The Effects of User Competition in Air Traffic Management Initiatives

Amy Kim
Civil and Environmental Engineering
University of Alberta
Edmonton, AB, Canada
amy.kim@ualberta.ca

Mark Hansen
Civil and Environmental Engineering
University of California, Berkeley
Berkeley, CA
mhansen@ce.berkeley.edu

Abstract—**The classic Air Traffic Flow Management literature has focused on the development of programs that typically aim to optimize total system-wide metrics in their resource allocation solutions. These system-optimal solutions are thought to be achievable because the airspace system is a highly controlled environment where aircraft ultimately cannot fly a given route without approval from the Air Navigation Service Provider (ANSP). We use a simple, classic traffic assignment principle to illustrate that if flight operators are asked to provide preference inputs that are used directly in a system-optimal resource allocation, a truly system-optimal solution may not be achievable. Because flight operators act in a way that is analogous to autonomous drivers attempting to minimize their individual expected travel costs, the actual solution may be closer to that of a user equilibrium. This finding has two major implications. Firstly, in ATFM programs such as CTOP where user preferences are considered, system-optimal solutions may be an unrealistic goal that cannot be attained. Secondly, resource allocation schemes that aim to provide greater equity by sacrificing efficiency may not be as inefficient as they first appear, when compared against schemes that aim to provide system-optimality.**

Keywords- Air traffic flow management (ATFM); Collaborative Trajectory Options Program (CTOP); en route resource allocation; strategic planning; Wardrop's principles; competitive behavior.

I. INTRODUCTION

The classic Air Traffic Flow Management (ATFM) literature has focused on the development of programs that have typically aimed to optimize system-wide metrics in their resource allocation solutions. These system-optimal (SO) solutions are thought to be achievable because of the fact that airspace is a highly controlled environment where aircraft ultimately cannot fly a given route without approval from the Air Navigation Service Provider (or ANSP, such as the FAA or EUROCONTROL). Flight operators cannot make autonomous routing decisions to minimize their travel costs, with respect to the time of travel and route taken, particularly under ATFM programs initiated to mitigate delays during times of heavy congestion. From this standpoint, the aviation environment is significantly different from road traffic operations, in that drivers are usually free to make autonomous travel choices in order to minimize their own individual travel costs. Under these conditions, Wardrop's first equilibrium condition – the user equilibrium (UE) – is shown to prevail, at which point no driver can reduce their travel costs by switching routes. System-optimal conditions can only exist through tolling and other similar controls on the roadway network. Flight operators are not always able to freely make routing decisions to minimize their travel costs, but desire to do so when possible and/or convenient.

The Collaborative Trajectory Options Program, or CTOP, is an ATFM initiative that is proposed to handle future projected demands. CTOP is similar to currently existing programs in that it aims to safely and efficiently meter aircraft flow through capacity constrained airspace regions. However, CTOP differs from current programs (such as the Airspace Flow Program, or AFP) in that it allocates alternate routes in addition to departure times. Most importantly, in exchange for this higher level of control, CTOP considers flight operators' preferences regarding the available resources in its resource allocation scheme. In this paper, we use a simple, classic traffic assignment principle to illustrate that if flight operators are asked to provide preference inputs that are used directly in a system-optimal resource allocation, a truly system-optimal solution may not be achievable; in fact, the actual solution may be closer to a user equilibrium solution. If flight operators are able to influence their resource allocations through their preference inputs, they will submit inputs that reflect their desire to minimize their expected travel costs, much like drivers on a roadway network make route choices to minimize their individual expected travel costs. This finding has two major implications. Firstly, in ATFM programs such as CTOP where user preferences are considered, a system-optimal solution may be an unrealistic goal that cannot be attained. Secondly, resource allocation schemes that aim to provide greater equity by sacrificing efficiency may not be as inefficient as they first appear, when compared against allocation schemes that aim to provide system-optimality.

In this paper, "operator" or "user" refers to ANSP customers such as commercial airlines and general aviation aircraft. "Traffic manager" refers to the agent responsible for allocating resources at the ANSP.

