Lump-Sum Markets for Air Traffic Flow Control with Competitive Airlines

Steven L. Waslander, *Member, IEEE,* Kaushik Roy, *Student Member, IEEE,* Ramesh Johari, *Member, IEEE,* Claire J. Tomlin, *Senior Member, IEEE*

***Abstract*—Air traffic flow control during adverse weather conditions is managed by the FAA in today's air traffic system, although it is the individual airlines that are in the best position to assess the costs of disruptions to scheduled operations. To improve the efficiency of resource allocation, a market mechanism is proposed which enables airlines to participate directly in the flow control decision making process. Since airlines can be expected to behave strategically, a lump-sum market mechanism is used for which existence of a Nash equilibrium and a bound on the worst-case efficiency loss have been shown for agents that anticipate the effects of their own bids on resource prices. The convergence properties of this mechanism are studied for a two-player game with linear utilities, which reveals that restricting the airline bid update step size can result in a wider range of stable bidding processes. The mechanism is then applied to an air traffic flow control scenario for multiple airports in the northeastern US, which demonstrates the feasibility of performing market-based resource allocation within the time horizon for reliable weather predictions.**

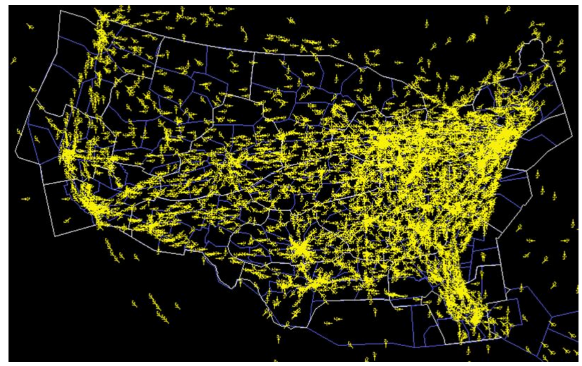

Fig. 1. Air traffic visualization for the US National Airspace, courtesy of NASA Ames and FACET [2].

I. Introduction

Air traffic management for the US National Airspace System (NAS) is one of the essential roles undertaken 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 disruptions result in significant financial burdens. Delays incurred by any individual flight can result in repercussions throughout the network, as 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 significant financial ramifications for the airlines.

S. Waslander is a Postdoctoral Scholar in the Department of Aeronautics and Astronautics, Stanford University, Stanford, CA. Email: stevenw@stanford.edu.

K. Roy is a Ph.D. Candidate in the Department of Electrical Engineering, Stanford University, Stanford, CA. Email: kroy1@stanford.edu.

R. Johari is an Assistant Professor in the Department of Management Science and Engineering and by courtesy in the Department of Electrical Engineering at Stanford University, Stanford, CA. Email: johari@stanford.edu.

C. Tomlin is a Professor in the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA, a Research Professor in the Department of Aeronautics and Astronautics and Director of the Hybrid Systems Laboratory at Stanford University, Stanford, CA. Email: tomlin@stanford.edu.

The FAA has three main avenues for reconfiguring scheduled operations in the face of airspace disruptions: ground delay programs (GDPs), miles-in-trail restrictions and flight reroutes. GDPs can be used to delay network traffic bound for specific constrained resources, such as destination airports and en route airspace. Miles-in-trail restrictions specify en route spacing on crowded airways to smooth flow into intersections and congested areas. Finally, flight reroutes ensure aircraft are redirected around weather disruptions and are managed through a national *playbook* of fixed procedures developed through years of experience. Currently, these methods address multiple stresses to the system individually. NAS-wide coordination of these actions is managed manually at the FAA Air Traffic Control System Command Center (ATCSCC), with verbal input from members of the Collaborative Decision Making (CDM) initiative [3], which currently enables a common situational awareness as well as airline participation in GDPs. The CDM initiative is especially important as a model program because it is the first for which incentives were designed to enhance airline participation [4], since rationing resources by schedule eliminates the penalty formerly faced by airlines when providing accurate information about flight changes.

As demand for resources increases, however, the need for the ATCSCC to provide efficient coordinated national flow control decisions will only increase, and the research community has therefore proposed a variety of air traffic modeling and control techniques that enable NAS level planning in real time [5], [6], [7], [8]. By aggregating individual flight plans into flows over a fixed network and defining capacity limits in both en route and terminal areas in terms of flow volumes, significant reductions in computational complexity are achieved over more detailed models that manage all flights individually.

Though the ability to synthesize trajectories for each aircraft is lost, appropriate traffic levels for sectors are computed to meet regional demands due to congestion and weather. This approach is in line with the hierarchical structure of the NAS, in which the discrete problem of aircraft separation is handled by sector controllers, while air traffic management at the center level often deals primarily with regional flow concerns. Another issue arises due to the fact that the costs incurred when flights must be delayed are primarily borne by the airlines, which are in the best position to assess the relative value of favoring one set of flights over another. Recent advances have focused on long-term allocation of airport slot resources [9], [10], or specifically on the GDP process [11], but do not consider disruptions to daily operations due to weather. To achieve efficient utilization of airspace resources, up to date private airline cost information must be included in the traffic flow decision making process, a need that is complicated by the competitive nature of airline operations.

In many engineering applications, such as Internet routing and electricity generation, market mechanisms have been proposed as a way to incorporate preference information into resource allocation decisions with competing agents [12], [13]. Market mechanisms assign an explicit price to each resource and agents balance the cost of purchasing resources with the value of those resources to them [14]. At each step in the process, the agents seek to minimize their costs by purchasing resources at market prices, after which the central coordinator updates resource prices in order to penalize excess demand. In previous work [8], a network flow model and market mechanism were presented for the air traffic control system, allocating airspace resources on a three-hour time horizon. The resources include take off, en route and landing traffic flows, and are allocated based on airline bids and centrally defined prices. Such a system could be implemented as a natural extension to the existing Collaborative Decision Making communications framework by requiring airlines to submit bids for traffic flow in areas affected by capacity restrictions due to weather.

Analysis of market mechanisms often relies on the assumption that agents treat prices as exogenous. For large numbers of agents with limited market share, such an assumption is reasonable. Unfortunately, applications with a small number of agents that command significant fractions of the available resources, such as the air traffic control problem, are susceptible to price manipulation by the agents. It becomes implausible to accept the assumption that agents would act without knowledge of how their actions affect market prices. Instead, agents must be thought of as *price-anticipating*, which is to say, capable of predicting the effect of changes in their own requests on the market. This can lead to arbitrary efficiency loss for certain market mechanisms [15], [16] and requires the definition of an alternate mechanism for which such a negative result can be avoided.

One such mechanism exists in the form of the lump-sum market, which was recently studied in the context of Internet congestion control [17], [18]. Instead of having agents request resource allocations directly, the lump-sum market, as its name suggests, requires agents to specify lump-sum payments they are willing to make for each resource, resulting in an elastic demand which can always be satisfied by adjusting prices relative to the cumulative payment per resource. Existence and uniqueness results for a Nash equilibrium have been established for price-anticipating agents in the single resource case [19], and existence of Nash equilibria in the case of multiple network resources [18]. What makes this market mechanism appealing is that a proportional bound on efficiency loss relative to the central solution has also been demonstrated for flow networks [18], limiting the possible losses that occur if agents act as price-anticipators.

The convergence properties of the lump-sum market are not well understood, but this work presents one set of preliminary results in this area. The two-player game for agents with linear utilities is considered in detail, and convergence properties are presented for continuous steepest ascent dynamics and for discrete best response dynamics. This investigation reveals marked differences between the convergence properties of the continuous and discrete dynamics, a fact that is not immediately obvious, as the discrete dynamics are not merely defined as discretized versions of the continuous dynamic (for which it is well understood that step size adversely affects convergence). Rather, the discrete dynamic is defined as a direct update to an agent's optimal bid, or best response, if all other agents are held at their current bids. As the difference in the marginal utility of agents increases, convergence of this dynamic fails due to the fact that small changes in one bid result in large changes in the other agent's best response.

For the air traffic control problem, these results present a method for improving the convergence of lump-sum markets with multiple airlines, arbitrary convex cost functions and numerous resources. If iterative best response bidding shows oscillatory pricing behavior in a specific situation, it is likely that fixing a maximum bid change between rounds will improve convergence, by mimicking the effect of requiring continuous dynamic updates. In practice, iterative best response bidding between price-anticipating airlines was observed to converge quite quickly, as is demonstrated in a flow control scenario presented in Section V. With eight airports, four airlines and a three hour time horizon, the scenario is indicative of the problem size that can be solved in real-time on a standard desktop computer. The mechanism can also be scaled to manage larger networks if distributed computation platforms are considered for each of the airlines.

Market mechanisms exhibit known deficiencies when implemented for engineering applications. The mechanisms do not consider the effect of repeated daily operation of the market on the airline strategies, nor do they consider the possibility that airlines might take into account the multiple stages of the bidding process. It will be necessary to ensure that price fluctuations do not result in unsafe situations, although this concern is mitigated by the three hour planning horizon. There are also issues relating to the adoption of a novel resource allocation method which represents a paradigm shift in US air traffic management, as most enroute resources are currently free of charge and airport fees are not slot dependent. Despite these open issues, this work presents a first step toward obtaining cost information from strategic airlines and incorporating it into air traffic control decision making. Resolving these deficiencies for resource allocation in air traffic control remains an area of future work.

The paper proceeds as follows. Section II defines a single resource lump-sum market and restates the main known results concerning existence and efficiency of a Nash equilibrium. The convergence properties of the lump-sum market for two agents with linear utilities are then investigated in Section III, for both iterative best response play and continuous steepest ascent dynamics. In Section IV, the air traffic network flow resource allocation problem is restated for completeness, as first defined in [8], and the lump-sum mechanism definition is extended to the network flow problem, as described in [18]. Simulation results are presented for an air traffic network consisting of eight airports in the Northeastern US over a three hour horizon with inclement weather affecting the destination airport in New York. Finally, some of the difficulties inherent in implementing the proposed mechanism and areas of future work are described in Section VI.

