proposition 1: Given the payment rule $v_j = \mu^\top A_j (y_j - y_j^p)$ each airline will always weakly prefer the market-based flow allocation to the schedule-based flow allocation; that is,

$$J_j(y_j^m) + v_j \leq J_j(y_j^p) \quad (20)$$

Proof: Assuming price-taking airlines, this result follows directly from the definition of the airline cost function:

$$\begin{aligned} J_j(y_j^m) + v_j &= J_j(y_j^m) + \mu^\top A_j (y_j^m - y_j^p) \\ &= \inf_{y_j \geq 0} \{J_j(y_j) + \mu^\top A_j (y_j - y_j^p)\} \\ &\leq J_j(y_j^p) + \mu^\top A_j (y_j^p - y_j^p) = J_j(y_j^p) \end{aligned}$$

□

For airlines that can exchange funds and measure costs in dollars, Proposition 1 ensures that a resource pricing solution exists such that

all airlines find the solution acceptable and that no subgroup is unhappy with the resulting allocation. Furthermore, because the market is closed, airlines with less expensive flights stand to reduce costs considerably over a schedule-based solution, because they will be compensated for flight delays that are caused by larger, more expensive, flights getting priority. By selling their allocated resources to more expensive flights in restricted periods, the airspace can be more efficiently used to reduce aggregate costs incurred by all airlines. Conveniently, the market approach is inherently distributed, because each airline is responsible for optimizing its own costs and determining the additional quantity of each resource to purchase at market prices above its schedule-based allocation. The role of the FAA transitions from that of central decision-maker to one of determining capacity constraints along links and setting market prices that equalize supply and demand.

V. Discrete Path Flow Model Simulation: Large Network

The simplicity of the discrete path flow model makes simulation of significant portions of the U.S. national airspace possible on a standard desktop computer. This section presents the implementation of a representative scenario, and the issue of problem complexity and computational tractability is addressed in Sec. VII.E.

A. Scenario

A simulation is presented for flow departing from 17 airports within approximately three hours of Chicago O'Hare. The simulation spans a 3-h time horizon and uses a 5-min time step. The nodes of the network are based on actual airport and airway intersection locations, links were generated artificially for convenience only, and a primary shortest path and secondary alternate path were generated over the network for each OD pair.

Four different airlines are considered, with different schedules and route preferences as outlined in Table 1, and are intended to capture some of the current differences among U.S. airlines. Mainline carriers tend to have peaked schedules that align scheduled arrivals and departures at their hub airports to attract business travelers who are most concerned with the timing of connecting flights. Low-cost carriers, on the other hand, keep "flat" schedules, with departures and arrivals scheduled consistently throughout the day, to maximize gate and personnel use and to keep costs down. The airlines also differ significantly in the amount of flow scheduled, which is representative of the different flow requirements caused by large mainline airlines versus smaller regional carriers. For this simulation, peaked schedules have tripled scheduled flow for the second quarter of the simulation, high costs are on average three times the low costs, and low traffic volume is 30% of high traffic volume.

En route capacity limitations are uniformly set to 90% of the peak scheduled flow along the busiest link, such that flow could not quite arrive entirely on schedule in good weather. Inclement weather is included as a polygon with fixed size and a known velocity that reasonably approximates standard weather pattern movements. Capacity is assumed to be restricted by 50% for all flows exiting a node inside the weather polygon, as is the allowable destination airport flow.

B. Results

Figure 2 shows the resulting aggregate flow over the network for all airlines at a sequence of times. The weather polygon and affected nodes are also shown, and link width is used to represent increasing

Table 1 Table of airline schedules and flow preferences

	Schedule type	Traffic volume	Costs
Airline 1	Peaked	High	High
Airline 2	Flat	High	Low
Airline 3	Peaked	Low	High
Airline 4	Flat	Low	Low

