

Toward Efficient and Equitable Distributed Air Traffic Flow Control[†]

Steven L. Waslander, Robin L. Raffard, Claire J. Tomlin,
 Department of Aeronautics and Astronautics, Stanford University*

Abstract—In order to ensure efficient and equitable use of the National Airspace, the problem of air traffic flow control is posed as a PDE constrained optimization program which explicitly incorporates distinct preferences for independent airlines. Novel cost metrics are introduced to capture the real-world cost of delays to the airlines, and control parameters are aligned with the actual inputs the FAA can manipulate in dealing with adverse weather conditions, which include ground-delay programs, miles-in-trail restrictions and flight reroutes. A market mechanism is proposed that allows for distributed computation of efficient and equitable solutions. The resulting solution preferentially reduces delays for expensive flights, but compensates less expensive flights via transfer payments.

I. INTRODUCTION

Air traffic flow control in the U.S. National Airspace (NAS) is an extremely complex and dynamic problem. There are currently as many as 6000 flights in the air at peak times, and weather-induced disruptions to NAS capacity affect many flights from many airlines simultaneously. Currently, adjustments are made through human intervention, in the form of experienced flow control specialists at the FAA Air Traffic Control System Command Center. There is a clear need for fast, scalable optimization tools that can ensure efficient use of the NAS when disruptions occur.

The FAA has three main avenues for controlling scheduled operations in real time, namely ground-delay programs, miles-in-trail restrictions and flight reroutes. Ground delays restrict the inflow bound for constrained resources, miles-in-trail reduce en-route vehicle velocity and flight reroutes ensure aircraft are redirected around weather disruptions. Since such adjustments directly affect airline flight costs, it is desirable to include the airline perspective in the formulation of optimal control strategies.

In previous work [1], a cooperative, distributable algorithm was proposed to solve a PDE constrained optimization program formulation of the air traffic flow control problem for multiple airlines with known cost functions. The aggregation of air traffic into a continuous flow over an Eulerian network ensures good scalability in the number of aircraft, as the exponential deterioration of computational complexity for combinatorial optimization is avoided. By agreeing to the Nash Bargaining Solution (NBS), an equitable outcome is ensured amongst the parties involved.

Similar approaches based on controlling aggregate flows have been proposed [2], [3], which suffer from a few key drawbacks. First, the methods do not include the ability to determine control actions for velocity restriction, which

reduces the set of possible outcomes and hence, limits the efficiency gain of the system. Second, the traffic flow is not necessarily represented as faithfully as the continuous PDE model. Finally, the issues involved with incorporating airline preferences are not addressed.

This paper seeks to advance our previous results in three significant directions. First, a novel airline cost function is introduced that translates the key metric of flight delay to the continuous domain. Second, a method for inflow control is added to match the triad of controls available to the FAA. Third, the impact of the method on airline incentives is addressed. This leads to the introduction of a publicly fair schedule-based allocation which ensures equitable allocation of available resources, and a market-based approach to maintain equity while achieving efficiency as well.

This paper proceeds as follows. A simplified resource allocation problem is considered in order to develop the issues involved in selecting a specific allocation method. Then, a more realistic model of the air traffic flow control problem is presented, which results in a PDE constrained optimization process that is shown to be distributable amongst individual airlines with centralized coordination through a market based algorithm. Finally, results are presented for a large scale simulation that highlights the benefits of the various extensions to the true air traffic flow control problem.

II. EFFICIENCY AND EQUITY IN ALLOCATION

The problem of managing air traffic flow is essentially one of allocating scarce resources to the airlines, all of whom have both prior claims to the resources based on the pre-approved flight schedule as well as private information concerning their preferences over allocations. This section develops a problem formulation that captures these elements, and for which a market-based approach can be designed.