II. Lump-Sum Market Mechanism

A. Lump-Sum Market for a Single Resource

Consider a system of a single resource with constrained quantity $C \in \mathbb{R}_+$. Competing for the resource are agents $j \in \{1, \ldots, J\} = \mathcal{J}$, who value their allocation, $x_j \in \mathbb{R}_+$ according to a utility function, $U_j(x_j) : \mathbb{R}_+ \to \mathbb{R}_+$. The agent utility functions are assumed to be both increasing and concave over the compact space defined by $\mathcal{X}_j = [0, C]$. A centralized formulation of the efficient allocation problem is defined in Problem P2.1.

Central Allocation Problem

$$\begin{aligned} \text{maximize} \quad & \sum_{j\in\mathcal{J}} U_j(x_j) \\ \text{subject to} \quad & \sum_{j\in\mathcal{J}} x_j \le C \end{aligned} \qquad \text{(P2.1)}$$

This convex program can readily be solved if all utility information is known, but, for systems with independent agents and a competitive environment, it is improbable that this information will be available to a central authority. Instead, market mechanisms can be employed to reach an agreement on resource allocation amongst agents without requiring full disclosure of utility functions.

A lump-sum market is defined in which each agent specifies a lump-sum payment, $w_j \in \mathbb{R}_+$, for the resource. This results in an elastic demand that balances demand with capacity using the pricing rule

$$\lambda = \frac{\sum w_j}{C}. \qquad (1)$$

The resource price is defined as $\lambda \in \mathbb{R}_+$ and the notation $\sum w_j$ is used when summing over all agents. An agent's lump-sum bid results in the allocation

$$x_j = \frac{w_j}{\sum w_j} C. \qquad (2)$$

In abstracting resource markets for analysis, it is essential to explicitly define models of agent information and behavior that accurately capture the likely actions of agents in the real system. For markets with large numbers of similarly sized participants, it is reasonable to assume that agents assume prices are beyond their influence. Such agents are said to be *price-taking*, and operate rationally with respect to the fact that they believe their bids have no effect on the market prices they experience. Restricting the price to be strictly positive (i.e. requiring that the resource is contested), a price-taking agent in a lump-sum market optimizes the following local utility,

$$U_j^{PT}(w_j; \lambda) = U_j\left(\frac{w_j}{\lambda}\right) - w_j. \qquad (3)$$

The notation for the arguments of the utility function, $U_j^{PT}(w_j; \lambda)$, employs a semi-colon to distinguish between arguments which are controlled by agent j, such as its bid, and those that are assumed by agent j to be determined exogenously, such as the price.

For markets with only a few participants, however, it is unreasonable to assume such naïvety on the part of competitive agents. Since the pricing rule for lump-sum markets can be assumed to be publicly available, it is likely that agents will anticipate the effect of updating their bids on market prices. This behavior is referred to as *price-anticipating*. A modified price-anticipating agent utility function is defined which incorporates the resource pricing rule and is expanded to account for the case in which all bids are zero.

$$U_j^{PA}(w_j; w_{-j}) = \begin{cases} U_j(\frac{w_j}{\sum w_j} C) - w_j & \text{if } w_j > 0 \\ U_j(0) & \text{otherwise} \end{cases} \qquad (4)$$

Again, for notational simplicity, let the summation, $\sum w_{-j}$ be over all agent bids other than agent j.

The appropriate equilibrium condition for a market with price-anticipating agents is that of the Nash equilibrium, where each agent prefers its current bid to any other, given that all other agents hold their bids fixed. A *Nash equilibrium* is defined as a vector of bids, $w^{NE} \ge 0$, such that

$$U_j^{PA}(w_j^{NE}; w_{-j}^{NE}) \ge U_j^{PA}(w_j; w_{-j}^{NE}), \quad \forall w_j \in \mathbb{R}_+. \qquad (5)$$

B. Equilibrium Properties for a Single Resource

Two key results in the analysis of lump-sum markets with price-anticipating agents are restated here, which establish existence and uniqueness of the Nash equilibrium as well as a bounded efficiency loss relative to the optimal solution to the Central Allocation Problem of no more than 25%.

The following theorem shows that a solution to the game exists in the form of a unique Nash equilibrium.

***Theorem** 1 (Existence, Uniqueness [19]): Assume that for $J > 1$, each agent j has a concave, increasing and continuously differentiable utility function, U_j. Then there exists a unique Nash equilibrium, $w^{NE} \ge 0$, which is the unique solution to the Modified Central Allocation Problem,*

Modified Central Allocation Problem

$$\begin{aligned} \textit{maximize} \quad & \sum_{j\in\mathcal{J}} \tilde{U}_j(x_j) \\ \textit{subject to} \quad & \sum_{j\in\mathcal{J}} x_j \le C \end{aligned} \qquad \textit{(P2.2)}$$

where

$$\begin{aligned}\tilde{U}_j(x_j) = \quad & \left(1 - \tfrac{x_j}{C}\right) U_j(x_j) \\ & + \tfrac{x_j}{C}\left(\tfrac{1}{x_j}\int_0^{x_j} U_j(y)dy\right), \quad \forall j \in \mathcal{J}\end{aligned} \tag{6}$$

and the allocation is defined by Equation (2).

Even though the modified utility function $\tilde{U}_j$ does not have a clear interpretation, it is a useful construct in that it enables the proof of existence and uniqueness of the Nash equilibrium. The following theorem ensures that efficiency loss in the game is bounded.

***Theorem** 2 (*Efficiency Loss [18]*): Let x^S be the unique solution to the Central Allocation Problem and x^G be the unique solution to the game, $(U_1^{PA}, \ldots, U_J^{PA})$ and hence to the Central Allocation Game. Then*

$$\sum_{j\in\mathcal{J}} U_j(x_j^G) \geq \frac{3}{4}\sum_{j\in\mathcal{J}} U_j(x_j^S). \tag{7}$$

Furthermore, for any set of agents, $\mathcal{J}$, there exist linear utilities for which this bound is tight.

This result is the main reason that lump-sum markets have recently been proposed as an alternate to resource markets for network resource allocation problems [18].

III. Convergence Properties of Lump-Sum Markets

In order for a mechanism such as the lump-sum market to be useful in an engineering application such as air traffic control, its convergence properties must be understood. Convergence results for best response dynamics are typically difficult to achieve, although some positive results have been established [20], [21]. By focusing on the two-player game with linear utilities, it is possible to gain some insight into the convergence properties of the lump-sum mechanism in general. To this end, the utility function for each price-anticipating agent is defined as

$$U_j^{PA}(w_j; w_{-j}) = c_j \frac{w_j}{\sum w_j} C - w_j \,,\; j = \{1, 2\}.$$

where c_j is agent j's constant marginal rate of return for additional units of the resource. Throughout this section, the symmetry of the problem is exploited to simplify notation by assuming that $c_2/c_1 \geq 1$, without loss of generality.

A. Best Response Functions

The best response functions for each agent can be analytically computed for linear utilities by solving the first order necessary conditions for optimality of agent j's optimization problem.

$$\frac{\partial}{\partial w_j} U_j^{PA} = \frac{c_j \sum w_{-j}}{(\sum w_j)^2} - 1 = 0\,, \quad \sum w_{-j} > 0. \tag{8}$$

Let $\mathcal{W} = \mathbb{R}_+^{N-1} \backslash \{0\}^{N-1}$. Solving Equation (8) for w_j yields the agent's best response function, $\beta_j : \mathcal{W} \to \mathbb{R}_+$,

$$\beta_j\left(w_{-j}\right) = \max\left(0, \sqrt{c_j \sum w_{-j}} - \sum w_{-j}\right), \tag{9}$$

which is defined for all values of $\sum w_{-j} > 0$. If the sum of the other agents' bids is zero, the best response function is undefined, as any positive bid will capture all available resources. Figure 2 presents the best response functions for two agents.

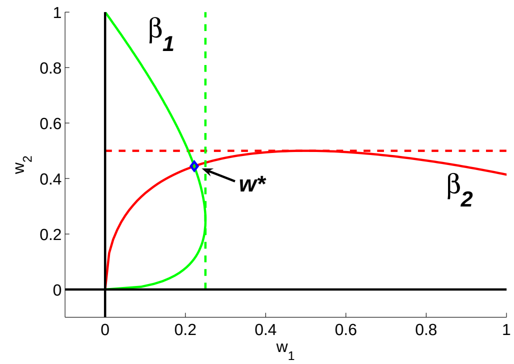

Fig. 2. Best response functions for two agents with utility slopes $c_1 = 1$ and $c_2 = 2$. Solid curves are best response functions, and dashed lines indicate maximum bids, with agent 1 in green (light gray), and agent 2 in red (dark gray). The equilibrium is denoted w^*.

B. Nash Equilibrium

At the Nash equilibrium of the game, denoted w^*, both agents bid their best responses: $w_1^* = \beta_1(w_2^*)$ and $w_2^* = \beta_2(w_1^*)$. The equilibrium can be found by taking the composition mapping $\beta_1(\beta_2(w_1))$ and setting it equal to w_1, to obtain

$$\begin{aligned} w_1 &= \beta_1(\beta_2(w_1)), \\ &= \sqrt{c_1(\sqrt{c_2 w_1} - w_1)} - (\sqrt{c_2 w_1} - w_1) \end{aligned} \tag{10}$$

which results in

$$c_1^2(c_2 w_1) = (c_1 + c_2)^2 w_1^2. \tag{11}$$

The equilibrium of the system occurs at

$$(w_1^*, w_2^*) = \left(\frac{c_1^2 c_2}{(c_1 + c_2)^2}, \frac{c_1 c_2^2}{(c_1 + c_2)^2}\right). \tag{12}$$

C. Discrete Dynamics

The discrete dynamic can now be defined by setting the order in which agents respond to changes in each others' bids, be it sequentially, concurrently or randomly. Let $k \in \mathbb{N}$ index the time step, then the concurrent discrete dynamic, $w(k+1) = \beta^c(w(k))$ is defined for each agent as

$$w_j(k+1) = \beta_j^c(w(k)) = \beta_j(w_{-j}(k))\,, \quad \forall\, j \in \mathcal{J}. \tag{13}$$

The sequential discrete dynamic $w(k+1) = \beta^s(w(k))$ is defined as

$$\begin{aligned} w_j(k+1) &= \beta_j^s(w(k)), \\ &= \begin{cases} \beta_j(w_{-j}(k)) & \text{if } k+1 \bmod j = 0, \\ w_j(k) & \text{otherwise.} \end{cases} \end{aligned} \tag{14}$$

Examples of the sequential update for the two-player game can be seen in Figure 3 for varying values of agent utilities. Note that as Agent 2's utility per unit resource increases relative to Agent 1's, the sequential update dynamics become less stable, resulting in limit cycle behavior in Figure 3 c) and instability in Figure 3 d).