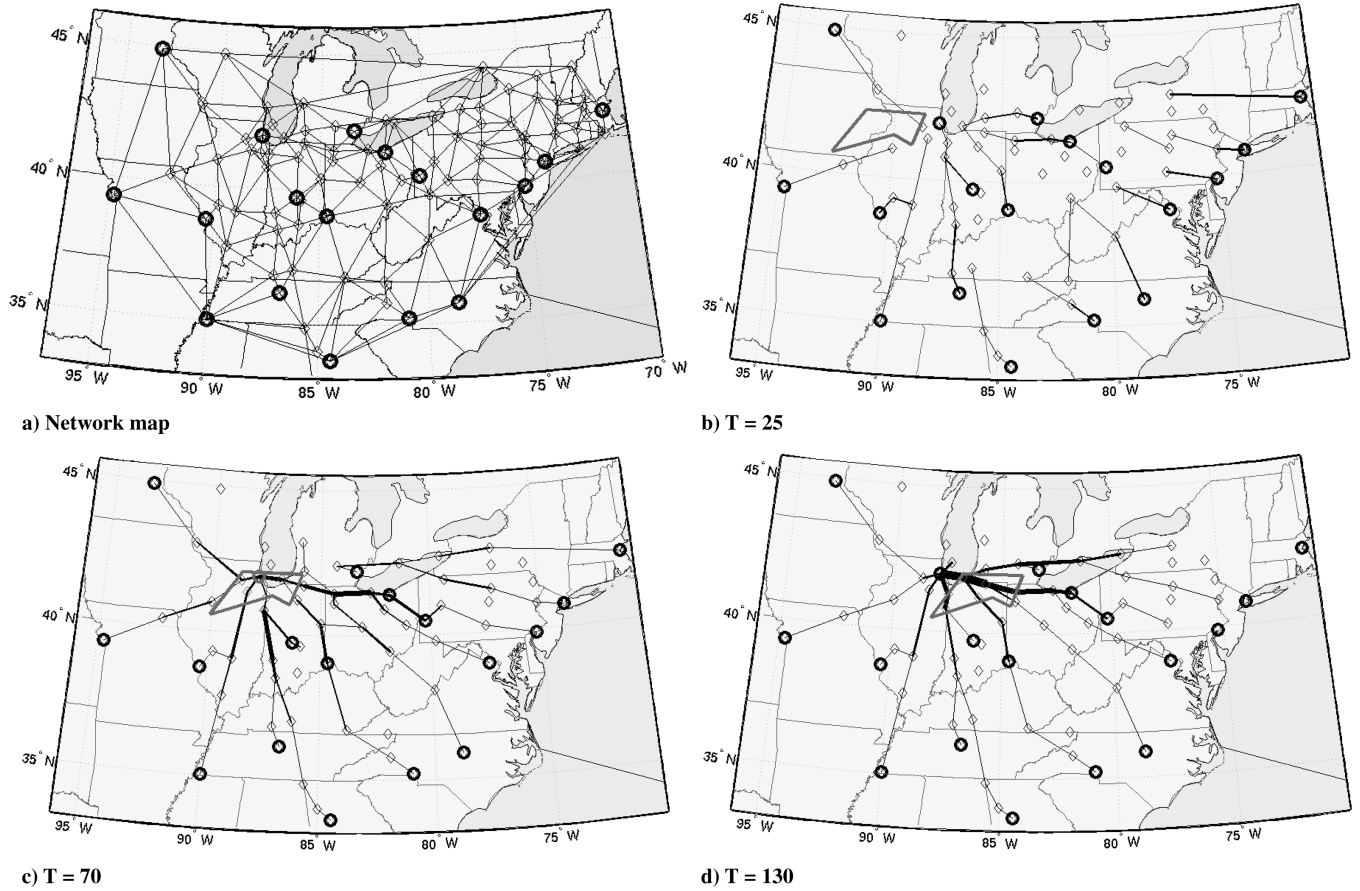


Fig. 2 Aggregate flow of traffic to Chicago O'Hare airport for all airlines; airports marked are with black circles, and all airway intersection nodes are depicted with diamonds.

flow volume. The weather-disruption polygon and affected nodes are represented in dark gray.

Figure 2b shows the relatively low-demand starting period of the simulation, in which most flows are taking the shortest path routes to the destination and none of the flow is close to capacity limits. It is possible to discern the effect of capacity reduction due to weather by comparing Figs. 2c and 2d, in which flow into Chicago is significantly reduced until the weather passes and capacity is restored. In Fig. 2d, the high demand on links near the destination reflects the convergence of many recovery flows on the airport.

The market-based flow allocation algorithm performance can be observed in Fig. 3, in which the difference between the primal and dual costs, or duality gap, is plotted. The slow rate of convergence of the duality gap is a direct result of the dual-pricing mechanism, which




Fig. 3 Duality gap at each iteration of the market-based flow allocation with discrete path flow model, as a percentage of the optimal cost.

relies on the subgradient method and can only be guaranteed to converge in an unbounded number of steps [28].

Figure 4 presents the arrival traffic pricing throughout the simulation, as well as the resulting flow and constraint satisfaction. From this figure, one can see the flow at the destination airport being reduced during the weather disturbance and then recovering once the disturbance has passed. Note that the price vectors for the true solution computed centrally and the market-based algorithm are quite similar, as are the resulting flow allocations.




Fig. 4 Optimal arrival flow at the destination airport for the large network scenario; airport pricing for market-based solution y^m (top) and arrival flow quantities (bottom); scheduled arrival traffic (dashed), airport capacity (solid), and market-based solution (\times).




Fig. 5 Deviation from desired cumulative arrival flow: schedule-based (solid line) vs market-based (dashed line) for each airline.

The deviation from the desired cumulative schedule is plotted for each airline in Fig. 5 for the market-based allocations. The high-cost airlines (1 and 3) have reduced their deviations in the market allocation, whereas the low-cost airlines (2 and 4) have increased their deviation. The biggest difference occurs when airport capacity becomes available after the weather passes; the high-cost airlines immediately begin to recover, whereas the low-cost airlines continue to increase their cumulative deviation for a period of time after the most severe restrictions have passed.

Finally, a comparison of schedule-based and market-based flow allocation costs is presented in Table 2, in which columns represent schedule-based costs, market-based costs, the actual delay cost incurred, the resulting payments made by the airlines, and the improvement noted by each airline in using the market-based allocation. The market-based costs are preferable to the schedule-based costs for all airlines, as noted by positive values for the improvement. Note that payments for the high-cost airlines are both positive, but they are negative for the low-cost airlines, indicating that high-cost airlines are reducing their delays by paying for more resources, whereas low-cost airlines are taking on more delays but are being compensated for it.

VI. Continuous-Time Air Traffic Model

Although the discrete path flow model conveniently captures many of the critical elements of the air traffic flow control problem in a convex formulation, it is limited in two distinct ways. First, the flow velocity is held fixed even though controllers can directly control traffic flow speed through miles-in-trail restrictions. Second, time is discretized before the formulation of the optimization program, and the resulting discretization error of link travel times, takeoff and landing times, etc., can accumulate and become significant. To improve on these limitations, a continuous-time formulation of the air traffic flow problem is now presented.

A. Continuous Link Flow Model

The network definition of Sec. III.A is maintained for the continuous link flow model. The definition for each link i is extended

Table 2 Cost comparison and resource payments for all airlines

	Schedule	Market	Delay cost	Payment	Improvement
Airline 1	2422.0	1814.9	635.4	1179.5	607.1
Airline 2	989.1	848.6	1603.8	-755.2	140.5
Airline 3	2061.8	1667.1	744.3	922.8	394.8
Airline 4	550.9	435.2	938.5	-503.3	115.7
Total	6023.8	3922.0	843.8	4765.8	1258.1

to include a length L_i , with the coordinate $x_i \in [0, L_i]$ used to refer to any point on link i . For every destination node $z \in \mathcal{Z}$ along each link in the network, each airline has an associated flow $q_{i,j}^z: [0, L_i] \times [0, T] \rightarrow \mathbb{R}_+$ of aircraft with density $\rho_{i,j}^z: [0, L_i] \times [0, T] \rightarrow \mathbb{R}_+$ and velocity, $v_{i,j}^z: [0, L_i] \times [0, T] \rightarrow \mathbb{R}_+$. The flow, density, and velocity obey**

$$q_{i,j}(x_i, t) = \rho_{i,j}(x_i, t)v_{i,j}(x_i, t) \quad (21)$$

for all airlines, at all points in space and time. For any diverging link $d \in \mathcal{D}$, the flow split parameter $\beta_{i,j,d}: [0, T] \rightarrow [0, 1]$ describes the j th airline's portion of the flow that travels from link i into link d . Continuity requires that the flow split among diverging links sums to one for each airline; namely, that

$$\sum_{d \in \mathcal{D}_i} \beta_{i,j,d}(t) = 1$$

Each source has associated with it aircraft inflow rates per airline of $q_{i,j}^{\text{in}}: [0, T] \rightarrow \mathbb{R}_+$, and each link has an initial link density $\rho_{i,j}^0: [0, L_i] \rightarrow \mathbb{R}_+$.