II. BACKGROUND

There are several programs that are used to meter traffic flow into constrained en route areas, including the Airspace Flow Program (AFP). In the AFP, the constrained airspace

The authors would like to thanks the Federal Aviation Administration for supporting this work.

region and the flights filed into this region during the time of reduced capacity are first identified. The reduced capacity is then allocated by assigning each impacted flight a delayed departure time on the original filed route. A flight can either accept the assigned departure time, or reject it and reroute around the constrained airspace (subject to traffic managers' approval). Slots to fly through the constrained region are vacated as flights are canceled or routed out, and the schedule is compressed such that remaining flights are moved up into earlier slots as available. The Collaborative Trajectory Options Program (CTOP) is a proposed concept that builds on the AFP to better coordinate rerouting under the anticipation of growing demands; it is designed to offer flight operators resources that consist of reroute options combined with delayed departure slots, while allowing operators to communicate their preferences regarding the offered resources. There is no program currently in place that applies reroutes and ground delays simultaneously – existing reroute programs cannot assign delays, and vice versa [1]. Optimization models that consider both rerouting and delay (on the ground and en route) decisions for constrained resource allocation have, however, been studied in the literature. One of the most well-known ATFM models [2] considers ground holding, air holding, and rerouting decisions in a static deterministic setting. It also requires the input of flight-specific air and ground-hold cost ratios. Hoffman, Burke, Lewis, Futer & Ball [3] develop an algorithm that allows for simultaneous rationing of ground and en route resources, as an alternative to using GDPs to handle en route constraints. Jakobovits, Krozel, & Penny [4] formulated an algorithm to schedule, reroute and airhold flights flying into and around constrained airspace. Mukherjee & Hansen [5] consider a variant of the single airport ground hold problem that considers reroutes for terminal airspace using a dynamic stochastic approach. The objective of many ATFM models is to minimize the system-wide cost of delay, i.e. maximize efficiency; however, providing equity between flights and/or flight operators is another important objective [6] [7].

For traffic managers to make resource assignment decisions that are of good value to flight operators, they should consider flight operators' resource preferences. Existing resource allocation programs such as GDPs and AFPs benefit from Collaborative Decision Making (CDM) [8], a joint government and industry initiative that aims to improve air traffic management by encouraging the exchange of up-to-date information between traffic managers and flight operators. Prior to the use of CDM for GDPs, airlines had little incentive to inform the FAA of their schedule updates, and in fact could inadvertently be penalized with higher delays by doing so. As a result, operators were reluctant to provide information updates and desirable slots would often go unused [9]. CDM provides operational incentives for airlines to keep their schedule changes up-to-date in the system. In addition, the system is transparent insofar as all flight operators can see what others are doing, thereby promoting a culture of accountability and self-regulation. The above features of CDM have greatly improved the efficiency of ATFM programs in which CDM is applied. However, operators' preferences are not explicitly communicated through CDM. ATFM concepts in which airlines do explicitly provide preference information to the FAA's resource allocation process have been studied [10] [11] but have yet to be implemented. CTOP is one of these concepts.

The literature on traffic assignment is extensive and well established. A user equilibrium results when drivers have the freedom to make their own personal travel choices to minimize their travel costs. Under congested conditions, a unique user equilibrium condition exists when drivers are homogeneous and (stochastic) errors in their travel time perceptions are accounted for. It has also been shown that for certain cases, the stochastic user equilibrium solution exists and is unique for heterogeneous driver classes [12] [13]. Konishi [14] extends Daganzo's work [12] to heterogeneous drivers with different utility functions. Konishi shows that an equilibrium solution is unique for a general class of utility functions on a simple network. There are many issues that arise in surface traffic assignment problems due to the nature of driver behavior, traffic controls and physical infrastructure characteristics on road networks. Many of these do not apply in ATFM due to fundamental differences in "driver" behavior, the physical organization of airspace, weather and operational conditions that shape aircraft flight patterns, and traffic management activities. However, from a traffic assignment perspective, the most significant difference is that aircraft cannot be in the airspace without permission from air traffic managers, and are always under their control. As a result, traditional ATFM models have focused on system-optimal traffic assignment solutions as mentioned above.