A. Simplified Resource Allocation

In order to better illustrate the issues involved, it is beneficial to reduce the problem to a simplified resource allocation model. The simplified model fixes a single source and destination, labeled $(\mathcal{S}, \mathcal{Z})$, a set of links $i \in \{1, \dots, N_{\mathcal{I}}\} = \mathcal{I}$ and a set of airlines $j \in \{1, \dots, N_{\mathcal{J}}\} = \mathcal{J}$. A decision must be made as to the flow $y_{i,j} \in \mathbb{R}_+$ to be allowed along each link for each airline¹.

The allocated flows must satisfy aggregate capacity constraints along each link,

$$\sum_{j \in \mathcal{J}} y_j \leq y^{max} : y^{max} \in \mathbb{R}_+^{N_{\mathcal{I}}} \quad (1)$$

¹Throughout the paper, $y_j = (y_{1,j}, \dots, y_{N_{\mathcal{I}},j})$, $j \in \mathcal{J}$, shall refer to the vector of flow for each airline and $y = (y_1^T, \dots, y_{N_{\mathcal{J}}}^T)$ is the matrix of all flows, with a similar form adopted for all other variables.

[†] This research was supported by ONR under the CoMotion MURI contract N00014-02-1-0720, by NASA JUP under grant NAG2-1564, and by NASA grant NCC2-5536

* {stevenw, rraffard, tomlin}@stanford.edu

The flow, y , shall also be referred to as an allocation, and the links as resources with capacity y^{max} . The schedule, $y^{des} \in \mathbb{R}_+^{N_I \times N_J}$, is set as a predefined allocation that does not necessarily satisfy capacity constraints. Finally, each airline is assumed to have a cost function, $J_j : (y_j) \in \mathbb{R}^{N_I} \rightarrow \mathbb{R}_+$ that defines a preference relation over all allocations. The vector of airline cost functions shall be referred to as $J = (J_1, \dots, J_{N_J})$.

The goal of the ensuing discussion is to determine an allocation rule $\mathcal{A} : (y^{max}, y^{des}, J) \mapsto y$, which takes the problem defined by the available resources, the schedule and the airline cost functions, and maps it to an allocation in a manner that makes efficient use of the airspace and is equitable to all airlines.

B. Schedule-Based Allocation

Without considering the airline cost functions, it is possible to define a schedule-based allocation rule that exhibits desirable properties in terms of equity amongst the airlines. An allocation rule that is impartial amongst the airlines and collusion-proof is defined by scaling all flows along contested links proportional to the pre-defined schedule,

$$y_{i,j}^p = \frac{y_{i,j}^{des}}{\sum_{k \in \mathcal{J}} y_{i,k}^{des}} y_i^{max} \quad \forall (i,j) \in \mathcal{I} \times \mathcal{J} \quad (2)$$

Both the strength and weakness of the proportional allocation rule is that it is independent of airline cost functions and private information. Since 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 likely that scheduled flights can vary dramatically in the costs that are incurred by delays. As such, any allocation that relies on schedule information alone cannot ensure no other solution exist that all airlines prefer. This condition is referred to interchangeably as Pareto optimality or efficiency.

Pareto Optimality: An allocation, $\mathcal{A}(y^{max}, y^{des}, J)$, is Pareto optimal if no airline's cost can be decreased without increasing the cost to at least one other airline.

Since the schedule-based allocation does not guarantee Pareto optimality, it is inherently problematic. The potential exists for discussion amongst airlines to result in dissatisfaction with an inefficient process, and side trades external to the allocation process might result. 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.

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

C. Cost-Based Allocation

A cooperative resource allocation method seeks to solve,

$$\begin{array}{ll} \underset{y}{\text{minimize}} & J_*(J_1(y_1), \dots, J_{N_J}(y_{N_J})) \\ \text{subject to} & \sum_{j \in \mathcal{J}} y_j \leq y^{max} \end{array} \quad (3)$$

where $J_* : \mathbb{R}^{N_J} \rightarrow \mathbb{R}$ is any function that defines a preference relation over all allocations. The problem is further restricted at this point to afford the possibility of uniquely defined allocation rules that exhibit both efficient and equitable properties. It is assumed that cost functions are convex and strictly decreasing, and that the allocation space is convex, which ensures that the problem is convex which guarantees existence of a unique solution. Note that these assumptions are relaxed for the model developed in Section III, and thus neither guarantee of existence or uniqueness of the solution is conserved.