Continuity results in the following PDEs:

$$\frac{\partial \rho_{i,j}}{\partial t} + \frac{\partial \rho_{i,j} v_{i,j}}{\partial x_i} = 0 \quad (i, j) \in \mathcal{I} \times \mathcal{J} \quad (22)$$

with initial and boundary conditions for each link and each airline summarized as

$$\begin{aligned} \rho_{i,j}(x_i, 0) &= \rho_{i,j}^0(x_i) \quad (i, j) \in \mathcal{I} \times \mathcal{J} \\ q_{i,j}(0, t) &= q_{i,j}^{\text{in}}(t) + \sum_{k \in \mathcal{M}_i} q_{k,j}(L_k, t) + \sum_{f \in \mathcal{F}_i} \beta_{f,j,i}(t) q_{f,j}(L_f, t) \end{aligned} \quad (i, j) \in \mathcal{I} \times \mathcal{J} \quad (23)$$

To capture vehicle and airspace capacity limitations, constraints must be imposed on the flow over the network. Aircraft from each airline must obey the same minimum and maximum velocity limits $v_i: [0, L_i] \rightarrow \mathbb{R}_+$ and $\bar{v}_i: [0, L_i] \rightarrow \mathbb{R}_+$, defined at each point in the network. A maximum density limit $\bar{\rho}_i: [0, L_i] \times [0, T] \rightarrow \mathbb{R}_+$ is imposed along each link to ensure safe use of the available airspace.**

$$\sum_{j \in \mathcal{J}} \rho_{i,j}(x_i, t) \leq \bar{\rho}_i(x_i, t) \quad i \in \mathcal{I} \quad (24)$$

B. Continuous Airline Model

As before, it is assumed that the airlines have agreed to a predefined schedule which can be represented as the desired number of aircraft arrivals per destination airport, $\eta_{i,j}^s: [0, T] \rightarrow \mathbb{R}_+$ ($i \in \mathcal{Z}$) as a function of time. Actual arrivals $\eta_{i,j}(t)$ can then be measured against the desired schedule and control decisions can be optimized to reduce overall delay for each airline. A characteristic cost function $J_j \in \mathbb{R}_+$ ($j \in \mathcal{J}$) can be defined as

$$J_j = -\sum_{i \in \mathcal{Z}} \int_0^T \omega_{i,j}(t) [\eta_{i,j}^s(t) - \eta_{i,j}(t)]^2 dt \quad (25)$$

where the function $\omega_{i,j}: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ defines a time-varying delay cost for each airline. The number of arriving aircraft can be written as an integral over the arrival flow:

$$J_j = -\sum_{i \in \mathcal{Z}} \int_0^T \omega_{i,j}(t) \left(\int_0^t q_{i,j}^s(L_i, s) - q_{i,j}(L_i, s) ds \right)^2 dt \quad (26)$$

Each airline therefore seeks to minimize the squared deviation in cumulative flow arriving at a destination airport with respect to the scheduled flow, in a continuous manner analogous to the discrete

**The dependence on z of all network variables is dropped to ease notation.

††Note that for multiple sinks, the flow density must be summed over all sinks as well as over all airlines.

path flow model airline cost. Once again, the quadratic nature of the cost function penalizes large deviations from the flight schedule more heavily than small deviations.

C. Central Optimization

By defining the continuous link flow model, three distinct flow control methods can be included in the optimization formulation, mirroring the options available to the FAA in practice. The GDP approach is embodied in control of inflow $q_{i,j}^{\text{in}}$, miles-in-trail restrictions are similar to controlling link velocity profiles $v_{i,j}$, and the flow splitting parameters $\beta_{i,j,d}$ allow for flight rerouting along alternate links. The multi-airline optimization problem can now be formulated as the central continuous flow allocation program for the continuous link flow model:

Minimize

$$\sum_{j \in \mathcal{J}} J_j \quad (27)$$

subject to Eqs. (22) and (23) and to the constraints

$$\begin{aligned} v_i &\leq v_{i,j} \leq \bar{v}_i & (i, j) \in \mathcal{I} \times \mathcal{J} \\ \sum_{j \in \mathcal{J}} \rho_{i,j} &\leq \bar{\rho}_i & i \in \mathcal{I} \\ \sum_{d \in \mathcal{D}_i} \beta_{i,j,d} &= 1 & (i, j) \in \mathcal{I} \times \mathcal{J} \end{aligned}$$

with the optimal solution denoted as $\rho_{i,j}^c$, $v_{i,j}^c$, $\beta_{i,j,d}^c$, and $q_j^{\text{in},c}$ for all links and airlines.

VII. Market Mechanism for the Continuous Link Flow Model

For the continuous link flow model, a market-based allocation is derived in three steps. The schedule-based allocation is defined as a baseline solution from which to assess market outcomes. Then the dual problem for the centralized optimization program is formed and decomposed into portions that are of interest to each airline and a portion that requires centralized coordination by the FAA. Finally, the adjoint method [29] is applied as a solution technique for the optimization problem to be performed by the airlines.

A. Schedule-Based Allocation

Similar to the discrete case, proportional allocation is used to define a baseline solution for the continuous link flow model. The inflow at each origin airport that is bound for a constrained resource is scaled back proportionally for all airlines until the resulting flow is feasible at each point on the network. The schedule-based proportional allocation density $\rho_{i,j}^p$ is once again used to define payments that result in acceptability of the market outcome.

B. Continuous Market Mechanism

Dual decomposition [30] is performed on the centralized aggregate-cost-minimization form of problem (27). Inspection of problem (27) reveals that only the density constraints couple the airline solutions together. The remaining constraints, including PDE and velocity constraints, can be considered airline-dependent domain constraints for the airline subproblems defined in Eq. (29). The Lagrangian is formed

$$\mathcal{L}(\rho, \lambda) = \sum_{j \in \mathcal{J}} (J_j + \langle \lambda, \rho_j \rangle) - \langle \lambda, \bar{\rho} \rangle \quad (28)$$

where the dual variable

$$\lambda: \prod_{i \in \mathcal{I}} [0, L_i] \times [0, T] \rightarrow \mathbb{R}_+$$

is referred to as the price vector, and $\langle \cdot, \cdot \rangle$ represents the inner product. The dual problem can be written as

$$\sup_{\lambda \geq 0} \{g(\lambda)\} = \sup_{\lambda \geq 0} \left\{ \sum_{j \in \mathcal{J}} \inf_{\rho_j} \{J_j + \langle \lambda, \rho_j \rangle\} - \langle \lambda, \bar{\rho} \rangle \right\} \quad (29)$$

which can be interpreted as a decoupled dual problem with $N_{\mathcal{J}}$ airline-specific subproblems that depend on a central price-update portion performed by the FAA. The price update is performed using a subgradient of g :

$$\frac{\partial g}{\partial \lambda} = \sum_{j \in \mathcal{J}} \rho_j^{\text{opt}} - \bar{\rho} \quad (30)$$

where ρ_j^{opt} is the result of each airline minimization. If the total request for flow along a link exceeds the link capacity, then the price is increased, otherwise it is decreased toward zero.