Resource allocation in general transportation networks has also been studied from a game-theoretic and/or market-based perspective. The one-player case corresponds to a classic system-optimal solution, while a many-player case yields the user equilibrium solution. Haurie & Marcotte [15] formulate a non-cooperative game where players, defined by their origin-destination pairs, must send flows along a congested network to serve demand at their destination node. The cost of sending flow along a given link is a function of the flow on that link (congestion effect). They show that the Nash-Cournot equilibrium corresponds to the user equilibrium. Wie [16] studies a dynamic extension of [15] where each player must make decisions (to minimize their cost) about sending a fixed volume of traffic from a single origin to a single destination over a network of routes. Players make simultaneous decisions over time, which is modeled using differential game theory, and Wie establishes a dynamic game theoretic interpretation of the user equilibrium condition. The author extends his work to account for two types of players – a user equilibrium player and a Cournot-Nash player [17]. The latter behaves to establish a system-optimal cost outcome for its set of network flow requirements.

This section has very briefly touched on the extensive body of research on ATFM models and transportation network analysis. This paper focuses on how the behaviors of flight operators in collaborative, user input-driven en route resource allocation programs like the CTOP might affect the results of these programs. We gain some insight into how system-optimal program designs may not align the objectives of the ANSP and flight operators, and with constrained user input based resource allocation, these program designs may result in more inefficient outcomes than aimed for or anticipated. This paper addresses the topic by using classic traffic assignment concepts not generally applied to ATFM problems.

III. Model Framework

A. Flight Cost Model

We use a simple modeling framework to illustrate our flight cost model. Say two fixes in en route airspace are connected by a nominal route, designated as such because it is the lowest cost path between the two points. Flights enter the nominal route at entry fix "A" and leave at exit fix "B". Under good conditions, all aircraft that are scheduled to use the nominal route can do so at their scheduled time without experiencing delay, meaning that the nominal route has sufficient capacity to serve the pre-constraint scheduled flight demand $D_0(t)$, in units of flights per hour. Suppose that a constraint develops along the nominal route, reducing its capacity such that the flight demand for this route cannot be accommodated without some queuing delay. The N flights originally scheduled to use this route (at a demand rate D_0) during the constrained period must either be rescheduled or re-routed to observe the reduced capacity. Flights are either given delayed departure times on the nominal route, or rerouted to one of $R-1$ alternate routes and possibly assigned a delayed departure time on that route. Each alternate route r is characterized by its capacity and travel time. The nominal route is assumed to have the lowest travel time, and therefore the lowest cost of travel. We assume that fixes A and B are not bottlenecks, and for the purpose of this analysis they can be thought of as the flights' origin and destination. Flight trajectories upstream of Fix A and downstream of Fix B are not considered in this analysis.

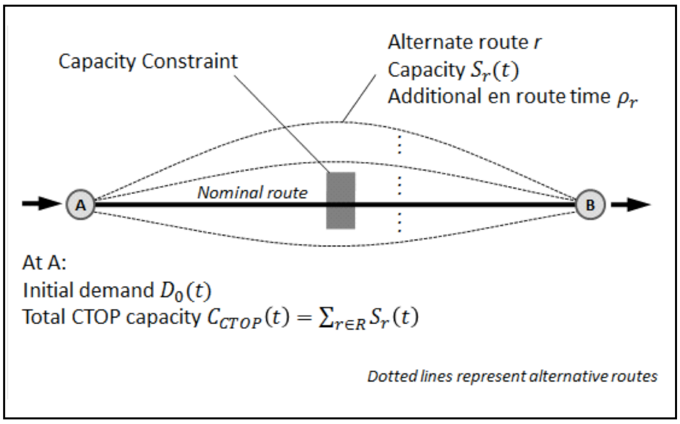

Figure 1. Model airspace geometry and select parameters

This research focuses on evaluating the added costs associated with greater en route time and ground delay due to the en route constraint. The flight cost function, $c_{n,j}$, represents the added cost of flight n taking departure slot j belonging to route r, due to constrained operating conditions. Over all the available routes $r = 1,2,\dots,R$, there are a total of J departure slots, where $r(j)$ indicates the route that slot j belongs to. $c_{n,j}$ is a function of the additional travel time of route r compared to the nominal route, time spent waiting on the ground for their assigned slot j on route r, and other factors. We assume it is the sum of the above components, and quantified in units of ground delay minutes.