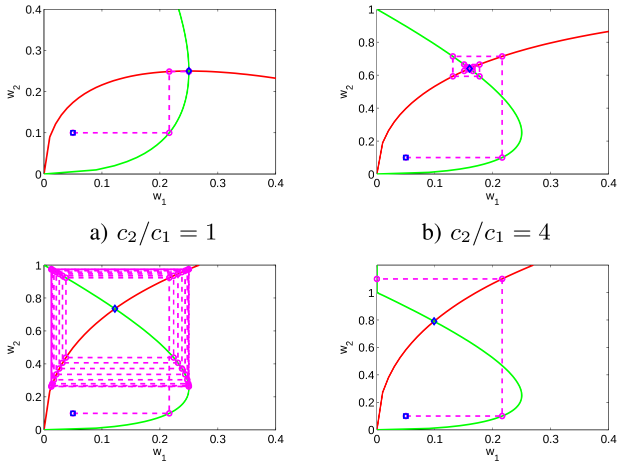

a) $c_2/c_1 = 1$ b) $c_2/c_1 = 4$

c) $c_2/c_1 = 6$ d) $c_2/c_1 = 8$

Fig. 3. Trajectories of sequential discrete dynamics for utility slope ratios $c_2/c_1 = 1, 4, 6, 8$. Best response curves are solid lines, discrete trajectories are dashed lines, the initial bids are denoted by a square, and the equilibrium is denoted by a diamond.

1) Local Stability: Local stability of the equilibrium can be determined by investigating the eigenvalues of the linearized dynamics, by the Hartman-Grobman Theorem, stated here in its original continuous form.

***Theorem** 3 (*Hartman-Grobman [22]*): If the Jacobian linearization, $Df(w)$, of the system at the equilibrium, w^*, has no zero or purely imaginary eigenvalues, then there exists a neighborhood $U \in \mathbb{R}^n$ and a homeomorphism $h : U \to \mathbb{R}^n$, which maps trajectories from the system to its linearization. In particular, $h(w^*) = 0$, and the homeomorphism can be chosen to preserve the parameterization by time.*

As a result, stability of the equilibrium with linearized dynamics applies to the nonlinear dynamics over some local region about the equilibrium. If the linearized dynamics have eigenvalues in the left half plane, the nonlinear dynamics are locally exponentially stable for an open neighborhood of the equilibrium. Similarly, if any eigenvalue of the linearized dynamics is in the right half plane, the nonlinear equilibrium is unstable. An analogous theorem exists for the discrete case [22], and local stability can be guaranteed by ensuring eigenvalues of the linearized discrete system are located inside the unit circle.

Starting with the concurrent discrete dynamic, the linearized dynamics about w^* are

$$\begin{aligned}(w(k+1) - w^*) &= D\beta^c(w)\,|_{w^*}\,(w(k) - w^*),\\ &= \begin{bmatrix} 0 & \dfrac{c_1 - c_2}{2c_2} \\ \dfrac{c_2 - c_1}{2c_1} & 0 \end{bmatrix} (w(k) - w^*),\end{aligned} \tag{15}$$

with eigenvalues

$$\lambda = \begin{bmatrix} \dfrac{c_1 - c_2}{2\sqrt{-c_2 c_1}} \\ \dfrac{c_2 - c_1}{2\sqrt{-c_2 c_1}} \end{bmatrix}, \tag{16}$$

which remain inside the unit circle for ratios of $c_2/c_1 < 3 + 2\sqrt{2}$. Above this ratio, the equilibrium is locally unstable. Note that the sequential update procedure results in the same linearization as the concurrent update, as can be demonstrated by considering the composition of two sequential updates.

2) Convergence: Depending on the ratio of utility slopes c_2/c_1, the discrete dynamics evolve in a number of manners. This section first identifies the ratios for which convergence is possible, guarantees convergence for a subset of these ratios, and then describes the behavior that results above these ratios.

The maximum bid of either agent is found by maximizing its best response function over all bids by the other agent. Setting the best response function derivative to zero for agent 1 yields

$$\begin{aligned}\frac{\partial \beta_1}{\partial w_2} &= \frac{\sqrt{c_1}}{2\sqrt{w_2}} - 1 = 0,\\ w_2^{max} &= c_1/4,\\ \beta_1(w_2^{max}) &= c_1/4.\end{aligned} \tag{17}$$

The maximum bid by agent 1 occurs when agent 2 bids $c_1/4$ and is also $c_1/4$. Convergence to the equilibrium is shown for all feasible initial bids, $w_0 \in \Lambda = (0, c_1/4] \times (0, c_2/4]$, for a limited range of utility slope ratios.

***Theorem** 4: Given utility slope ratios of $c_2/c_1 \in [1, 9/4)$, dynamics defined by Equation (14) and any initial bid pair $w_0 \in \Lambda$, the system converges to the equilibrium*

$$w^* = \left(\frac{c_1^2 c_2}{(c_1 + c_2)^2}, \frac{c_1 c_2^2}{(c_1 + c_2)^2} \right).$$

Please refer to the Appendix for the details of the proof.

For ratios of $c_2/c_1 > 3 + 2\sqrt{2}$, it is possible to demonstrate that the best response function of agent 2, $\beta_2(w_1)$, is positive for all $w_1 \in (0, w_1^{max}]$. This ensures that the only games in which either agent will post a bid value of zero occur when agent 2 bids above the maximum value at which agent 1 will participate. Setting $\beta_1(w_2) = 0$ yields $w_1^{pmax} = c_1$ as that maximum value, and then setting $\beta_2(w_1^{max}) = w_1^{pmax}$ and solving for the ratio c_2/c_1 gives a limit of $c_2/c_1 = 25/4$. Therefore, for all $c_2/c_1 > 25/4$, the discrete dynamic outcome may be undefined, and for all $c_2/c_1 \in (3 + 2\sqrt{2}, 25/4)$, the system is not able to converge to the equilibrium, as it is locally unstable, nor is it able to exit the region of feasible bids, $(0, c_1/4] \times (0, c_2/4]$. Instead, all trajectories must continue throughout the feasible region as t goes to infinity.

In summary, the two-player game with discrete dynamics can be subdivided into four types based on the ratio of utility slopes, c_2/c_1.

- For ratios of $c_2/c_1 < 3 + 2\sqrt{2}$, the equilibrium is locally stable, and convergence to the equilibrium can be guaranteed for all initial values of $w \in (0, c_1] \times (0, c_2]$ and utility slopes $c_2/c_1 < 9/4$.
- For slopes $9/4 \le c_2/c_1 < 3 + 2\sqrt{2}$, no guarantee is provided, although it is our conjecture that all trajectories of the discrete dynamics do converge to the equilibrium.
- If the ratio of $c_2/c_1 \ge 3 + 2\sqrt{2}$, two further possibilities exist. First, for ratios of $3 + 2\sqrt{2} < c_2/c_1 < 25/4$, convergence to the equilibrium is not possible and a limit cycle occurs.
- Then, for ratios $c_2/c_1 > 25/4$, the possibility arises that one agent is forced to bid zero, at which point the best

response function for the other agent is undefined, and the outcome is undefined.

Returning to Figures 3 a)-d), it can be seen that in scenarios a)-b), the discrete dynamic achieves the equilibrium, in c), a limit cycle results, and in d), a zero bid is placed by agent 1, resulting in an undefined outcome. As expected, the ratios of c_2/c_1 for which each of these behaviors occurs fall into the regions described in this section.

This result is interesting from the point of view of applying the mechanism to the air traffic flow control problem, as it demonstrates that even with only two price-anticipating airlines it is possible that iterative best response bidding may not converge to an equilibrium, leaving controllers without a traffic flow solution to implement. The following section performs a similar analysis for a continuous airline bidding update model with more favorable results.

D. Continuous Dynamics

Instead of allowing an agent to bid its best response directly, it is possible to define continuous dynamics for the game whereby each agent seeks to locally improve in the direction of steepest ascent of its value functions. Additionally, we can define an aggressiveness factor for each agent, κ_j, which sets the rate at which an agent reacts to its value function gradient. The continuous dynamic in terms of bids is defined as

$$\begin{aligned} \dot{w}_j(t) &= f_j(w(t)) = \kappa_j \left(\frac{\partial}{\partial w_j} v_j \left(w_j(t); \sum w_{-j}(t) \right) \right), \\ &= \kappa_j \left(\frac{c_j \sum w_{-j}(t)}{(\sum w_j(t))^2} - 1 \right), \qquad (18) \end{aligned}$$

which is well-defined for all $\sum w_j > 0$, and is defined as ∞ for $\sum w_j = 0$. Recall also that $w_j \geq 0$, which introduces a discontinuity in the definition. The dynamics for the complete system are summarized as $\dot{w} = f(w)$, and example trajectories are presented in Figure 4 for the two-player game for varying levels of aggressiveness.