The market mechanism defined for the continuous link flow model is therefore analogous to the algorithm defined for the discrete path flow model (Algorithm 1), in which λ replaces μ and ρ replaces y .

C. Airline Continuous Link Flow Optimization

With the market pricing for restricted resources determined, a payment $\gamma_j \in \mathbb{R}$ is defined once again for each airline:

$$\gamma_j = \langle \lambda, \rho_j - \rho_j^p \rangle \quad (31)$$

To ensure that the resulting allocation is acceptable, the payment charges airlines for the difference between the resources selected and those received by the airline from the proportional allocation. The airline market cost $C_j = J_j + \gamma_j \in \mathbb{R}$ is therefore defined as the sum of the delay costs and the resource payment γ_j .

The airline continuous link flow optimization problem, defined in problem (32), consists of minimizing the total cost (equal to the sum of the delay cost and the resource trade costs) subject to the PDE flow propagation dynamics. Therefore, it consists of a PDE optimal control program that can be solved efficiently using an interior point method, in which the gradient of the cost with respect to the control variables is computed via the adjoint method [29] and is defined as follows:

Minimize

$$C_j = J_j + \gamma_j$$

subject to Eqs. (22) and (23) and to the constraints

$$\begin{aligned} v_i &\leq v_{i,j} \leq \bar{v}_i & i \in \mathcal{I} \\ \sum_{d \in \mathcal{D}_i} \beta_{i,j,d} &= 1 & i \in \mathcal{I} \end{aligned} \quad (32)$$

The interior point method used to solve the preceding optimization program can be summarized in three main steps.

1) The first step consists of eliminating the equality constraints on the flow split parameters $\beta_{i,j,d}$ through a substitution of variables and of incorporating all inequality constraints as logarithmic barrier terms in an augmented cost function.

2) The second step consists of computing the gradient of the augmented cost function with respect to the control variables. Applying the calculus of variations, the gradient is computed by a first-order variation of the augmented cost function, which is a function of the variation of the control variables δq_j^{in} , δv_j , and $\delta \beta_j$ and the variation of the state $\delta \rho_j$. Because the variation of the state cannot be directly controlled, terms in the variation of the cost that depend on the variation of the state are replaced by a function of the variation of the control. This is achieved by selecting an appropriate adjoint variable (also referred to as the costate), which is governed by a PDE. Then propagation of both the state and adjoint PDEs allows for the evaluation of the gradient of the cost function.

3) The third step consists of deriving the Newton step of the optimization program. The Newton step is the minimizer of the quadratic approximation of the cost function at the current solution and is found by solving a linear system of PDEs that results from additional applications of the adjoint method.

The airline continuous link flow optimization proceeds by solving both the PDE and adjoint PDE to form the gradient and then by solving the linear system of PDEs to derive the Newton step. All control variables are then updated in the direction of the Newton step. Further details of the optimization procedure are summarized in the Appendix, and a complete description of the methodology, as well as other applications of adjoint-based optimization, are described by Raffard [31].

D. Acceptability

It is once again possible to ensure participation of the airlines in the market mechanism for the continuous link flow model. As described for the discrete path flow model in Sec. III.D, the airlines will always weakly prefer the market-based solution to the schedule-based solution. Therefore, it will always be in the best interest of the airlines to participate in the market. The proof of acceptability for the continuous link flow model identically follows that of Proposition 1.

E. Problem Size

The size of the air traffic flow control problem depends primarily on the number of OD pairs, time steps, and link resources. The U.S. NAS contains 60 major airports^{**} for which flow scheduling could be considered, which results in as many as 3540 OD pairs. A 1-min time interval over a 14-h day yields 840 time steps, and 5-min intervals result in 168 time steps. The total number of link resources can be approximated to be 2500, by assuming approximately seven links across each of 400 sectors that cover the U.S. For an average sector width of 200 miles, the total length of all links in the network can be estimated at 500,000 miles, or 25,000 link sections, with an underlying space discretization of 20 miles.

The main difference in computational complexity between the discrete path flow model and the continuous link flow model is in the local optimization, because the subgradient price-update technique behaves similarly for both. For the path flow model, the path flow variables scale with the product of time steps N_T , OD pairs $|OD|$, and alternate routes R , and the number of constraints scale with N_T and the number of links I . For the continuous link flow model, the underlying space discretization along each link dx must be included, whereas aircraft rerouting is taken into account by the flow split parameters β . Therefore, both the flow variables and constraints scale with the product of N_T , I , the link space grid dimension, and the number of destination airports. For accurate flow propagation, both space and time discretization must be sufficiently small.

For a path flow model with five alternate routes per OD pair and 5-min time intervals, the flow allocation problem could reach sizes on the order of 1 million flow variables and one-half million resource constraints. At this size, exact second-order techniques are impractical and convex program solution times in the tens of minutes for each local optimization are likely on standard computing hardware. In comparison, the continuous link flow model would result in problems of up to eight million variables per destination airport. More tractable problems can be generated by including only those resources for which congestion or weather have resulted in a desired schedule that currently strains capacity limits. Similarly, the problem size is significantly improved by limiting the planning time horizon to three hours, which is in line with the time horizon for which sufficiently accurate weather predictions are available. The simulation results presented in the next section are characteristic of problem sizes for which real-time market-based allocation is feasible.

VIII. Flow Model Comparison Simulation: Small Network

To contrast the strengths and weaknesses of the aggregate flow models described earlier, a second scenario is presented. It contains three airports over a simple five-link network with flights originating

and departing from all airports, and both the continuous link flow and discrete path flow models are implemented for this network so that the efficiency of the results may be compared.