Aggregate Cost Solution: It is possible to assume that all costs are defined in a common unit of measure, or currency, and that payments between airlines are possible. In this case, the standard of comparison is fixed by the units of the cost functions themselves. With these assumptions, an allocation rule that is efficient must minimize the Aggregate Cost (AC),

$$J_{AC} = \sum_{j \in \mathcal{J}} J_j(y_j) \quad (4)$$

D. Non-Cooperative Incentives

One of the key benefits of AC allocation is that there exist methods for ensuring that the outcome remains efficient, even in the face of private information. Two further definitions from equity theory [4] are needed,

Weak Consistency: The cost incurred by every subset of airlines in applying the same allocation rule to its portion of the total allocation is unchanged if all other airlines' allocations are held fixed, although the allocation may not be unique.

Weak consistency implies that no subset of airlines will be dissatisfied with their allocation if the same rule were applied to a reduced problem including only the subset. Full consistency further requires that if the subset of airlines agree to an alternate allocation that satisfies the rule, then the full allocation still obeys the allocation rule. Instead of requiring full consistency, the combination of weak consistency as described above, and replicability is sufficient for the main result of this section, Theorem 1.

Replicability: An allocation method is replicable if for every problem for which y is an allocation, the matrix, $\text{diag}(y, \dots, y)$ is an allocation for the problem expanded to include k copies of the original problem.

One further definition from welfare economics is required,

Competitive Allocation: A competitive allocation is an allocation for which a price vector, $p \in \mathbb{R}_+^{N_I}$, exists over all links such that each airline prefers its allocation to one any other allocation that could be purchased at market prices.

The following result [5], presented without proof, shows that a competitive allocation from proportional division exists and ensures both efficiency and equity are addressed in the resource allocation problem, for a restricted set of problems.

Theorem 1: (Thomson, 88) Assume that all airline cost functions, J_j , are differentiable, strictly decreasing, strictly quasi-convex and that the marginal cost approaches $-\infty$ as the allocation of a good tends to 0. Then, competitive allocation from proportional division is acceptable, Pareto

optimal, weakly consistent and replicable and conversely, all allocation methods that exhibit these properties are competitive allocations from proportional division.

For airlines that can exchange funds and measure costs in dollars, this result 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, since the market is closed, airlines with less expensive flights stand to reduce costs considerably over a schedule-based solution, as 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, as each airline is responsible for optimizing their own costs and determining the additional quantity of each resource to purchase at market prices above their schedule based allocation. The role of the FAA transitions to one of determining capacity constraints along links and setting market prices that equalize supply and demand.

III. DISTRIBUTED AIR TRAFFIC FLOW CONTROL

In order to develop a distributed, market based resource allocation method for air traffic flow control, a realistic model is introduced which captures the problem's complex nature while ensuring tractability of the algorithm. This section proceeds with the definition of a PDE constrained aggregate flow model. A market-based optimization method is then presented in which airlines calculate the control strategy in a distributed manner.

Currently, commercial air traffic is required to file flight plans that are defined by fixed waypoints. As first described in [6], inspection of Enhanced Traffic Management System (ETMS) data reveals that a fixed network topology matches well with actual flight data. It is possible to define an integer optimization program that includes individual decisions on routes and ground delays for each flight in the NAS [2], however, the problem is NP-hard and of sufficient size to be intractable. Instead, air traffic is approximated as a continuous flow over the links of a fixed network, as first introduced in [6]. The model can be interpreted as incorporating a continuity equation for each link with appropriate flow boundary conditions at the nodes.