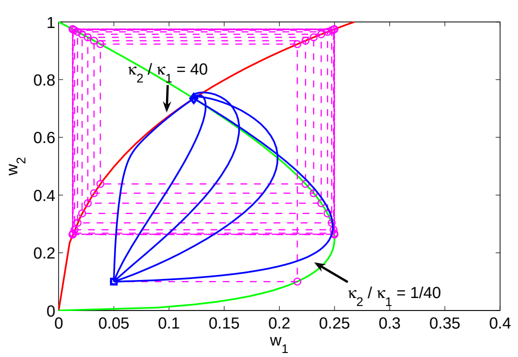

Fig. 4. Trajectories of discrete sequential and continuous dynamics for utility slope ratio $c_2/c_1 = 6$. Continuous trajectories are presented for aggressiveness ratios varying from $\kappa_2/\kappa_1 = \{1/40, 1/3, 1, 3, 40\}$. Best response curves are light solid lines, discrete trajectories are dashed lines, continuous trajectories are dark solid lines, the initial bids are denoted by a square, and the equilibrium is denoted by a diamond.

1) Local Stability: The linearized continuous dynamics are $\dot{w} \approx Df(w)|_{w*}(w - w^*)$. At $w^* = \left(\frac{c_1^2 c_2}{(c_1+c_2)^2}, \frac{c_1 c_2^2}{(c_1+c_2)^2} \right)$, the eigenvalues of the linearized system can be shown to be negative for all values of $\kappa_1, \kappa_2, c_1, c_2 > 0$, but are omitted for space considerations.

2) Convergence of Continuous Dynamic: To demonstrate convergence of the continuous dynamic, a series of definitions and theoretical results are required. This section relies on an established body of work in nonlinear analysis [22].

First, a trajectory is defined as $\gamma(t, w_0)$ where $\gamma : \mathbb{R}_+ \times \mathbb{R}^2 \to \mathbb{R}^2$, $t \in \mathbb{R}_+$ and $w_0 \in \mathbb{R}^2$. An *invariant set*, $\mathcal{S} \subseteq \mathbb{R}^2$, is one for which no trajectory that starts in $\mathcal{S}$ at $t_0 = 0$ exits the set $\mathcal{S}$. A point w is an *ω-limit point* of a trajectory γ if there exists a sequence of times t_n such that $t_n \to \infty$ as $n \to \infty$ for which $\lim_{n\to\infty} \gamma(t_n, w_0) = w$, and can be interpreted as any point that is visited infinitely often. An *ω-limit set*, $\omega(w_0)$, of point w_0 is the set of all ω-limit points of the trajectory $\gamma(t, w_0)$. *Closed orbits* are defined as periodic trajectories that do not contain any equilibria. Finally, *heteroclinic (resp. homoclinic) orbits* are defined as the union of multiple equilibria (resp. a single equilibrium) and trajectories connecting them.

To complete this brief overview of the necessary concepts, two well known results are restated. The first, known as Andronov's theorem, further characterizes ω-limit points.

***Theorem** 5 (*Andronov's Theorem [23]*): Every ω-limit point must be either an equilibrium, an element of a closed orbit, or an element of a heteroclinic or homoclinic orbit.*

Next, Bendixson's theorem can be used to determine the lack of existence of closed orbits inside a simply connected region.

***Theorem** 6 (*Bendixson's Theorem [22]*): Suppose $D \subseteq \mathbb{R}^2$ is a simply connected region such that the divergence, $div(f) = \frac{\partial f_1}{\partial w_1} + \frac{\partial f_2}{\partial w_2}$ is not identically zero in D and does not change sign anywhere in D. Then D contains no closed orbits of $\dot{w} = f(w)$.*

It is possible to demonstrate, in three steps, convergence of the continuous dynamic defined in Section II-A over a set of initial conditions. The first step is to demonstrate that the set of feasible initial conditions defines an invariant set for any values of κ_1, κ_2 and for values of $c_2/c_1 < 25/4$. The second step involves showing that closed orbits do not exist inside the region of interest, which is done using Bendixson's theorem for planar systems. Finally, it is possible to demonstrate that the equilibrium is the unique element of the ω-limit point for all points inside the invariant set. By Andronov's theorem, it is known that all ω-limit points must be either equilibria, closed orbits or heteroclinic/homoclinic orbits. Closed orbits can be eliminated by use of Bendixson's theorem, heteroclinic orbits by the fact that only one equilibrium exists inside the invariant region, and homoclinic orbits by the local stability of the equilibrium. This guarantees that all trajectories must eventually reach the equilibrium. Since the equilibrium is locally stable, a trajectory cannot leave the equilibrium once it has been reached, and convergence is ensured.

An invariant set for the continuous dynamics can be defined by the set of feasible initial conditions, which is all points where both agents bid no more than the maximum of their best response functions, $(w_1, w_2) \in [0, c_1/4] \times [0, c_2/4]$. For cases where $c_2/c_1 > 4$, the region can be further restricted to ensure that neither agent is priced out of the market on the initial bid, namely that $w_2 \in [0, c_1]$. The region under consideration can be defined by four separate curves, as depicted in Figure 5, labeled Ω, and defined in Table I.

Number	Curve	Normal
g_1	$w_1 = 0$ $w_2 \in [0, \min(\frac{c_2}{4}, c_1)]$	$\{-1, 0\}$
g_2	$w_1 \in [0, \frac{c_1}{4}]$ $w_2 = \min(\frac{c_2}{4}, c_1)$	$\{0, 1\}$
g_3	$w_1 = \frac{c_1}{4}$ $w_2 \in [0, \min(\frac{c_2}{4}, c_1)]$	$\{1, 0\}$
g_4	$w_1 \in [0, \frac{c_1}{4}]$ $w_2 = 0$	$\{0, -1\}$

TABLE I
INVARIANT SET DEFINITION.

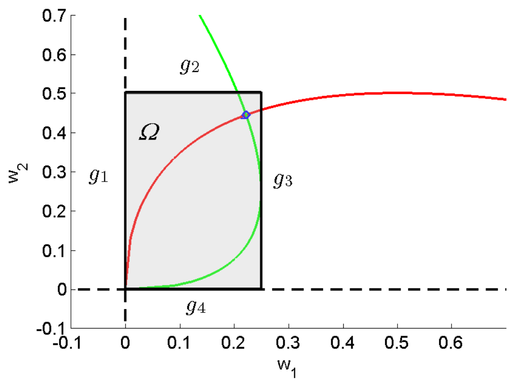

Fig. 5. Invariant region for utility slopes, $c_1 = 1$, $c_2 = 2$.

Lemma 1: The region Ω defined in Table I is an invariant set for the system defined in Equation (18).
Please refer to the Appendix for details of the proof.

It is now possible to state the main convergence result for continuous dynamics with two agents.

***Theorem** 7 (Convergence to Nash Equilibrium): Given utility slope ratios of $c_2/c_1 \in [4/25, 25/4]$, dynamics defined by Equation (18) and any initial bid pair $w_0 \in \Omega$, the system converges to the equilibrium*

$$w^* = \left(\frac{c_1^2 c_2}{(c_1 + c_2)^2}, \frac{c_1 c_2^2}{(c_1 + c_2)^2} \right).$$

Please refer to the Appendix for details of the proof.

Although this result is limited to utility slope ratios $c_2/c_1 < 25/4$, it is our conjecture that convergence holds beyond this threshold as well. The difficulty in extending the result lies in the discontinuity in the definition of the continuous dynamics when a zero bid is placed by either agent in reaction to a bid placed by one agent above the maximum at which the other agent will participate, which eliminates the ability to apply either Bendixson's or Andronov's theorem.

The analysis of convergence properties in this section reveals that the continuous update dynamics are stable for a much larger range of agents than iterative best response bidding. This result is in contrast with the desire of mechanism designers to allow the more natural mechanism definition where agents bid iteratively and react to current market prices. The lesson for the air traffic control application is therefore that if issues arise with limit cycles or zero bids in the iterative setting, convergence to the unique Nash equilibrium can be improved by approximating continuous ascent dynamics. In practice, this can be done by enforcing a maximum bid adjustment step size.

The methods employed for the convergence analysis rely heavily on the fact that only two agents were considered. In extending these results to more than two agents that exists for air traffic control, a novel approach such as a Lyapunov stability argument would most likely be needed. Nonetheless, the insight provided by the two agent scenario should be applicable to problems with more agents.

IV. APPLICATION TO AIR TRAFFIC FLOW CONTROL

The lump-sum market mechanism described above can be seen to satisfy three crucial properties required for the inclusion of airline preference information in the air traffic flow control problem. First, an equilibrium solution is known to exist even when airlines anticipate the effect of their bids on prices. Second, this equilibrium is known to exhibit a worst-case loss in efficiency of 25%, ensuring that strategic airline behavior will not result in arbitrarily poor system performance. Third, the convergence results from the previous sections reveal that it is possible to achieve this equilibrium through a distributed bidding process. Although it is true that other properties for this mechanism can also be considered essential, these remain in the realm of future work.

This mechanism is applied to an air traffic network flow model first defined in [8]. The model aggregates flow separately for each airline between origin and destination pairs, and seeks to avoid excessive densities of aircraft at any one resource (jetway or airport runway) at any specific point in the planning horizon, based on resource capacity estimates derived from weather forecasts. The flow control approach does not develop trajectories for individual aircraft but rather determines appropriate traffic levels for individual links to meet regional goals such as rerouting traffic flows around congestion or weather. This approach is reasonable in light of the hierarchical structure of the NAS, where more localized tasks of aircraft separation assurance are handled by TRACON and sector controllers, while regional concerns are in the realm of traffic flow management at the center level. Also, since aircraft operate most efficiently at their desired cruising speed, and since the flow abstraction necessarily reduces the model fidelity, the flows are assumed to operate at fixed velocities, which greatly reduces the computational complexity of the flow allocation problem.