A. Scenario

Four airlines are competing on a network composed of three airports: New York (NYC), Washington (DC), and Chicago (ORD) (see Fig. 6a). Flights departing from the east coast bound for Chicago may take a direct route, which becomes saturated due to the excess demand caused by merging flow from both NYC and DC. Alternatively, they may take a southern or northern route, which are longer and therefore will result in more delay with respect to the desired schedule. Symmetrically, flights from Chicago to either of the east coast airports may take a central path or the alternate, longer, route. The time frame for the scenario is fixed to four hours. Among the four airlines, there are two large airlines and two small airlines, with desired terminal landing rates of one aircraft every minute (airlines 1 and 2) and every 4 min (airlines 3 and 4). The airline costs per minute delay are once again of two types: one that is a constant value of 1 (airlines 2 and 4) throughout the scenario, representing small aircraft for a low-cost carrier with few connecting passengers, and one that is 2 and then doubles in cost at the halfway point in time (airlines 1 and 3), representing a main line, hub, and spoke schedule with large aircraft. The airline definitions are summarized in Table 3.

Initial flow conditions are assumed to be zero throughout the network, and the maximum density along each link is restricted to one aircraft every 8 n mile, due to ongoing bad weather. Both delays in flights and capacity constraint violations result along the central direct links in both directions.

B. Results

The continuous link flow model optimization results are summarized in Figs. 6 and 7, followed by a comparison of the deviation incurred in Fig. 8 and resulting costs for both models in Tables 4 and 5. As seen in Sec. V, convergence of the price-based subgradient method can be quite slow [28], and in practice, optimization of the scenario takes on the order of 3000 iterations to achieve near-feasible solutions for both flow models. For the continuous link flow model, the optimization results in a market-based allocation that is feasible everywhere except at specific peak times along certain portions of the central links in both directions, for which the violation reaches a maximum of 11%. Nevertheless, based on this solution, it is possible to scale down the inflows to meet the maximum density requirement without significantly altering the outcome of the market-based approach.

A sequence of three images is presented in Figs. 6b–6d, which depict the airline-specific flights along the various links at chronological instants in the simulation.^{**} The first image shows flights entering the network, and subsequent images show the evolution of the solution in favor of the more expensive airlines. The less expensive airlines end up rerouting to the longer routes, whereas the costlier flights consume all available capacity on the central link in both directions.

In addition to plotting the cumulative deviation from the schedule for the example in Fig. 8, the resulting minutes of delay for incoming flights at each point in time are displayed in Fig. 7 for the continuous link flow model. The delay per unit flow is calculated from the cumulative flow results by determining the time that elapses between when the cumulative flow amount was scheduled to arrive and when it actually does arrive. Because the rate of inflow onto the network is held constant over time in this scenario, the result of the schedule-based allocation is to assign the same delay to all airlines. Comparing the results for the default solution and the outcome of the market, it is clear that for all airlines, the inflow-restricted schedule-based delays are significantly worse than the market-optimized solution. It is also interesting to note that the large airline with high costs is able to eliminate almost all delay minutes and that the smallest airline with

^{**}A major airport is defined as one that supports greater than 0.25% of total passenger enplanements in the U.S.

^{**}Videos of ideal, schedule-based, and market-based flows can be seen at <http://glebe.stanford.edu/~stevenlw/ATFC>.

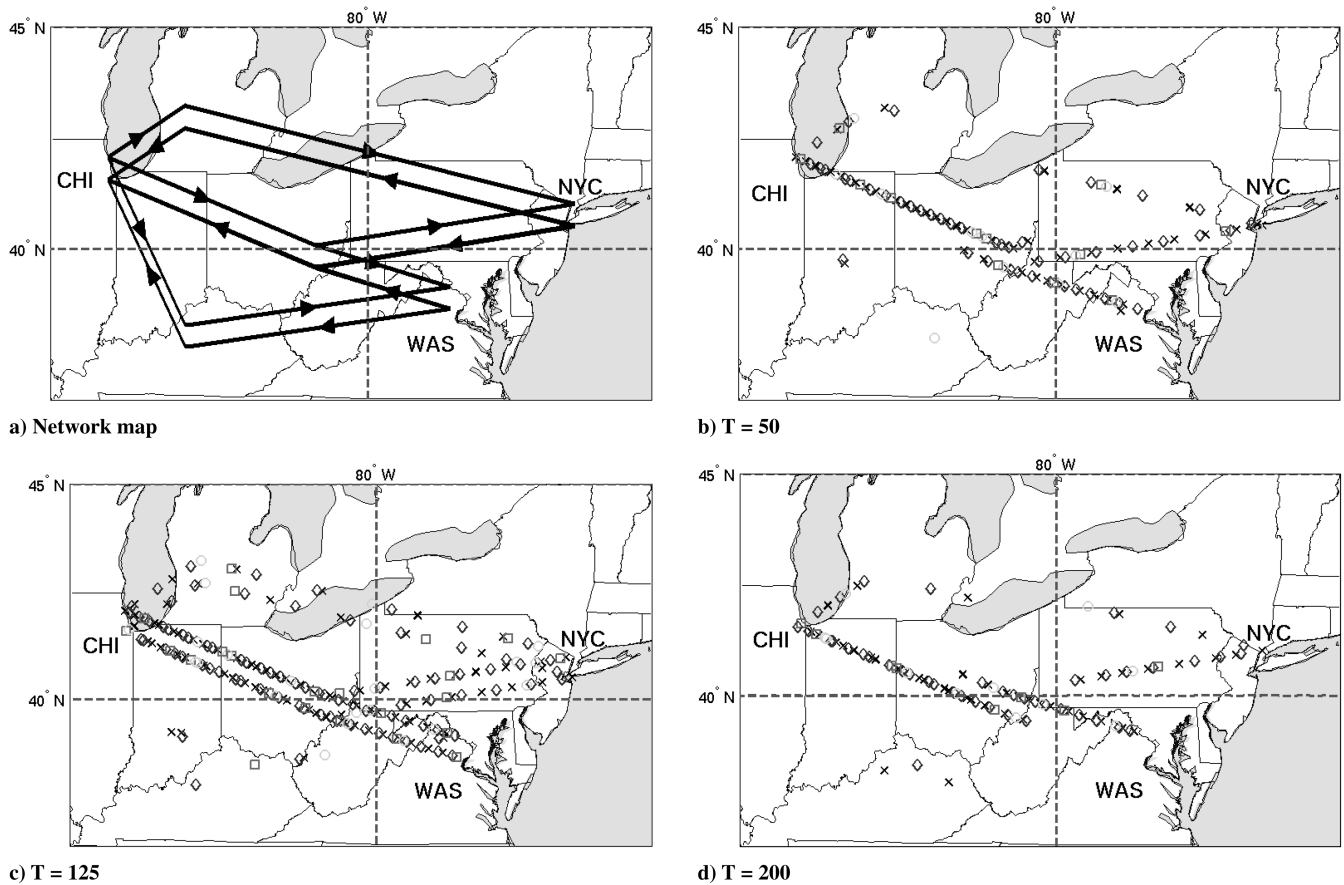


Fig. 6 Market-based optimized traffic flow patterns; individual flights from airlines 1, 2, 3, and 4 are coded \times , \diamond , $*$, and \circ , respectively.

the lowest cost achieves the worst delays in the market-based solution.