$$c_{n,j} = \alpha_n \rho_{r(j)} + d_j - g_{0,n} + \varepsilon_{n,r(j)} \tag{1}$$

Where:

- $c_{n,j}$ represents the added cost of flight n taking departure slot j (which belongs to route r);
- α_n is the ratio of flight n's unit airborne time and ground delay costs;
- $\rho_{r(j)}$ is the additional en route time on route r compared to the nominal route $(\rho_{r(j)} \geq 0)$;
- d_j is the departure time on slot j at fix A;
- $g_{0,n}$ is flight n's original pre-CTOP scheduled departure time at fix A, and
- $\varepsilon_{n,r(j)}$ is an independently and identically distributed random term, assumed to be normally distributed such that $\varepsilon_{n,r(j)} \sim N(0, \sigma^2)$.

The random term $\varepsilon_{n,r(j)}$ is meant to capture flight n's route (r) and situation-specific cost factors that are not accounted for in the deterministic part of the model. It therefore represents the difference between the total cost $c_{n,j}$ and the deterministic cost. $\varepsilon_{n,r(j)}$ may be positive or negative; in the latter case it represents unknown cost-mitigating factors. $c_{n,j}$ accounts for direct costs including additional fuel, crew time, and equipment maintenance, and indirect costs such as passenger satisfaction, gate time, flight coordination, costs related to other airline internal business objectives, and others. We assume that there is no air holding, such that all anticipated delay is taken on the ground.

The unit cost of airborne delay exceeds that of ground delay such that $\alpha_n \geq 1$. If $\alpha_n = 2$, every one minute flight n spends in the air is equivalent in cost to n spending two minutes on the ground. $\rho_{r(j)}$ is non-negative, assuming the nominal route has the shortest flying time. Ground delay, or $(d_j - g_{0,n})$, is non-negative as well because aircraft cannot depart before their original scheduled time.

As noted above, the random term in (1) represents the part of airlines' route-specific flight costs about which traffic managers have little to no information. The specification of the random term, and its role in the allocation process, are key determinants of the performance of each allocation scheme. As elaborated further below, in some schemes the preference inputs provided to the traffic manager includes the information contained in the random term, while in others it does not. Clearly this will determine whether the objective function used by the traffic manager fully reflects flight operator costs, or does so only partially. The assumption of iid normality for the random term is made primarily for modeling convenience. Although the (deterministic and stochastic) costs of flights operated by the same airline may be correlated, we assume here that intra-airline flight differences are so pronounced that this correlation can be ignored.

B. System-Optimal Resource Allocation

The resource allocation scheme presented here incorporates flight operators' preference information in system-optimal allocations of reroutes and/or delayed departure times. It gives the operators flexibility in expressing their flights' route cost/preference information to traffic managers. The allocation cost calculation is best shown graphically as done in Hoffman, Lewis, & Jakobovits [11]. An illustration is shown in Figure 2. Suppose a flight n has three route options ($R = 3$), and the operator of flight n submits their inputs about the available options to traffic managers. The information contained in these inputs may differ from one resource allocation scheme to another. The inputs are used to construct $\Delta_{n,r}$, which is the cost of flight n traveling on route r before the ground delay cost is added, measured in units of ground delay minutes. Thus, if $\Delta_{n,1} = \Delta_{n,2} + k$, the operator of flight n would be indifferent to having n assigned to route 1 with no ground delay or route 2 with a ground delay of k minutes. $\Delta_{n,r}$ values contain all the operators' flight cost information available to the traffic managers; the traffic managers use this information to assign constrained resources to flights through the adopted allocation mechanism. The $\Delta_{n,r}$ values ensure that with any route and ground delay slot assigned to a flight, traffic managers have received some indication regarding the relative value of that route/slot combination to the flight operator. The total cost (as perceived by traffic managers) for flight n on route r is a function of the departure slot, and resulting ground delay, assigned to n on r. Suppose that, based on flight availability, the ground delay flight n must take is $\left(d_{j(1)} - g_{0,n}\right)$ on route 1, $\left(d_{j(2)} - g_{0,n}\right)$ on route 2, and $\left(d_{j(3)} - g_{0,n}\right)$ on route 3. The traffic managers would then determine that it would cost $c_{n,j(1)}$, $c_{n,j(2)}$, and $c_{n,j(3)}$ for flight n to take each of these routes. The resource that flight n ultimately receives will depend on the mechanism used to assign resources to flights participating in the CTOP.