It is important to note that other models exist for the NAS resource allocation problem [7], [24], which enable variable velocity flow modeling, or tracking of individual aircraft. This specific flow model was selected, however, for its convex and scalable formulation, which makes it ideally suited for use with

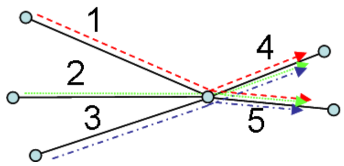

Fig. 6. A simple air traffic network with 5 links is used to illustrate the various components of the definition, where allowable flows are indicated by dashed and dotted lines. The source set is $\mathcal{S} = \{1, 2, 3\}$ and the sink set is $\mathcal{Z} = \{4, 5\}$. For this network, there are no merge sets, as 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 $1, 2$ and 3 can be split into links 4 and 5.

the lump-sum market mechanism and for application to large portions of the NAS. This section proceeds by first presenting an overview of the preferred flow modeling approach. The lump-sum mechanism is then adapted to the network flow allocation setting, following its definition in [18], which is followed by a discussion of the computational complexity of the optimization required of the airlines. The mechanism is then implemented on a multi-airport scenario in Section V.

A. Network Definition

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

A subset of links, $\mathcal{S} \subseteq \mathcal{I}$, is defined as the set of sources from which flow can originate and $\mathcal{Z} \subseteq \mathcal{I}$ is defined as the set of sinks from which flow exits the network. Each link $i \in \mathcal{I}$ has a set of links $\mathcal{M}_i$, whose entire flow merges into link i, and a set of links $\mathcal{D}_i$, whose members each receive a portion of flow from link i. Let $\mathcal{M} = \cup_{i \in \mathcal{I}} \mathcal{M}_i$ represent the set of all merging links and define $\mathcal{D}$ similarly. Then, for each diverging link $d \in \mathcal{D}$, there exists a non-empty set $\mathcal{F}_d$ of inflow links for which a portion of the flow is able to flow into link d. See Figure 6 for a small example of the elements of this common network definition.

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 simplified by discretizing time with time step δt, as $t \in \{0, \delta t, ..., T\} = \mathcal{T}$ 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 resources, $r = (i, t) \in \mathcal{R} = \mathcal{I} \times \mathcal{T}$, with $R = |\mathcal{R}|$ the total number of resources. Each resource has associated with it an en route capacity limit, summarized by an en route resource capacity limit vector, $C^e \in \mathbb{R}^R_+$. Once again, the assumption that these limits can be determined safely and in advance is implicit in the definition of the capacity of resources.

In the Discrete Path Flow model, airlines plan flights for specific origin-destination (OD) pairs, $(n_o, n_d) \in \mathcal{OD} \subseteq \mathcal{N} \times \mathcal{N}$. Each flight must follow a route over the network, which can be started at any point in the planning horizon (i.e. delayed arbitrarily). A route/time pair is referred to as a path, p, which can be represented in terms of the resources the flights will consume, and implicitly defines the speed at which the aircraft travel over the network, $p = (r^1, \ldots, r^{N_p})$, where $N_p \in \mathbb{N}$ is the length of the path. Note that each path is defined for a specific airline[1], so that $p \in \mathcal{P}_j$, $\mathcal{P} = \bigcup_{j \in \mathcal{J}} \mathcal{P}_j$, and $P = |\mathcal{P}|$. Each 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,

$$A^e(r, p) = \begin{cases} 1 & \text{path } p \text{ consumes resource } r, \\ 0 & \text{otherwise.} \end{cases} \tag{19}$$

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 $A^e y \leq C^e$.

Similarly, airport departure and arrival flow can be restricted to satisfy runway capacity limitations. Since multiple routes can be defined for each OD pair, airport arrival and departure flow can be determined by summing over the corresponding path flow allocations. Define the matrices $\Psi^o, \Psi^d, \in \{0, 1\}^{V \times P}$, to map any path flow allocation, y, to origin-destination pair inflows, $y^o \in \mathbb{R}^V_+$, and outflows, $y^d \in \mathbb{R}^V_+$, where $\mathcal{V} = \mathcal{OD} \times \mathcal{T}$ is the set of OD/time pairs, of cardinality V. Then,

$$\Psi^o(v, p) \quad = \quad \begin{cases} 1 & \text{path } p \text{ starts at OD/time pair } v, \\ 0 & \text{otherwise,} \end{cases} \tag{20}$$

$$\tag{21}$$

$y^o = \Psi^o y$ and y^d is defined similarly. The airport departure and arrival flow can be defined by summing over all OD pairs that arrive at or depart from a specific airport. Similar to Equation (19), let $A^o \in \{0, 1\}^{(|\mathcal{S}| \times N_T) \times V}$, $A^d \in \{0, 1\}^{(|\mathcal{Z}| \times N_T) \times V}$ be the origin and destination matrices, respectively. Then airport capacities are constrained by $A^o y^o \leq C^o$ and $A^d y^d \leq C^d$, where $C^o \in \mathbb{R}^{|\mathcal{S}| \times N_T}$ and $C^d \in \mathbb{R}^{|\mathcal{Z}| \times N_T}$ are the limits on safe airport operation.

C. Discrete Airline Model

In order to react strategically to unpredictable phenomena such as weather, it is necessary to define an airline cost function over possible outcomes. For airlines, aircraft arrival delays relative to the predefined flight schedule are an excellent indicator of the costs to be incurred by air traffic control decisions, and are in fact monitored by the FAA to assess air traffic control performance [1]. Therefore, the airline costs are modeled by minimizing the quadratic deviation from the

[1]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.

cumulative scheduled arrival flow, which penalizes not only deviations from the scheduled flow but also how long those deviations persist. This cost structure does not exactly penalize individual flight delays, however, as this information is lost in the flow aggregation process. The quadratic nature of the cost function does penalize large deviations from the flight schedule more heavily than small ones, which mimics the fact that larger disruptions tend to impact subsequent scheduled flights more severely, driving a domino effect of cost escalation as delays persist. 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, non-decreasing 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. Future work will aim to include the case of airline flight cancellations as well, which essentially define an outside option as resource prices increase.

To simplify notation, let $Y^o \in \mathbb{R}_+^V$ and $Y^d \in \mathbb{R}_+^V$ be the cumulative origin inflow and destination outflow for all OD/time pairs, respectively. Cumulative departure inflow is defined by $Y^o = \Gamma y^o$, where $\Gamma \in \{0,1\}^{V\times V}$ is a block diagonal matrix with lower triangular blocks of dimension N_T, and cumulative arrival inflow is defined similarly. Let y^{so} be the scheduled departure flow and y^{sd} be referred to as the scheduled arrival flow, with Y^{so} and Y^{sd} defined equivalently to Y^o and Y^d. A representative airline cost function can be defined as,

$$J_j(y_j) \quad = \quad ||\Omega_j \max\left(\mathbf{0}, Y_j^{sd} - Y_j^d\right)||_2^2, \tag{22}$$

where $\Omega_j \in \mathbb{R}_+^{V\times V}$ is a diagonal weighting matrix that captures each airline's preferences over OD pairs at each time epoch. Note that the zero vector of dimension V is represented by $\mathbf{0}$, and the $\max$ function is performed elementwise; that is, $\max(\mathbf{0},\cdot)$ operates on a vector such that positive entries are unchanged and all other entries are set to zero. Equation (22) defines an example of an airline cost function that is convex and non-increasing, since it is a positive linear combination of convex, non-increasing functions in the path flow variables, y_j.

D. Centralized Optimization

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

$$A = \begin{bmatrix} A^e \\ A^o \\ A^d \end{bmatrix}, \quad C = \begin{bmatrix} C^e \\ C^o \\ C^d \end{bmatrix}. \tag{23}$$

Note that the matrix A can be divided into $N_{\mathcal{J}}$ submatrices, $A = [A_1|\ldots|A_{N_{\mathcal{J}}}]$, which map airline path flow allocations, y_j, to the resources consumed. The air traffic path flow allocation problem can now be seen to take a similar form to the Central Allocation Problem defined in Problem P2.1,

Central Path Flow Allocation

$$\begin{array}{ll} \underset{y\geq 0}{\text{minimize}} & \sum_{j\in\mathcal{J}} J_j \\ \text{subject to} & Ay \leq C \end{array} \tag{P4.1}$$

whose 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 and SeDuMi in the Matlab environment [25], [26], [27].

E. Lump-Sum Market Mechanism for Air Traffic Flow Control

The definition of the lump-sum market for this problem formulation requires additional information of the bidders, in order to avoid situations where a Nash Equilibrium might not exist if demand for a resource is insufficient to keep prices away from zero [18]. Bidders are also required to submit a link flow request $\hat{x}_j \in \mathbb{R}_+^I$ which is the amount of flow they would like to have allocated when the price for a resource is zero. Let $z_j = (w_j, \hat{x}_j)$ be defined as the the *airline strategy*, with $z = (z_1, \ldots, z_J)$ the complete strategy vector. The flow allocation process is then augmented for each resource as follows,

$$x_{ij}(z) = \begin{cases} \dfrac{w_{ij}}{\sum\limits_{j\in\mathcal{J}} w_{ij}} C_i & \sum\limits_{j\in\mathcal{J}} w_{ij} > 0, \\ \hat{x}_{ij} & \sum\limits_{j\in\mathcal{J}} w_{ij} = 0\,, \sum\limits_{j\in\mathcal{J}} \hat{x}_{ij} \leq C_i, \\ 0 & \sum\limits_{j\in\mathcal{J}} w_{ij} = 0\,, \sum\limits_{j\in\mathcal{J}} \hat{x}_{ij} > C_i. \end{cases} \tag{24}$$

where $x_{ij}(z)$ is the allocation of flow on link i assigned to airline j resulting from the airline strategy z. Note that if the total requests for flow on a link exceed the link capacity but the sum over all bids on that link is zero, no airline receives any flow allocation as the link has become contested.

The local airline cost functions are defined as,

$$J_j^{PA}(z_j; z_{-j}) = J_j(y_j(x_j(z))) + \sum_{i\in\mathcal{I}} w_{ij}, \tag{25}$$

Note that the airlines not only determine the lump sum bids, w_j, to place for each resource, but also identify the path flows, y_j, which maximize the value of the flow allocations over each of the links, x_j, which are determined by the lump sum bids. In effect, the requirement to identify path flows in addition to resource bids represents an embedded network flow maximization problem in the argument of the airline cost.