Comparison of the cumulative deviation from schedule market results for the continuous link flow model and the discrete path flow model is possible in Fig. 8, in which the improvement obtained from enabling variable velocity profiles along each link allows for reductions in the delays incurred by all airlines. For all airlines, the continuous flow optimization is able to recover almost entirely from the capacity restrictions by increasing velocity along the outside links, whereas the discrete path flow model is unable to take advantage of the ability of flight to increase speed to reduce delays. The size of the various costs and improvements is visible by comparing Tables 4 and 5.

One interesting element to note in Table 5 is the market cost for airline 4 in the discrete path flow model, which is in fact negative. This implies that the smallest, least expensive, airline is able to sell the resources allocated to them so profitably that they prefer the reduced-capacity situation to the unconstrained scenario in which their ideal cost would be zero. This type of outcome has also been observed in other scenarios for either model and relies on a large discrepancy in the valuation of resources amongst the airlines to occur.

From this example, it is possible to discern the tradeoff between using the continuous link flow model and the discrete path flow model as the basis for air traffic flow control. Because of the

additional ability to control link velocity in the continuous model, the resulting solution is able to recover from a congestion situation more efficiently by routing flow along the longer, alternate, route at maximum velocity, whereas a certain amount of delay cannot be avoided in the discrete model solution. The disadvantages of the continuous link flow model, however, are twofold. First, the computational complexity, as described in Sec. VII.E, increases significantly for the continuous link flow model, which, for the small network example, results in computation times approximately double those of the discrete path flow model. Second, implementation of the continuous variation in velocity over the network would represent a departure from current controller-pilot communication techniques and would rely on the implementation of more complete 4-D flight-plan definitions. The additional gains in efficiency possible from the continuous link flow model may therefore not be entirely achievable in practice until such advances are implemented. It is for these reasons that both models are presented, so that the benefits and drawbacks of each may be considered.

XI. Conclusions

Air traffic flow control seeks to address the large-scale problem of aircraft routing in the U.S. National Airspace System on an aggregate level in the face of unexpected disruptions to capacity, so that the control actions required for collision avoidance at the individual aircraft level remain manageable for experience air traffic sector controllers. The two models presented in this paper advance aggregate modeling for air traffic control in significant, albeit different, directions. The discrete path flow model defines a convex optimization program that allows predefined rerouting and inflow control of traffic flows and can be solved for considerable portions of the U.S. airspace in practical time. The continuous link flow model incorporates inflow, velocity, and arbitrary rerouting control and accurately propagates the flow over the network. By capturing as many of the characteristics of the underlying problem as possible in

Table 3 Summary of airline properties; desired total inflow per departure airport in aircraft per minute, and cost incurred per unit deviation from schedule

	Airline 1	Airline 2	Airline 3	Airline 4
Inflow: $q^*(0, t)$ (aircraft/min)	1	1	1/4	1/4
Flow cost: $\omega_j(t)$	2 if $t < T/2$ 4 if $t \geq T/2$	1	2 if $t < T/2$ 4 if $t \geq T/2$	1




Fig. 7 Delay as a function of time experienced at ORD for the four different airlines using the continuous link flow model. The schedule-based proportional allocation is depicted by the solid curve, and the market-based allocation is depicted by the dashed curve.

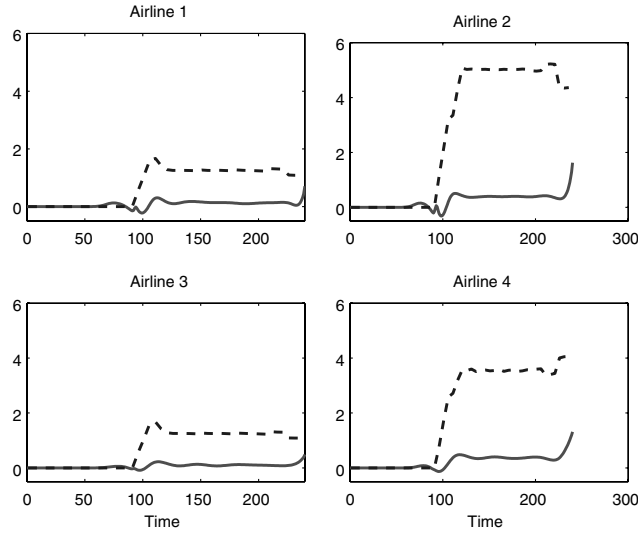


Fig. 8 Deviation in cumulative flow (in minutes) arriving at Chicago for the continuous link flow model (solid) and the discrete path flow model (dashed).

the aggregate flow model, these methods provide realistic solutions for recovering from adverse weather disruptions.

Interestingly, the NAS is relied upon by a large variety of users who incur a wide range of benefits and costs based on the daily control decisions made by airspace controllers. The assurance of safe operation is of critical importance to all users, but so too is efficient and equitable allocation of the available resources, because without it, the costs of travel could be prohibitively expensive and result in slowed economic growth for the nation. It is therefore strategically important to not only increase the traffic capacity limitations needed to ensure safe operations, but also to ensure the allocation of resources is performed in a manner that is agreeable to all users.

Table 4 Costs and resource payments for continuous link flow model

	Schedule	Market	Delay cost	Payment	Improvement
Airline 1	28,337	12.7	6.4	6.3	28,331
Airline 2	7116	10.7	5.6	5.1	7111
Airline 3	1771	8.6	4.4	4.2	1767
Airline 4	445	7.6	4.1	3.5	442
Total	37,669	40.0	21.0	19.0	37,651

Table 5 Costs and resource payments for the discrete path flow model

	Schedule	Market	Delay cost	Payment	Improvement
Airline 1	28,583	572	97	475	28,012
Airline 2	7174	292	383	-91	6882
Airline 3	1786	66	101	-36	1720
Airline 4	448	-161	186	-347	609
Total	37,991	767	1	769	37,223

To this end, this paper evaluates the potential for real-world implementation of market-based flow allocation. Because it cannot be assumed that all airlines will willingly divulge private cost information to allow the FAA to determine an efficient resource allocation, a market-based approach is presented that allows the airlines to iteratively converge to flow control strategies that avert excess demand on restricted elements of the NAS affected by severe weather. Although the resulting allocation will inevitably afford larger, more expensive, aircraft the right of way in times of restriction, it will compensate the smaller aircraft owners for the inconvenience and will unlock the maximum aggregate amount of savings from the available airspace resources, rewarding all participating airlines with reduced costs.