A. Network Description

Consider an air traffic network, as in Fig. 1, defined as above by a fixed set of links $i \in \mathcal{I}$ and a fixed number of airlines $j \in \mathcal{J}$. Each link, i , is of length L_i , with the coordinate $x_i \in [0, L_i]$ used to refer to any point on link i . Each link has an associated airline specific aircraft density, $\rho_{i,j} : (x_i, t) \in [0, L_i] \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$, and airline specific aircraft velocity profile over each link, $v_{i,j} : (x_i, t) \in [0, L_i] \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$. Where convenient, the flux, $q_{i,j} : (x_i, t) \in [0, L_i] \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$, shall be used to refer to the product of density and velocity.




Fig. 1. Link-based network depiction for a subsection of the NAS on the eastern seaboard.

A subset of links $\mathcal{S} \subseteq \mathcal{I}$, are defined as sources and, $\mathcal{Z} \subseteq \mathcal{I}$, are defined as sinks. Each source has associated with it aircraft inflow rates per airline of $q_{i,j}^{in} : t \in \mathbb{R}_+ \rightarrow \mathbb{R}_+$.

Each link $i \in \mathcal{I}$ has a set of links \mathcal{M}_i , which merge into link i and a set of links \mathcal{D}_i , which diverge from link i . Let $\mathcal{M} = \cup_{i \in \mathcal{I}} \mathcal{M}_i$ represent the set of all merging links and $\mathcal{D} = \cup_{i \in \mathcal{I}} \mathcal{D}_i$ represent the set of all diverging links.

For any diverging link $d \in \mathcal{D}$, the flow split parameter, $\beta_{i,j,d} : t \in \mathbb{R}_+ \rightarrow [0, 1]$ describes the j^{th} airline's portion of the flow which travels from link i into link d . Continuity requires that the flow split amongst diverging links sum to one for each airline, namely that $\sum_{d \in \mathcal{D}_i} \beta_{i,j,d}(t) = 1$. For each link in the diverging set, $d \in \mathcal{D}$, there exists a set of inflow links with at least one element, $f \in \mathcal{F}_d \neq \emptyset$.

In order 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, $v_i^{\min}(x_i)$ and $v_i^{\max}(x_i)$, defined at each point in the network. A maximum density limit, $\bar{\rho}_i : (x_i, t) \in [0, L_i] \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is imposed along each link to ensure safe use of the available airspace.

Continuity results in the following PDEs,

$$\frac{\partial \rho_{i,j}}{\partial t} + \frac{\partial q_{i,j}}{\partial x_i} = 0, \quad \forall (i, j) \in \mathcal{I} \times \mathcal{J} \quad (5)$$

Completing the definition of the PDE model, initial densities are defined for all airlines along all links in the network and continuity conditions are fixed for all time between merging and dividing links.

$$\begin{aligned} \rho_{i,j}(x_i, 0) &= \rho_{i,j}^0(x_i), \quad \forall (i, j) \in \mathcal{I} \times \mathcal{J} \\ q_{i,j}(0, t) &= q_{i,j}^{in}(t) + \sum_{k \in \mathcal{M}_i} q_{k,j}(L_k, t) \\ &\quad + \sum_{f \in \mathcal{F}_i} \beta_{f,j,i}(t) q_{f,j}(L_f, t), \end{aligned} \quad (6)$$

$$\forall (i, j) \in \mathcal{I} \times \mathcal{J}$$

B. Cost Formulation

In order to react strategically to unpredictable phenomenon such as weather, it is necessary to define a preference relation over possible outcomes, *a priori*. A key metric in determining desirable allocations of NAS resources during disrupted periods is that of cumulative arrival delay incurred, as additional costs incurred by airlines are directly related to the increased flight time and schedule disruptions. Although deviations from the planned routes, departure delays and en-route velocity changes can impact airline costs, it is assumed that arrival delays are the most prominent concern.