The definition of a Nash equilibrium of Equation (5) is augmented to include the flow requests,

$$J_j^{PA}(z_j^{NE}; z_{-j}^{NE}) \leq J_j^{PA}(z_j; z_{-j}^{NE})\,, \quad \forall z_j \in \mathbb{R}_+. \tag{26}$$

The existence of a Nash equilibrium and its efficiency loss can now be established for the multi-link network problem. Because the strategy space of the airlines has been expanded to include link flow requests, the requirements on the costs can be relaxed to no longer require decreasing, differentiable costs, although the uniqueness of an equilibrium can no longer be guaranteed. In establishing these results for airlines with costs instead of utilities, the following assumption is required.

Assumption 1: The airline cost functions, J_j are assumed to be convex, non-increasing and continuous. Furthermore, the cost of a null allocation for each resource is assumed to be finite, and is referred as the worst case cost, $d_j = J_j(0)$.
As a result of the worst case cost function, the airline cost functions can be posed as utility functions, $U_j(x_j(z)) = d_j - J_j(x_j(z))$, which are nonnegative, concave, continuous and nondecreasing for all flow allocations, and evaluate to $U_j(0) = 0$ for a null allocation. This assumption allows for a modest generalization of the results of [18] to problems in which airlines measure preference in terms of cost, and for which null allocations do not result in zero costs/utilities.

***Theorem** 8 (*Existence [18]*): Assume that for $J > 1$, each airline j has a cost function, J_j that satisfies Assumption 1. Then there exists a Nash equilibrium, $z^{NE} \geq 0$ of the game defined by $(J_1^{PA}, \dots, J_J^{PA})$, if the allocation is defined by Equation (24).*

The efficiency loss bound for lump-sum markets also holds in the current context.

***Theorem** 9 (*Efficiency Loss [18]*): Let z^S be the unique efficient solution to the Central ATC Allocation problem and z^G be a Nash equilibrium of the game, $(J_1^{PA}, \dots, J_J^{PA})$. Then,*

$$\sum_{j\in\mathcal{J}} d_j - J_j(x_j^G(z^G)) \geq \frac{3}{4}\left(\sum_{j\in\mathcal{J}} d_j - J_j(x_j^S(z^S))\right). \quad (27)$$

Furthermore, for any set of airlines, $\mathcal{J}$, there exist linear utilities for which this bound is tight.

The local airline optimizations can now be defined for both price-taking and price-anticipating airlines. For airlines that act as price-takers, the airline subproblems are defined as

Price-Taking Airline Subproblem

$$\begin{aligned} &\text{minimize} && J_j(y_j) + \sum_{i\in\mathcal{I}} w_{ij} \\ &\text{subject to} && \lambda_i A_{i,j} y_j \leq w_{ij},\ \forall i \in \mathcal{I} \end{aligned} \quad \text{(P4.2)}$$

The notation $A_{i,j}$ denotes the i^{th} row of the j^{th} sub-matrix, A_j, of the matrix A defined in Equation (23). $A_{i,j}$ is a row vector of the same length as y_j, the path flow column vector.

For airlines that act as price-anticipators, a simplifying assumption is made that there exists a minimium bid, $\epsilon_w > 0$, for every resource, ensuring that only the first pricing rule in Equation (24) is applicable. The airline subproblems are defined as

Price-Anticipating Airline Subproblem

$$\begin{aligned} &\text{minimize} && J_j(y_j) + \sum_{i\in\mathcal{I}} w_{ij} \\ &\text{subject to} && w_{ij} \geq \epsilon_w \\ & && A_{i,j} y_j \leq C_i \left(\frac{w_{ij}}{\sum_{j\in\mathcal{J}} w_{ij}}\right),\ \forall i \in \mathcal{I} \end{aligned} \quad \text{(P4.3)}$$

In both cases, the airline subproblem is a convex optimization program and can therefore be solved efficiently with standard optimization tools. The airline cost functions are convex, and for Problem P4.2, the constraints are linear. The flow constraints in Problem P4.3 are also convex, which can be seen by rewriting them as

$$A_{i,j} y_{ij} \leq C_i \left(1 - \frac{\sum_{k\neq j} w_{ik}}{w_{ij} + \sum_{k\neq j} w_{ik}}\right), \quad (28)$$

and noting that for constant $a \in (0, \infty)$, $w \in (0, \infty)$, the function $f(w) = a/(w + a)$ is convex in w.

The lump-sum market mechanism is adapted for the air traffic control problem in Algorithm 1, where iterations are indexed by k.

Algorithm 1 Discrete dynamic algorithm for lump-sum markets.

Initial bids, z_j^0, submitted $\forall j \in \mathcal{J}$.
$k = 0$
repeat
  $\lambda^k = \frac{\sum w_j^k}{C}$, central price update.
  Allocations $x_j^k(z))$ are set as defined in Equation (24).
  Agents update their responses z_j^{k+1} by solving Problem P4.2 or P4.3.
until $\sum ||z_j^{k+1} - z_j^k|| < \epsilon_z$

F. Problem Size

The size of the air traffic flow control problem defined above depends primarily on the length of the planning horizon and the subset of the US NAS to be included. The US NAS contains 60 major airports[2] for which flow scheduling might be considered, which results in as many as 3,540 OD pairs. The total number of link resources can be approximated to be 2,500, by assuming approximately 7 links across each of 400 sectors that cover the US. For a path flow model with 5 alternate routes per OD pair and 5 minute time intervals, the flow allocation problem could reach sizes on the order of one million flow variables and one half million resource constraints. At this size, exact second order techniques are currently impractical and convex program solution times in the tens of minutes for each local optimization are likely on standard computing hardware. 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 Section V are representative of problem sizes for which real-time market based allocation converges in a realistic time frame for mitigating schedule disruptions due to weather.

[2]A major airport is defined as one that supports $> 0.25\%$ of total passenger enplanements in the US.

V. Simulation: Air Traffic Flow Control

A. Scenario

A scenario based on the north-eastern United States is used to investigate the relative performance of price-taking and price-anticipating airlines in the lump-sum market mechanism. For this scenario, flow departing from seven airports within approximately three hours of and destined for New York is modeled, as depicted in Figure 7 a). Once again, the simulation spans a three hour time horizon and uses a five minute time step. The nodes of the network are based on actual airport and airway intersection locations available from the FAA [28], 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 airlines are considered for this scenario, as summarized in Table II. They are defined to represent the various types of airlines in existence, including low-cost and mainline carriers, as well as large volume and small volume airlines. The schedule definitions for mainline carriers tend to be peaked, meaning that many flights are scheduled to arrive at their hub airports at the same time, putting a strain on the airport and increasing their cost of operations. In contrast, low-cost airlines try to maintain flat schedules which utilize gates and surface crews evenly, thereby reducing their costs.

	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

TABLE II
Scenario: Airline schedules and flow preferences.

En route capacity limitations are uniformly set to 90% of the peak scheduled flow along the busiest link, and 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 with any subsection of the link inside the weather polygon at each time step.

B. Results

Figures 7 b)-d) 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 again used to represent flow volume along a link. The weather disruption polygon and affected links and nodes are represented in red (dark gray).

Figure 7 b) shows the relatively low demand starting period of the simulation, where most flows are taking the shortest path routes to the destination, and none of the flow is close to the link capacity limits. Once again, weather has a significant impact on flow to the destination and in particular, Figure 7-d) shows a surge of recovery flow once the weather has passed.

Figure 8 presents the arrival traffic and pricing throughout the simulation for both price-taking and price-anticipating airlines, 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 recovery does not occur immediately, as not only was the destination airport affected by the weather disruption, but also the links leading to the destination and so recovery was stalled until flows along the final legs of the flight paths could again reach the airport.

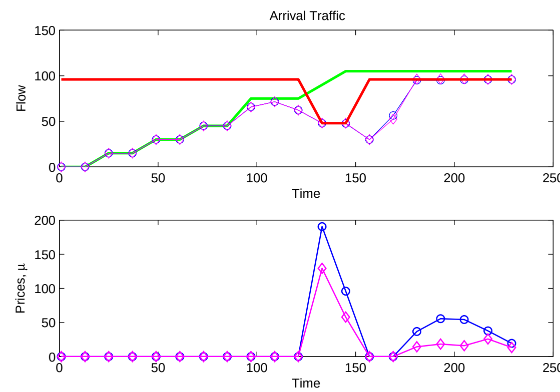

Fig. 8. Arrival flow at the destination airport. Top) Arrival flow quantities: Scheduled arrival traffic (light green line), airport capacity (dark red line), price-taking solution (dark blue o) and price-anticipating solution (light magenta ◇). Bottom) Airport pricing for price-taking solution (dark blue o), and price-anticipating solution (light magenta ◇).

The main distinction in the solutions of the price-taking airlines and the price-anticipating airlines is that the prices that result in the price-anticipating solution are lower than with price-taking airlines. This effect is caused by the airlines that command a significant portion of the available resources, which bid lower than they otherwise would in order to keep prices down. In the extreme with two price-anticipating airlines and a single resource, if one airline bids zero, the other is incentivized to bid as close to zero as possible, as it will still receive all the available resource regardless of the magnitude of its bid. By contrast, the price-taking airline would bid such that the price equals the negative of its marginal cost when it receives all the available resource, which is some finite non-negative value for strictly increasing, convex cost functions.

The deviation from the desired cumulative schedule is plotted for each airline in Figure 9, for both the price-taking and price-anticipating airlines. The high cost, high volume airline (Airline 1) has increased its deviations in the price-anticipating solution, whereas the other airlines (Airlines 2-4) have decreased their deviation. Essentially, the high cost, high volume carrier commands enough of the available resources to depress market prices to its advantage, and so incurs some additional delay as a result.