Appendix: Continuous Airline Link Flow Optimization

As described in Sec. VII.C, the interior point method used to solve the airline's continuous link flow optimization [problem (32)] is based on three main steps, the descriptions of which are expanded in this Appendix.

1) All local constraints can be either eliminated or incorporated into the augmented cost function as follows. The equality constraint

$$\sum_{d \in D_i} \beta_{i,d}(t) = 1$$

is incorporated by eliminating a variable from the problem formulation. If \bar{d}_i denotes the last diverging link from link $i \in \mathcal{F}$, then

$$\beta_{i,\bar{d}_i}(t) = \left[1 - \sum_{d \in C_i} \beta_{i,d}(t) \right]$$

in which $C_i = D_i \setminus \{\bar{d}_i\}$. Thus, the variable β_{i,\bar{d}_i} is eliminated and the constraints $\beta_{i,d}(t) \geq 0$ ($d \in D_i$) and

$$\sum_{d \in D_i} \beta_{i,d}(t) = 1$$

become $\beta_{i,d}(t) \geq 0$ ($d \in C_i$) and

$$\sum_{d \in C_i} \beta_{i,d}(t) \leq 1$$

The inequality constraints are handled by using a logarithmic barrier, with M being the barrier parameter. As M approaches ∞ , problem (32) becomes equivalent to the unconstrained airline continuous link flow optimization defined in problem (A1):

Minimize

$$\begin{aligned} I = & C - \frac{1}{M} \sum_{i \in \mathcal{I}} \int_0^T \int_0^{L_i} \log\{\bar{v}_i(x_i) - v_i(x_i, t)\} [v_i(x_i, t) \\ & - v_i(x_i)] dx_i dt - \frac{1}{M} \sum_{i \in \mathcal{S}} \int_0^T \log\{\bar{q}_i^{\text{in}}(t) - q_i^{\text{in}}(t)\} q_i^{\text{in}}(t) dt \\ & - \frac{1}{M} \sum_{i \in \mathcal{F}} \int_0^T \left[\sum_{d \in C_i} \log \beta_{i,d}(t) \right] + \log \left[1 - \sum_{d \in C_i} \beta_{i,d}(t) \right] dt \end{aligned} \quad (\text{A1})$$

subject to Eqs. (22) and (23).

Algorithm 2 Airline continuous link flow optimization algorithm

Start Guess a feasible control $u = (v, q^{\text{in}}, \beta)$: take q^{in} small, v homogeneous over space, and $\beta_{i,d} = 1/|D_i|$. Furthermore, choose two tolerance criteria ϵ_1 and ϵ_2 .

- 1) Repeat.
- 2) Repeat.
- 3) Solve PDEs (22) and (23).
- 4) Solve adjoint PDEs (A3) and (A4).
- 5) Form gradient $\nabla_u I$ using Eq. (A2).
- 6) Solve the auxiliary optimization program to obtain the approximate Newton step δu (see Raffard [31]).
- 7) Line search: compute $\alpha > 0$ such that $I(u + \alpha \delta u)$ is minimized.
- 8) Update control: $u := u + \alpha \delta u$.
- 9) Until $|\langle \nabla_u I, \delta u \rangle| < \epsilon_2$.
- 10) Increase logarithmic barrier size M .
- 11) Until $1/M < \epsilon_1$.
- 12) Return $u_{\text{optimal}} = u$.

2) The adjoint method yields an expression of the gradient of the augmented cost function as a function of the control, the state, and the costate. The gradient of the original cost function C with respect to the control $u = (q^{\text{in}}, v, \beta)$ is summarized as follows:

$$\begin{aligned} \nabla_{q^{\text{in}}} C(t) &= \rho^*(0, t) & \nabla_v C(x, t) &= \rho(x, t) \frac{\partial \rho^*(x, t)}{\partial x} \\ \nabla_{\beta_{i,d}} C(t) &= \rho_i(L_i, t) v_i(L_i, t) \left[\rho_{d_i}^*(0, t) - \rho_d^*(0, t) \right] \end{aligned} \quad (\text{A2})$$

where

$$\rho^*: \prod_{i \in \mathcal{I}} [0, L_i] \times [0, T] \rightarrow \mathbb{R}$$

is the costate variable. The gradient terms due to the logarithmic barrier for each inequality constraint have been omitted due to space restrictions. The costate ρ^* is selected to eliminate dependence of the gradient calculation on the state variables and must satisfy the adjoint PDE

$$\frac{\partial \rho^*}{\partial t} + v \frac{\partial \rho^*}{\partial x} = \lambda \quad (\text{A3})$$

with boundary conditions

$$\frac{d\rho^*(L, t)}{dt} = \int_0^t q^s(L, s) - q(L, s) ds, \quad \rho^*(L, T) = 0 \quad (\text{A4})$$

3) Derivation of the Newton step is omitted for space considerations. The interested reader is referred to the work of Raffard [31] for a derivation of the Newton step on this and other similar problems.

Finally, the local optimization procedure is summarized in Algorithm 2.

Acknowledgements

This work is supported by the Office of Naval Research (ONR) under the CoMotion Multi-University Research Initiative (MURI) contract N00014-02-1-0720, by NASA Joint University Program (JUP) under grant NAG2-1564, and by NASA grant NCC2-5536.

References

- [1] *Transportation Statistics Annual Report*, U.S. Dept. of Transportation, Research and Innovative Technology Administration, Bureau of Transportation Statistics, http://www.bts.gov/publications/transportation_statistics_annual_report/2006/ [retrieved 14 Sept. 2007].
- [2] "National Airspace System Operational Evolution Plan," Ver. 8, Federal Aviation Administration, 2006.
- [3] Vossen, T. V., and Ball, M. O., "Slot Trading Opportunities in Collaborative Ground Delay Programs," in *Transportation Science*, Vol. 40, No. 1, Feb. 2006, pp. 29–43.
- [4] Ball, M. O., Hoffman, R., Knorr, D., Wetherly, J., and Wambsganss, M., "Assessing the Benefits of Collaborative Decision Making in Air