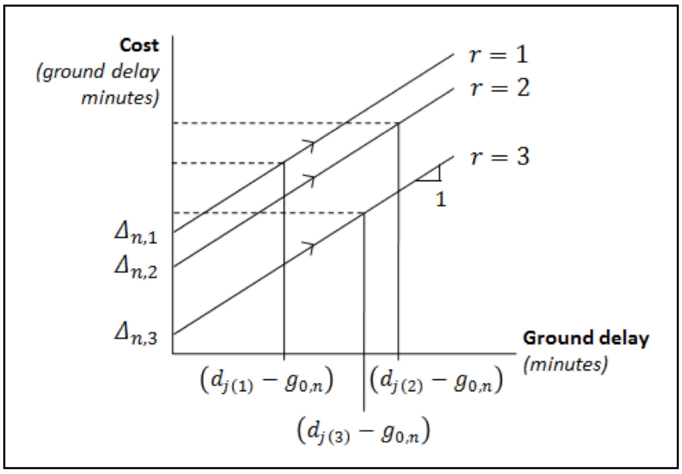

Figure 2. Illustrating the use of inputs $\Delta_{n,r}$ to determine flight costs.

When the CTOP is announced, traffic managers provide all operators of impacted flights with information about the constrained airspace (start time, duration, location, etc.) and the reroute options available. Operators are then asked to submit the requested route preference inputs to traffic managers by some pre-specified deadline: the operator of flight n submits $\Delta_{n,r}$ for each route r available in the CTOP. Traffic managers allocate all resources simultaneously using the submitted information, with the objective of minimizing the total (known) operator cost of the program, without explicit consideration for equity between flights and/or operators. This allocation scheme does not have any mechanism to reward or penalize flight operators for submitting their inputs. Operators do not know which route and slot the traffic managers will assign their flight(s); although operators do know what routes are available, they have no information about the ground delay that will be assigned to their flight on a given route. We assume that each operator would calculate the additional cost of a flight reassignment option assuming the flight cost model of (1). In one version of the system-optimal models, the Full Information, System-Optimal (FISO), flight operators submit complete information about their flights; they submit $\Delta_{n,r} \forall r$ which consist of the following parts of the flight cost model:

$$\Delta_{n,r} = \alpha_n \cdot \rho_r + \varepsilon_{n,r}, \varepsilon_{n,r} \sim N(0, \sigma^2) \tag{2}$$

Based on the illustration of Figure 2, flight n could be assigned any one of the three routes. Traffic managers will identify the minimum cost assignment based on all $\Delta_{n,r}$ submitted by all the flights for all routes, plus the ground delay associated with departure time slots. Due to the fact that the flight operators' complete route preference information is available for decision making through the information they offer to traffic managers, the random term of the flight cost model (representing other proprietary airline route preferences not accounted for in the deterministic part of the cost function) is available to the resource allocation process, and is therefore included within the objective function that traffic managers optimize. FISO is highly idealized, in that flight operators may not necessarily know or be able to provide this highly detailed and specific information in a convenient or timely manner, particularly in the absence of incentives (resource or equity guarantees). However unrealistic it is, in principal the FISO model yields the most efficient system performance that can be achieved from any CTOP allocation scheme. FISO can be solved numerically as an assignment problem.