As the timescale for this problem is on the order of hours, a finite time horizon, $T \in \mathbb{R}_+$, is set. It is assumed that the airlines have agreed to a predefined schedule, which can be represented as the desired number of aircraft, $n_{i,j}^{des} : t \in \mathbb{R}_+ \rightarrow \mathbb{R}_+, \forall i \in \mathcal{Z}$, arriving at the destination airports. Actual arrivals, $n_{i,j}(t)$, can then be measured against the ideal and control decisions can be optimized in order to reduce overall delay for each airline. A characteristic cost function, $J_j(n_j) \in \mathbb{R}_+, \forall j \in \mathcal{J}$, can be defined as,

$$J_j(n_j) = - \sum_{i \in \mathcal{Z}} \int_0^T \alpha_{i,j}(t) (n_{i,j}^{des}(t) - n_{i,j}(t))^2 dt \quad (7)$$

where the function $\alpha_{i,j} : (t) \in \mathbb{R}_+ \rightarrow \mathbb{R}_+$ defines a time varying delay cost for each airline. In the context of the PDE flow model, the number of arriving aircraft can be converted to an integral over the arrival flux,

$$J_j(q_j) = - \sum_{i \in \mathcal{Z}} \int_0^T \alpha_{i,j}(t) \left(\int_0^t q_{i,j}^{des}(L_i, s) - q_{i,j}(L_i, s) ds \right)^2 dt \quad (8)$$

Each airline therefore seeks to minimize the squared deviation in cumulative flow arriving at a destination airport with respect to the scheduled flow. If airlines have different types of aircraft with different costs, separate flows must be maintained for each type during the local optimization process. The quadratic nature of the penalty penalizes large deviations from the flight schedule more heavily than small ones.

C. Network Flow Optimization

In keeping with the FAA control authority over NAS, three distinct flow control methods are included in the optimization formulation. The GDP approach is embodied in control of inflow, $q_{i,j}^{in}(t)$, miles-in-trail restrictions are similar to controlling link velocity profiles, $v_{i,j}(t)$, and the flow splitting parameters, $\beta_{i,j,d}(t)$ allow for flight rerouting along alternate links. The multi-airline optimization problem can now be formulated as a multi-objective optimization,

$$\begin{array}{ll} \min. & \sum_{j \in \mathcal{J}} J_j(q_j) \\ \text{s. t.} & (5) \text{ and } (6) \\ & v_i^{\min} \leq v_{i,j} \leq v_i^{\max}, \forall (i,j) \in \mathcal{I} \times \mathcal{J} \\ & \sum_{j \in \mathcal{J}} \rho_{i,j} \leq \bar{\rho}_i, \quad \forall i \in \mathcal{I} \\ & \sum_{d \in \mathcal{D}_i} \beta_{i,j,d} = 1, \quad \forall (i,j) \in \mathcal{I} \times \mathcal{J} \end{array} \quad (9)$$

D. Market Based Allocation Method

Market based allocation is derived in four steps.

1) *Schedule-Based Allocation*: Due to the added complexity of the full air traffic flow model relative to the simplified model of Section II, it is no longer possible to scale allocations directly for flows along a contested link. Instead, inflows that are implicated in the excess demand are scaled proportional to the scheduled flow, and the possibility of rerouting or metering the flow is not considered. For each

contested link, the schedule-based allocation, $q_{i,j}^s$ with link density profile $\rho_{i,j}^s$ is the result of scaling back inflows until a feasible solution is achieved.

2) *Dual Decomposition*: Next, a dual decomposition [7], [1] is performed on the centralized, aggregate cost minimization form of Eq. (9). Inspection of Eq. (9) 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 local optimizations. The Lagrangian is formed,

$$\mathcal{L}(\rho, \lambda) = \sum_{j \in \mathcal{J}} (J_j + \lambda^\top \rho_j) - \lambda^\top \bar{\rho} \quad (10)$$

where the dual variable, $\lambda : (x, t) \in \prod_{i \in \mathcal{I}} [0, L_i] \times [0, T] \rightarrow \mathbb{R}_+$, is referred to as the price vector. The dual problem can be written as,

$$\max. g(\lambda) = \sum_{j \in \mathcal{J}} \inf_{\rho} \{J_j + \lambda^\top \rho_j\} - \lambda^\top \bar{\rho} \quad (11)$$