A comparison of price-taking and price-anticipating airline costs is presented in Table III, where columns represent the actual delay cost incurred, the resulting payments made by the airlines and the total market cost that results for each airline. Three interesting elements of the results are visible in this table. First, the overall efficiency of the solution is better with price-taking airlines vs. price-anticipating airlines, as visible from the total delay cost for each solution. This

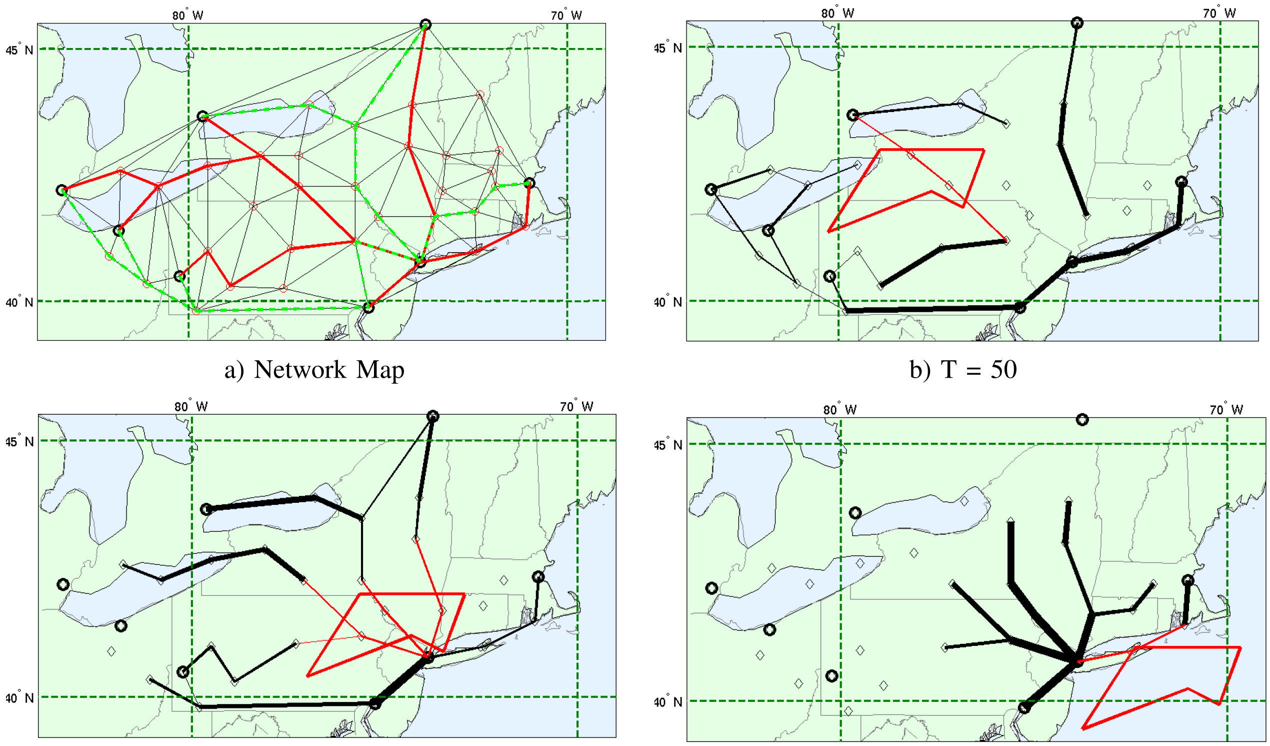

a) Network Map b) T = 50

c) T = 100 d) T = 150

Fig. 7. Aggregate flow of traffic to New York airports for all airlines. Each figure has airports marked with black circles, and all airway intersection nodes depicted with diamonds. Subfigure a) presents the network map. Subfigures b)-d) depict flow volume at times $t = 50, 100, 150$ minutes, and use link width to represent flow volume. The weather disruption polygon is shown in red (dark gray), as are all links and nodes affected by weather.

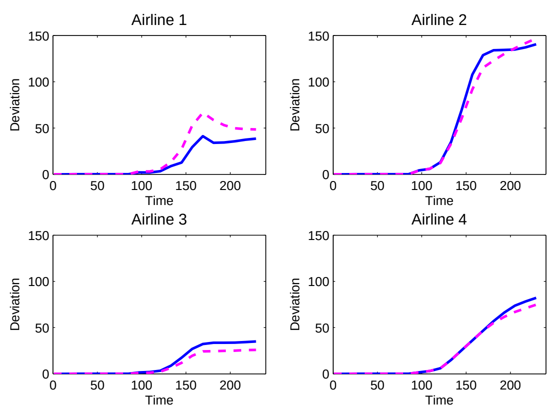

Fig. 9. Deviation from desired cumulative arrival flow: Price-taking solution (solid blue line) vs price-anticipating solution (dashed magenta line) for each airline. The high cost, high volume airline (Airline 1) incurs more delay as a price-anticipator and other airlines (Airlines 2-4) incur less delay.

Price-Taking Airlines			
	Delay Cost	Payment	Market Total
Airline 1	573	6442	7015
Airline 2	2026	3392	5418
Airline 3	504	848	1352
Airline 4	454	65	519
Total	3557	10747	14304

Price-Anticipating Airlines			
	Delay Cost	Payment	Market Total
Airline 1	1374	3835	5209
Airline 2	1846	2381	4227
Airline 3	268	741	1009
Airline 4	400	91	491
Total	3888	7048	10936

TABLE III
COST COMPARISON AND RESOURCE PAYMENTS FOR ALL AIRLINES IF ACTING AS PRICE-TAKERS OR PRICE-ANTICIPATORS.

follows from the fact that the lump-sum market with price-taking airlines converges to the efficient solution, whereas the Nash equilibrium for price-anticipating airlines need not be efficient. Second, since the resource prices are lower with price-anticipating airlines, so to are the payments they make (for all but Airline 4), and finally, all airlines prefer their market totals in the price-anticipating scenario over those in the price-taking scenario.

Finally, the market mechanism convergence performance can be observed in Figure 10. The slow rate of convergence of the market mechanism with price-taking airlines is a direct result of the pricing update mechanism, which must maintain a small step size to avoid excessive oscillation of market prices. Subgradient techniques can be used for the price update, but the algorithm can only be guaranteed to converge in an unbounded number of steps [29]. In contrast, with price-anticipating airlines, there is a significant improvement in the number of iterations required for the market prices to converge, which is indicative of the anticipatory effect that airlines have by including the effect of their bids on market prices in their local decision making. In practice, the local optimizations

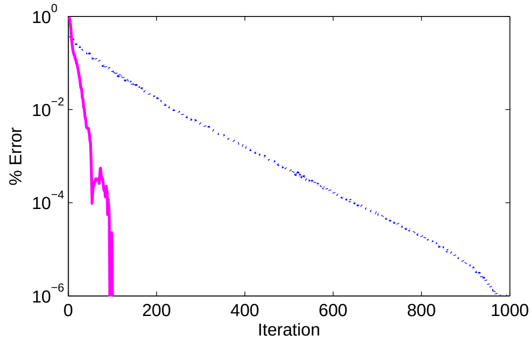

Fig. 10. Convergence of lump-sum market based flow allocation with the Discrete Path Flow model for both price-taking (dotted blue) and price-anticipating (solid magenta) airlines. Error as a percentage of the final cost.

involve convex nonlinear constraints on the airline bids, and therefore take longer to perform than the equivalent price-taking optimization, so some of the benefit of the faster market convergence is lost. For a market mechanism that requires hundreds of iterations to converge, local airline optimizations cannot consume longer than ten seconds if time constraints on the overall computation are to be met.

VI. Conclusion

In summary, this paper proposes the use of a lump-sum market mechanism for the allocation of resources amongst price-anticipating airlines competing for air traffic network flow, in order to improve system efficiency when weather disruptions affect capacity. This method allows for the fact that airlines are in direct competition and will attempt to gain an advantage over their competitors.

This work includes an investigation of convergence properties of the mechanism, to strengthen the justification for its implementation. By restricting the class of problems to two agents with linear utilities, both discrete best response and continuous steepest descent dynamics can be classified into regions of guaranteed stability. Most interestingly, the continuous dynamic is stable over a broader range of agent types than the discrete dynamic, which implies that for systems where convergence of the discrete best response dynamic is too slow or non-existent, a limited step size approach may improve the situation.

The effect of price-anticipation by the airlines on the market outcome is demonstrated for a representative flow problem, and both moderate efficiency loss combined with lower resource pricing result, and in fact all airlines prefer their outcome with this mechanism over the efficient market solution that results from price-taking strategies. Although no convergence guarantee exists for price-anticipating airlines on network flow problems, in practice, the convergence results were in fact significantly better when airlines predict the effect of their bids on prices.

A. Future Work

Many fascinating aspects of the implementation of market mechanisms for disruption management in air traffic control remain. The notion that airspace resources have inherent value, and that allocation to the highest bidder is the most desirable method for managing excess demand is certainly not universally accepted. Great care will need to be taken to address fears from players such as the general aviation community, but ultimately both safety and efficiency concerns will be advanced if busy routes and airports are preferentially reserved for larger vehicles with higher resource valuations. Similarly, safety concerns may result if convergence to feasible solutions cannot be guaranteed within the planning window, and further work is needed in defining market performance guarantees before their use in safety critical applications can be considered. The CDM initiative provides a useful benchmark for how to appropriately implement a market mechanism with airline participation. The ration by schedule approach for GDPs, which assigns delays based on the existing flight schedule, results in an incentive for airlines to provide accurate flight information, thereby reducing the frequency of unused landing slots. Although the mechanisms are different, the lump-sum market approach proposed in this work does not inherently give airlines resource rights implied by the existing flight schedule. It is important for future work to address this limitation of the lump-sum market theory to ensure various stakeholders have proper incentives and are satisfied with the equity of the approach to ensure participation in a coordinated flow control solution.