In another version of the system-optimal models, or the Parametric System-Optimal (PASO) scheme, we envision that flight operators provide flight cost parameter inputs to a centrally-managed FAA database. Operators would be encouraged to update their parameters as necessary, and at any time. However, at the time a CTOP is announced, typically several hours prior to its actual start time [11], strategic resource allocation decisions must be made. The parameter values contained in the database at that time will be used to determine route and ground delay assignments for all impacted flights. If we assume that traffic managers have adopted the flight cost model of (1), the parameter requested of flight operators would be the air-to-ground cost ratio α. Therefore traffic managers do not receive complete information (as per (1)) about the operators' flight costs in PASO, but rather, $\Delta_{n,r} = \alpha_n \rho_r$. These inputs are used to perform a system-optimal resource allocation, albeit one based on incomplete information because the random term is not included in the input (recall that in FISO, the private route preference information represented by the random term is provided through $\Delta_{n,r}$).

PASO has two main features of interest. Firstly, the parametric input is very flexible in that it can be used to estimate the cost of any routing option. In FISO, $\Delta_{n,r}$ are submitted specifically for the available route alternatives in a particular CTOP, because they contain the random term and therefore the additional information it contains about route preferences. The advantage of PASO is that even if a flight operator does not have complete information about all the routes available in CTOP, traffic managers can still use the operator's parametric input to identify a desirable option that may not originally have been available or they might not have been aware was available. Secondly, in traffic management programs like the AFP and CTOP, decisions must be made quickly, and operators may not have the resources to provide highly detailed information about their flights (as represented by the random term) in a convenient or timely manner. By providing α values to the database, operators are ensured that the ANSP has at least some generic information – not necessarily particular to a specific capacity constraint situation – about their flights and cost structure.

If the random term in (1) and (2) has a small variance (i.e. σ^2 is small), the PASO resource allocation will be efficient, because the deterministic portion of the flight cost model is a good reflection of actual costs. If, however, the random term has a large variance, PASO resource allocations will be less efficient. We would like to ascertain how PASO performs as the variance of the random term – and hence the inability of the cost function to capture information about flight operators' route preferences – increases.

PASO is formulated as an assignment problem like FISO, but with the objective function consisting only of the deterministic part of the flight cost function. Both FISO and PASO, in addition to several other schemes that assign en route resources within the CTOP paradigm using different allocation rules, are further described and investigated in [18].

IV. Implications of Competitive Behavior in System-Optimal Resource Allocation Schemes

The commercial airline industry is a highly competitive environment, and flight operators are often competing for resources that are constrained due to weather and other operational limitations. Flight operators will do what they can to ensure they are treated equitably, and obtain the resources they require to best fulfill their business objectives. This has sometimes resulted in a lack of updated information provision to the ANSP, as well as gaming behaviors, which are both of concern with respect to the efficiency of collaborative ATFM programs [19]. As a result, we must understand how these types of behaviors might arise, and be encouraged, in particular designs of user input-driven resource allocation mechanisms where the objectives of the ANSP and the flight operators could be misaligned. We must also understand the resulting performance of each allocation mechanism with respect to efficiency and equity objectives. In this section we present an analysis of the potential implications of competitive behaviors within the system-optimal resource allocation schemes presented in Section 3.2.

A. *Original versus Revised Parameter Inputs*

In both system-optimal resource allocation schemes, traffic managers minimize the total flight costs that have been communicated to them by the flight operators. In a system-optimal assignment, the resources allocated to flights typically vary in cost, such that some flights end up with more desirable resources than others. As a result, over time and many repeated manifestations of similar constraint situations and resource allocation outcomes, rational flight operators are likely to minimize the expected cost of their allocations through their preference communications. We anticipate that this behavior will be exhibited by all (rational) flight operators, and thus lead to equilibrium conditions where no flight can expect to lower their expected assignment cost by changing their strategy.