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

$$\frac{\partial g}{\partial \lambda} = \sum_{j \in \mathcal{J}} \rho_j^{opt} - \bar{\rho} \quad (12)$$

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

3) *Market Definition*: With the market pricing for restricted resources determined, a payout $p_j \in \mathbb{R}$ is defined for each airline,

$$p_j = \sum_{i \in \mathcal{I}} \lambda_i^\top (\rho_{i,j} - \rho_{i,j}^s) \quad (13)$$

In order to ensure that the resulting allocation is acceptable, the payout charges airlines for the difference between the resources selected and those received by the airline from the schedule-based allocation. The airline market cost, $C_j = J_j + p_j, \in \mathbb{R}$, is therefore defined as the sum of the delay costs and the resource trade costs, p_j . Acceptability follows from the fact that each airline is optimizing an improvement over the cost incurred by a schedule-based allocation.

$$\begin{aligned} J_j(q_j^s) &= J_j(q_j^s) + \lambda^\top (\rho_j^s - \rho_j^s) \\ &\geq \inf_{\rho_j} (J_j(q_j) + \lambda^\top (\rho_j - \rho_j^s)) \\ &= \inf_{\rho_j} C_j \end{aligned} \quad (14)$$

4) *Local Optimization*: The local optimization problem can be solved by applying the adjoint method[8] to the airline flow control problem. Note that the control variables for this optimization problem are q^{in} , v and β , which implies that the airlines are deciding the strategies to be taken in relieving excess demand, while the FAA assumes a more supervisory role of setting prices for resources.

$$\begin{array}{ll}
\min. & C_j \\
\text{s. t.} & (5) \text{ and } (6) \\
& v_{i,j}^{\min} \leq v_{i,j} \leq v_{i,j}^{\max} \quad \forall i \in \mathcal{I} \\
& \sum_{d \in \mathcal{D}_i} \beta_{i,j,d} = 1 \quad \forall i \in \mathcal{I}
\end{array} \quad (15)$$

A descent direction is needed to perform the above optimization, which can be found by calculating a first order variation of the objective function, the PDE constraint and the inequality constraints and applying the adjoint method to eliminate the state variable variations. Dropping dependence on j , the resulting gradient of the local objective function with respect to the control variables is,

$$\begin{aligned}
\nabla_{q^{in}} C(t) &= \rho^*(0, t) \\
\nabla_v C(x, t) &= \rho(x, t) \frac{\partial \rho^*(x, t)}{\partial x} \\
\nabla_{\beta_{i,d}} C(t) &= \rho_i(L_i, t) v_i(L_i, t) (\rho_k^*(0, t) - \rho_d^*(0, t))
\end{aligned} \quad (16)$$

where, $\rho^* : (x_i, t) \in \prod_{i \in \mathcal{I}} [0, L_i] \times [0, T] \rightarrow \mathbb{R}^{N_{\mathcal{I}}}$, is the costate variable and the index k refers to the diverging link $k \in \mathcal{D}_i$ that has the highest index, and therefore receives flow portion $1 - \sum_{r \neq k} \beta_{i,r}$. This technical detail allows the elimination of an equality constraint from the local problem formulation. The gradient terms due to the inequality constraints have been omitted due to space restrictions. The costate, ρ^* , must satisfy the following adjoint equation,

$$\frac{\partial \rho^*}{\partial t} + v \frac{\partial \rho^*}{\partial x} = \lambda \quad (17)$$

with the following boundary condition

$$\begin{aligned}
\frac{d\rho^*(L, t)}{dt} &= \int_0^t q^{des}(L, s) - q(L, s) ds \\
\rho^*(L, T) &= 0
\end{aligned} \quad (18)$$

The local optimization proceeds by solving the PDE and costate PDE, and then forming the gradient to determine a descent direction in which to update all control variables.

IV. IMPLEMENTATION

A. Scenario

Four airlines are competing on a network (see Fig. 1), composed of 3 airports: New York (NYC), Washington (DC) and Chicago (ORD). 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 more delay with respect to 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 200 minutes.