- Traffic Management," *Proceedings of the 3rd USA/Europe Air Traffic Management R&D Seminar*, http://atmseminar.eurocontrol.fr/all-seminars/3rd-usa-europe-atm-r-d-seminar/paper_020/attachment_download/file [retrieved 14 Sept. 2007].
- [5] Ball, M. O., Donohue, G. L., and Hoffman, K., "Auctions for the Safe, Efficient and Equitable Allocation of Airspace System Resources," *Combinatorial Auctions*, edited by P. Cramton, Y. Shoham, and R. Steinberg, MIT Press, Cambridge, MA, 2006, Chap. 22, pp. 507–538.
- [6] Wilson, R., "Architecture of Power Markets," *Econometrica : Journal of the Econometric Society*, Vol. 70, No. 4, July 2002, pp. 1299–1340.
- [7] Falkner, M., Devtsikiotis, M., and Lambadaris, I., "An Overview of Pricing Concepts for Broadband IP Networks," *IEEE Communications Surveys and Tutorials*, Vol. 3, No. 2, 2000, pp. 2–13; also available online at <http://dl.comsoc.org/cocoon/comsoc/servlets/GetPublication?id=1228373> [retrieved 14 Sept. 2007].
- [8] Mas-Colell, A., Whinston, M. D., and Green, J. R., *Microeconomic Theory*, 2nd ed., Oxford Univ. Press, New York, 1995, pp. 545–572.
- [9] Kelly, F. P., Maulloo, A. K., and Tan, D. K. H., "Rate Control for Communication Networks: Shadow Prices, Proportional Fairness and Stability," *Journal of the Operational Research Society*, Vol. 49, No. 3, 1998, pp. 237–252.
- [10] Johari, R., and Tsitsiklis, J. N., "Efficiency Loss in a Network Resource Allocation Game," *Mathematics of Operations Research*, Vol. 29, Aug. 2004, pp. 407–435. doi:10.1287/moor.1040.0091
- [11] Odoni, A. R., "The Flow Management Problem in Air Traffic Control," *Flow Control and Congested Networks*, edited by A. R. Odoni and Z. Szego, Springer-Verlag, Berlin, 1987.
- [12] Terrab, M., and Odoni, A. R., "Strategic Flow Management for Air Traffic Control," *Operations Research*, Vol. 41, No. 1, Jan.–Feb. 1993, pp. 138–152.
- [13] Bertsimas D., and Stock Patterson, S., "The Air Traffic Flow Management Problem with En Route Capacities," *Operations Research*, Vol. 46, No. 3, May–June 1998, pp. 406–422.
- [14] Bertsimas D., and Stock Patterson, S., "The Traffic Flow Management Rerouting Problem in Air Traffic Control: A Dynamic Network Flow Approach," *Transportation Science*, Vol. 34, No. 3, Aug. 2000, pp. 239–255. doi:10.1287/trsc.34.3.239.12300
- [15] Hunter, G., Ramamoorthy, K., Cobb, P., Huang, A., Blake, M., and Klein, A., "Evaluation of Future National Airspace System Architectures," AIAA Modeling and Simulation Technologies Conference and Exhibit, San Francisco, AIAA Paper 2005-6492, Aug. 2005.
- [16] Krozel, J., Penny, S., Prete, J., and Mitchell, J. S., "Comparison of Algorithms for Synthesizing Weather Avoidance Routes in Transition Airspace," AIAA Conference on Guidance, Navigation, and Control, Providence, RI, AIAA Paper 2004-4790, Aug. 2004.
- [17] Lighthill, M. J., and Whitham, G. B., "On Kinematic Waves, 2: A Theory of Traffic Flow on Long Crowded Roads," *Proceedings of the Royal Society of London, Series A: Mathematical and Physical Sciences*, Vol. 229, No. 1178, 1955, pp. 317–345.
- [18] Richards, P. I., "Shock Waves on the Highway," *Operations Research*, Vol. 4, No. 1, June 1956, pp. 42–51.
- [19] Menon, P., Sweriduk, G., and Bilimoria, K., "New Approach for Modeling, Analysis and Control of Air Traffic Flow," *Journal of Guidance, Control, and Dynamics*, Vol. 27, No. 5, 2004, pp. 737–744.
- [20] Bayen, A. M., Raffard, R., and Tomlin, C. J., "Eulerian Network Model of Air Traffic Flow in Congested Areas," *Proceedings of the 2004 American Control Conference*, Vol. 6, Inst. of Electrical and Electronics Engineers, Piscataway, NJ, 30 June–2 July 2004, pp. 5520–5526.
- [21] Bayen, A. M., Raffard, R., and Tomlin, C. J., "Adjoint-Based Constrained Control of Eulerian Transportation Networks: Application to Air Traffic Control," *Proceedings of the 2004 American Control Conference*, Vol. 6, Inst. of Electrical and Electronics Engineers, Piscataway, NJ, 30 June–2 July 2004, pp. 5539–5545.
- [22] Sun, D., Yang, S. D., Strub, I., and Bayen, A. M., "Eulerian Trilogy," AIAA Guidance, Navigation, and Control Conference and Exhibit, Keystone, CO, AIAA Paper 2006-6227, Aug. 2006.
- [23] Young, H. P., *Equity in Theory and in Practice*, Princeton Univ. Press, Princeton, NJ, 1993, pp. 66–79.
- [24] Arrow, K. J., and Debreu, G., "Existence of an Equilibrium for a Competitive Economy," *Econometrica : Journal of the Econometric Society*, Vol. 22, No. 3, July 1954, pp. 265–290.
- [25] CVX, Software Package, Ver. 0.85, Stanford Univ., Stanford, CA; also available at <http://www.stanford.edu/~boyd/cvx/> [retrieved 14 Sept. 2007].

- [26] SeDuMi: Self Dual Minimization, Software Package, Ver. 1.1, Stanford Univ., Stanford, CA; also available at <http://sedumi.mcmaster.ca/> [retrieved 14 Sept. 2007].
- [27] Matlab, Software Package, Ver. 7.1.0, The Mathworks, Inc., Natick, MA.
- [28] Shor, N. Z., *Minimization Methods for Non-Differentiable Functions*, Springer Series in Computational Mathematics: 3, Springer-Verlag, New York, 1985, pp. 22–44.
- [29] Jameson, A., “Aerodynamic Design via Control Theory,” *Journal of Scientific Computing*, Vol. 3, Nov. 1988, pp. 233–260. doi:10.1007/BF01061285
- [30] Bertsekas D. P., *Nonlinear Programming*, 2nd ed., Athena Scientific, Belmont, MA, 1995.
- [31] Raffard, R. L., “Optimal Control of Systems Governed by Differential Equations with Applications in Air Traffic Management and Systems Biology,” Ph.D. Thesis, Stanford Univ., Stanford, CA, Dec. 2006.