According to the flight cost model and the resource allocation schemes presented previously, flight operators can exercise some control over their expected resource assignment costs through the information they provide to traffic managers. We want to know what flight operators might do when faced with this decision, and we start by exploiting the properties of a basic traffic assignment analysis [13]. We present a highly simplified setup where each flight in a set of N flights must be assigned to one of two available routes ($R = 2$). Each route has slots spaced at identical headways (g), and route 1's en route time is greater than that of route 2 (such that $\rho_1 > \rho_2$). All N flights have identical original scheduled departure times ($g_0 \approx 0$), air-to-ground cost ratios (α), and no additional unknown route preferences (such that $\varepsilon_{n,r} = 0 \; \forall n, r$). The results of applying the PASO and FISO schemes are therefore identical; the system-optimal allocation will result in an assignment of X_1^* flights to route 1, and $X_2^* = N - X_1^*$ flights to route 2. Equation (3) gives us the values of X_1^* and X_2^*:

$$X_1^* = \frac{\alpha\cdot(\rho_2-\rho_1)+gN}{2g}; \;\; X_2^* = N - X_1^* = \frac{\alpha\cdot(\rho_1-\rho_2)+gN}{2g} \qquad (3)$$

where α is the ratio of air-to-ground cost for all flights, and all other parameters are as described previously.

Figure 3 is a graphical representation of the N flights' expected route assignment costs. The lightweight dotted lines represent the average cost that flights would expect to incur by being assigned to a route, as a function of α and the total number of flights assigned to that route. Although in theory each flight is assigned to a distinct slot, after a flight is assigned to a given route, in this case slot assignments on that route are completely arbitrary. As a result, a flight can only know the expected cost of being assigned to a certain route as a function of the total flights assigned to that route and their original α value. Similarly, the light solid lines represent each route's expected marginal cost curves with α. The system-optimal assignment with α (3) is found from the point where the expected marginal cost curves of routes 1 and 2 are identical, represented by point A in Figure 3. Since $\rho_1 > \rho_2$, it follows that $X_1^* < X_2^*$.

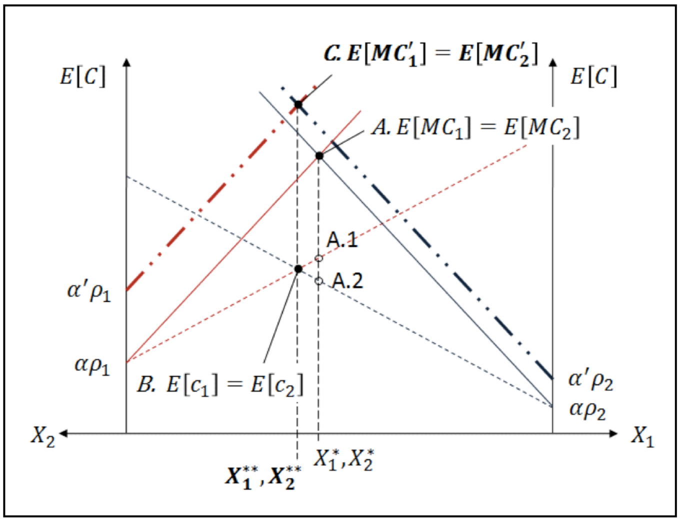

Figure 3. Original versus revised system-optimal solutions

Also, if one draws a vertical line at point A, one can observe that the expected cost of being on route 1 (point A.1 in the figure) is higher than the expected cost of being on route 2 at the system-optimal assignment (A.2). As a result, flight operators will want to maximize the probability of being assigned to route 2 instead of route 1 using whatever strategic handle is available to them, which in this case are the air-to-ground cost ratios they submit to traffic managers. If a flight operator should submit an air-to-ground cost ratio value reflecting the original, non-competitive situation (α), the probability of their flight being assigned to the lower cost route (route 2) is X_2^*/N. Instead, imagine they submit a revised cost ratio $\alpha' > \alpha$, because it is inherently more beneficial for them to be on the lower cost route (route 2) rather than route 1. By submitting α', their flight will be assigned to route 2 with probability 1 if all others were to submit as if they were not in this competitive situation. However, it is unlikely that a submission scenario like this would occur: all flight operators are aware of the capacity constraint, and all would benefit from having their flight assigned to the second route. The question is, is there a unique air-to-ground cost ratio that is eventually submitted by flight operators in order to maximize their probability of assignment to route 2, and what is its value at equilibrium?