Among the four airlines, there are two large airlines and two small airlines, with desired terminal landing rates of 1 aircraft every minute (A/L 1,2) and every 4 minutes (A/L 3,4). The airline costs per minute delay are of two types, one that is a constant value of 1 (A/L 2,4) 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 (A/L 1,3), representing a main-line, hub and spoke schedule with large aircraft.

Initial flow conditions are assumed to be adversely offset from the desired flow, and the maximum density along each link is restricted to 1 aircraft every 8 nautical miles due to ongoing bad weather. Both delays on flights and capacity constraint violations result along the central, direct links.

The market algorithm was implemented in a distributed manner which faithfully represents the airlines as individual decision makers with private information. Separate computing nodes were assigned local optimizations for each airline and destination airport. The central node collected the results upon completion and updated prices. Relying on standard desktop computers, average node computation times were under 20 seconds.

B. Results

Convergence of the price-based subgradient method can be quite slow[7], and in practice, optimization of the scenario took on the order of 3000 iterations to achieve near-feasible solutions. The optimization resulted in a market-based allocation that is feasible everywhere except at some peak times along some 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 so as to meet the maximum density requirement without significantly altering the outcome of the market based approach.

A sequence of three images is presented in Fig 2, which depict the airline-specific flights along the various links at chronological instants in the simulation. The first image shows a near uniform distribution of flights which were assumed as initial conditions, and subsequent images show the evolution of the solution in favor of the larger airlines. The airlines with less aircraft end up rerouting to the longer routes while the airlines with more aircraft and especially with larger planes end up consuming all available capacity on the central link.

The resulting delays for incoming flights at each point in time are displayed in Fig. 3. Note that in all cases, the inflow-restricted schedule based delays are significantly worse than either the ideal or the optimized solutions. It is also interesting to note that the large airline with high costs matches their ideal delay most closely and that all other airlines incur at least slight increases in delay over their ideal result.

In the market driven process, the airlines pay the FAA if they want to reduce the delays incurred by the inflow-restricted schedule, assuming a link is at capacity. This phenomenon is represented in Table I, in which Airlines 1 and 2 have bought flow from Airlines 3 and 4. A net payment is made to the FAA which arises from the improvement in efficiency that can be had over an inflow-restricted schedule.

V. CONCLUSIONS

Air traffic flow control is an extremely complex process for which models that accurately reflects its true nature are