Specific to the market mechanism presented in this work, it is important to note that the price-anticipating airline model does not capture every possible action airlines might take to attempt to manipulate the allocation mechanism. Specifically, the repeated nature of the allocation problem allows many additional strategic behaviors to be executed by the airlines that are not captured in this model. Well known negative results demonstrate that every allocation can be sustained as an equilibrium in repeated games, through strategies that punish deviation over multiple epochs. Even in the single market case, each set of bids submitted by the airlines is evaluated separately in this model, without any notion that multiple bidding rounds are required for convergence. This is a strict limitation on the strategies considered for the airlines, and may therefore not capture the true outcome if the mechanism were implemented.

From the air traffic flow modeling perspective, further issues also remain. In particular, the conversion of the equilibrium flow solution into modifications to the existing aircraft flight plans has not been addressed. The definition of the network flow problem is inherently continuous, which is a requirement of the mechanism under consideration, whereas the existence of individual flights imposes an integer constraint on modifications. The result is that the flow solution may be difficult to achieve in practice, despite the fact that airlines are required to make payments for their flow allocations. In each case, these issues represent interesting areas of future research, with broad implications not only for air traffic management, but for many other multi-agent engineering applications.

References

[1] Department of Transportation, *Bureau of Transportation Statistics*,

2006. [Online]. Available: http://www.bts.gov/

[2] K. Bilimoria, B. Sridhar, G. Chatterji, K. S. Sheth, and S. Grabbe, “FACET: Future ATM concepts evaluation tool,” *Air Traffic Control Quarterly*, vol. 9, no. 1, pp. 1–20, March 2001. [Online]. Available: http://technology.arc.nasa.gov/SOY2006/SOY_FACET/

[3] CDM Working Group, *Collaborative Decision Making 2005 Structure*, May 2005, last viewed April 7th, 2008. [Online]. Available: http://cdm.fly.faa.gov/whatscdm/cdmdocs.html

[4] T. V. Vossen and M. O. Ball, “Slot trading opportunities in collaborative ground delay programs,” *Transportation Science*, vol. 40, no. 1, pp. 29–43, February 2006.

[5] D. Bertsimas and S. Stock Patterson, “The air traffic flow management problem with enroute capacities,” *Operations Research*, vol. 46, no. 3, pp. 406–422, May-June 1998.

[6] P. Menon, G. Sweriduk, and K. Bilimoria, “A new approach for modeling, anaylsis and control of air traffic flow,” in *Proceedings of the AIAA Guidance, Navigation and Control Conference and Exhibit*, Monterey, CA, August 2002.

[7] D. Sun, S. D. Yang, I. Strub, and A. M. Bayen, “Eulerian trilogy,” in *Proceedings of the AIAA Guidance, Navigation and Control Conference and Exhibit*, Keystone, CO, August 2006.

[8] S. L. Waslander, R. L. Raffard, and C. J. Tomlin, “Market-based air traffic flow control with competing airlines,” *AIAA Journal of Guidance, Dynamics and Control*, vol. 31, no. 1, pp. 148–161, Jan-Feb 2008.

[9] M. O. Ball, G. Donohue, and K. Hoffman, “Auctions for the Safe, Efficient and Equitable Allocation of Airspace System Resources,” in *Combinatorial Auctions*, Y. S. Cramton, P. and R. Steinberg, Eds. Cambridge, MA: MIT Press, 2006, ch. 22, pp. 507–538.

[10] A. Czerny, P. Forsyth, D. Gillen, and H.-M. Neimeier, *Airport Slots: International Experiences and Options for Reform*. Ashgate Publishing, 2008.

[11] T. Vossen and M. Ball, “Optimization and mediated bartering models for ground delay programs,” *Naval Research Logistics*, vol. 53, p. 75 90, 2006.

[12] M. Falkner, M. Devtsikiotis, and I. Lambadaris, “An overview of pricing concepts for broadband IP networks,” *IEEE Communications Surveys*, Second Quarter 2000.

[13] R. Wilson, “Architecture of power markets,” *Econometrica*, vol. 70, no. 4, pp. 1299–1340, July 2002.

[14] A. Mas-Colell, M. D. Whinston, and J. R. Green, *Microeconomic Theory*, 2nd ed. New York, NY: Oxford University Press, 1995.

[15] B. Hajek and S. Yang, “Strategic buyers in a sum-bid game for flat networks,” *IMA Workshop 6: Control and Pricing in Communication and Power Networks*, March 2004.

[16] R. Johari and J. N. Tsitsiklis, “A scalable network resource allocation mechanism with bounded efficiency loss,” *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 5, pp. 992–999, May 2006.

[17] F. Kelly, A. Maulloo, and D. Tan, “Rate control for communication networks: Shadow prices, proportional fairness and stability,” *Journal of the Operational Research Society*, vol. 49, no. 3, pp. 237–252, March 1998.

[18] R. Johari and J. N. Tsitsiklis, “Efficiency loss in a network resource allocation game,” *Mathematics of Operations Research*, vol. 29, no. 3, pp. 407–435, August 2004.

[19] B. Hajek and G. Gopalakrishnan, “Do greedy autonomous systems make for a sensible internet?” in *Proceedings of the Conference on Stochastic Networks*, Stanford University, CA, 2002.

[20] J. S. Shamma and G. Arslan, “Dynamic fictitious play, dynamic gradient play, and distributed convergence to Nash equilibria,” *IEEE Transactions on Automatic Control*, vol. 50, no. 3, pp. 312–327, March 2005.

[21] D. Fudenberg and D. K. Levine, *The Theory of Learning in Games*. Cambridge, MA: The MIT Press, 1998.

[22] S. Sastry, *Nonlinear Systems: Analysis, Stability, and Control*. New York, New York, USA: Springer Verlag, 1999.

[23] J. Guckenheimer and P. Holmes, *Nonlinear Oscillations, Dynamical Systems and Bifurcations of Vector Fields*. Berlin, Germany: Springer Verlag, 1983.

[24] 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, August 2000.

[25] M. Grant, S. Boyd, and Y. Ye, *cvx: Version 0.85*, 2006. [Online]. Available: http://www.stanford.edu/~boyd/cvx/

[26] J. F. Sturm, O. Romanko, and I. Polik, *SeDuMi: Self Dual Minimization Version 1.1*, 2006. [Online]. Available: http://sedumi.mcmaster.ca/

[27] Mathworks, *Matlab Version 7.1.0*, 2006. [Online]. Available: http://www.mathworks.com/products/matlab/

[28] Federal Aviation Administration, *Air Traffic Control System Command Center*, 2007. [Online]. Available: http://www.fly.faa.gov/

[29] N. Z. Shor, *Minimization Methods for non-differentiable functions*, ser. Springer Computational Mathematics: 3, translated from Russian by K.C. Kiwiel and A. Ruszczinski, Eds. New York, USA: Springer Verlag, 1985.

APPENDIX

First, the proof for Theorem 4 is presented.

Proof: First, there exists a region, Λ_0, of increasing bids, where both agents will always increase bids until Λ_0 is exited. This region can be identified by inspecting the best response functions and noting that $\beta_2(w_1)$ is increasing in w_1 throughout Λ, while $\beta_1(w_2)$ is increasing for $w_2 \in (0, c_1/4]$. Therefore, any bid pair starting in the region $(0, c_1/4] \times (0, c_1/4]$ must result in a trajectory that exits Λ_0 at some step k.

For the remainder of the feasible set, $\Lambda \backslash \Lambda_0 = (0, c_1/4] \times (c_1/4, c_2/4]$, the worst case slopes in terms of the magnitude of the eigenvalues of a linearized dynamic occur at $w_1 = \epsilon > 0$, where $\partial\beta_1/\partial w_2 \rightarrow -1/2$ as $\epsilon \rightarrow 0$ and $w_2 = \beta_2(c_1/4)$, where $\partial\beta_2/\partial w_1 = \sqrt{c_2/c_1} - 1$. The resulting eigenvalues of this worst case linearization are uniformly bounded inside the unit circle for all $c_2/c_1 < 9/4$. Therefore, all trajectories that start in the feasible set Λ converge to the equilibrium, w^*, for ratios of $c_2/c_1 < 9/4$ and $c_2/c_1 > 1$. ■

Next, the proof for Lemma 1 is presented.

Proof: Invariance is demonstrated by taking the dot product of the outward normal to the boundary and the dynamics, which is shown to be every non-positive. Along g_1,

$$\begin{aligned}(\dot{w}_1, \dot{w}_2) \cdot (-1, 0) &= -\dot{w}_1 \\ &= -\kappa_1 \left(\frac{c_1}{w_2} - 1 \right) \leq 0\,, \quad \forall\, w_2 \leq c_1.\end{aligned} \tag{29}$$

The remaining curves can be evaluated similarly. ■

Finally, the proof for Theorem 7 is presented.

Proof: The region Ω defined in Figure 5 is a simply connected region which is invariant to the flow defined in Equation (18) for ratios of $c_j/c_k \in [1, 25/4]$. Therefore, any trajectory starting from a point $w_0 \in \Omega$ cannot exit Ω and must spend infinite time in the region. By Andronov's theorem, every ω-limit point must be either an equilibrium, an element of a closed orbit or an element of a heteroclinic or homoclinic orbit. There exists only one equilibrium point in Ω, as shown in Section III-B, and it is stable, therefore heteroclinic and homoclinic orbits are not possible in Ω, and so all ω-limit points must be either elements of a closed orbit or the unique equilibrium. The existence of closed orbits can be ruled out using Bendixson's theorem since the divergence of the flow field is negative for all $w \in \Omega$:

$$\begin{aligned} div(f) &= -2 \left(\frac{\kappa_1 c_1 w_2}{(w_1 + w_2)^3} + \frac{\kappa_2 c_2 w_1}{(w_1 + w_2)^3} \right), \\ &< 0\,, \quad \forall \{w_1, w_2\} \in \Omega. \end{aligned} \tag{30}$$

Therefore, the unique ω-limit point in Ω is the equilibrium, and so all trajectories starting in Ω must converge to this equilibrium. Since the equilibrium is locally stable, a trajectory cannot leave the equilibrium once it has been reached. ■