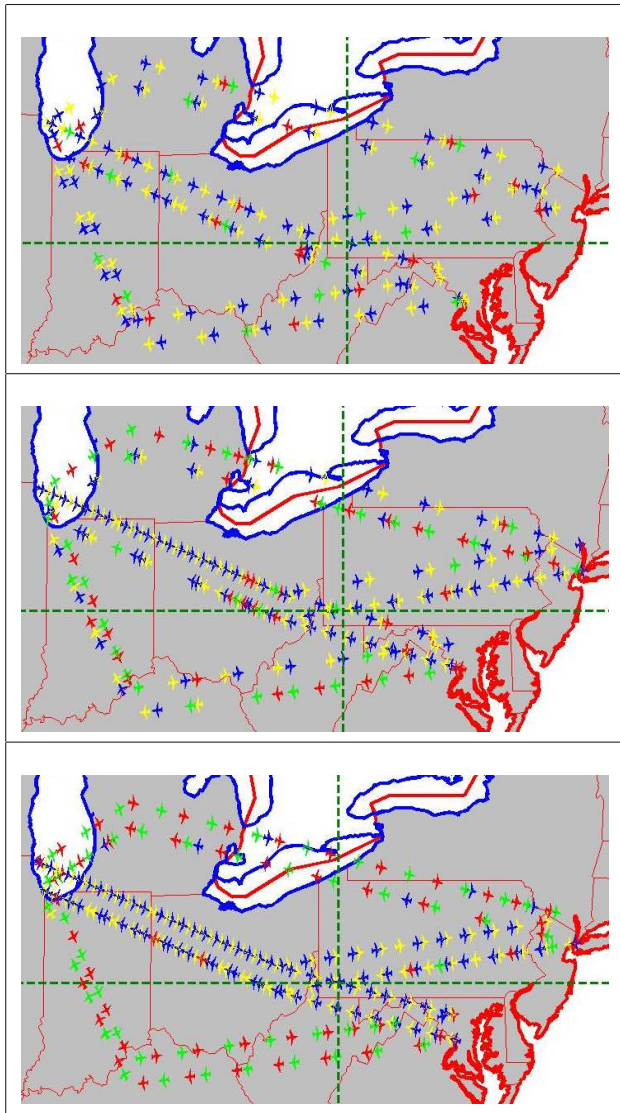


Fig. 2. Market-based optimized traffic flow at three different times in the simulation: $t = (0, 50, 100)$ minutes. Airlines 1,2,3 and 4 are colored blue, yellow, red and green, respectively. As time progresses, the central link becomes filled to capacity with flights primarily from the large and expensive airlines. Videos of ideal, schedule based and market-based flows can be seen at <http://cherokee.stanford.edu/~raffard/ATF>.

difficult to develop. In making the PDE constrained flow model more realistic, this paper extended our previous modeling in two key directions. The addition of inflow restriction to the control choices available resulted in greater flexibility to accommodate weather disruptions. The reformulation of the method in terms of the variable cost arrival delays ensured that the results of the optimization align with the true interests of the airlines.

Most importantly, this paper evaluated the potential for real-world implementation of such a system. Since it cannot be assumed that all airlines will willingly divulge private cost information in order to allow the FAA to determine an equitable resource allocation, a market based approach was presented that allows the airlines to iteratively converge to flow control strategies that avert excess demand on restricted




Fig. 3. Delay as a function of time experienced at ORD for the 4 different airlines. The ideal delay is represented by the dash-dotted green line, the inflow-restricted delay by the solid blue curve, and the market-based delay by the dashed red curve. Airlines 1 and 2 experience the most delays, while airline 3 and 4 have smaller delays (no delay for the last 50 minutes).

	A/L 1	A/L 2	A/L 3	A/L 4
Scaled inflow cost	4633	1194	142	39
Optimized cost	719	384	31	-10
Payment to FAA	96	95	-1	-25

TABLE I. In the market process, all 4 airlines reduce their cost by a factor of 2.5 to 7. Airline 1 makes the largest payment to the FAA, while Airline 3 and 4 receive money from the FAA.

elements of the NAS. 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 unlock the maximum aggregate amount of savings from the available airspace resources, rewarding all airlines with reduced costs.

REFERENCES

- [1] 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 *Proceedings of the AIAA Guidance, Navigation and Control Conference and Exhibit*, San Francisco, Aug 2005.
- [2] D. Bertsimas and S. Stock Patterson, "The traffic flow management rerouting problem in air traffic control: A dynamic network flow approach," *Transportation Science*, vol. 34, no. 3, pp. 239–255, Aug 2000.
- [3] P. K. Menon, G. D. Sweriduk, and K. Bilimoria, "A new approach for modeling, analysis and control of air traffic flow," in *Proceedings of the AIAA Conference on Guidance, Navigation and Control*, August 2002.
- [4] H. P. Young, *Equity in Theory and in Practice*. Princeton, NJ: Princeton University Press, 1993.
- [5] W. Thomson, "A study of choice correspondences counterspeculation, auctions, and competitive sealed tenders," *Journal of Finance*, vol. XVI, 1988.
- [6] A. M. Bayen, R. L. Raffard, and C. J. Tomlin, "Eulerian network model of air traffic flow in congested areas," in *American Control Conference*, December 2004.
- [7] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, England: Cambridge University Press, 2004.
- [8] A. Jameson, "Aerodynamic design via control theory," *J. of Scientific Computing*, vol. 3, pp. 233–260, November 1988.