Given that all flights are identical and their operators aim to minimize their expected costs, we imagine that this goal is reflected in their submitted cost ratio values, which in turn will push the assignment towards a user equilibrium (UE) based on their original, non-competitive cost ratios α. At this UE assignment with α, the original expected cost of a flight being assigned to either route are equal, or $E[c_1] = E[c_2]$. At this assignment, flights have no incentive to change their inputs as they cannot lower their original expected assignment costs by doing so. The expected cost at the user equilibrium is represented by the following expression:

$$E[c_1] = E[c_2] \Rightarrow \alpha\rho_1 + 0.5gX_1^{**} = \alpha\rho_2 + 0.5gX_2^{**} \qquad (4)$$

X_1^{**} and X_2^{**} are the user equilibrium flight assignments to routes 1 and 2, respectively, with α (Point B in Figure 3). We can imagine that over many repeated occurrences of this CTOP, flights will submit α' values that yield a system-optimal assignment with X_1^{**} and X_2^{**} – the assignment at the non-competitive (with α) user equilibrium of (4). This system-optimal assignment is found when:

$$E[MC_1'] = E[MC_2'] \Rightarrow \alpha'\rho_1 + gX_1^{**} = \alpha'\rho_2 + gX_2^{**} \qquad (5)$$

The value of α' can be determined by drawing a vertical line through point B, and then finding the point along this vertical line where the marginal cost curves of (5) intersect. It is identified as point C in Figure 3. We can now find an expression for α', the revised (due to competition) air-to-ground cost ratio submitted by flights to obtain the original user equilibrium solution, by solving (4) and (5):

$$\alpha' = 2\alpha \qquad (6)$$

The equilibrium (where all flights have identical expected assignment costs) is reached when all flights submit α' that is twice that of their original non-competitive cost ratio α. Given the assumption that all flight operators are rational – they make choices to minimize their expected costs – we see that under a system-optimal allocation scheme, flights will need to adjust their inputs to reflect the competitive nature of the allocation in their flight cost structures. The revised inputs submitted to traffic managers are twice the original values in cases with two and three routes. The calculation is contained in the Appendix.

In the classic traffic assignment literature, travel time perception errors are accounted for in a stochastic formulation of the above. Each route (or link) has an actual cost of travel, but drivers may perceive this actual cost erroneously. Their choices may be affected by this error, which is represented by a random term ε. Therefore, at the SUE, no driver can reduce his perceived travel time by changing routes. All route choice decisions are made based on this "erroneous" perception of link cost. In our flight cost model, the random term does not represent perception error but rather, the airlines' true preferences for routes. As a result, at the user equilibrium, no flight operator can reduce his true travel cost by changing routes. If we assume that the random term is Gumbel distributed with scale parameter ν, the expression for α' – the revised competitive air-to-ground cost ratio submitted by flights to obtain a non-competitive user equilibrium – becomes:

$$exp\left(\frac{1}{2\nu}(2\alpha^T - \alpha^L)(\rho_1 - \rho_2)\right) = \frac{Ng_1 - \alpha^L(\rho_2 - \rho_1)}{Ng_2 + \alpha^L(\rho_2 - \rho_1)} \qquad (7)$$

When $\alpha = 2$, preliminary numerical investigations show that as $\nu \to 0, \alpha' \to 4$. As ν grows larger, α' is asymptotic to a value that can be smaller or larger than $\alpha'_{\nu\to 0}$. In fact, in some cases, as ν grows large, $\alpha' < \alpha$. The behavior of α' is to be further investigated in future work. A closed form solution like (7) cannot be obtained for a normally distributed random term; in the following section, to find the total cost solutions for both FISO under a non-competitive user equilibrium, we use the method of successive averages.

This simple analysis suggests that traffic managers may not be able to sustain a system-optimal allocation when aiming for one based on inputs that reflect a non-competitive situation; rather, flight operators will create a user equilibrium by virtue of the fact that they are trying to operate their flights in a competitive situation. A similar result may be true in the case of heterogeneous flights. It has been shown in past research that unique user equilibriums do exist under certain conditions for heterogeneous commuters in both network and single bottleneck models [20] [12] [14]. However, there may be instabilities due to other gaming behaviors and flight characteristics, and further analysis must be pursued in future research.

B. Comparison to other Resource Allocation Algorithms

In [18], numerical examples are used to illustrate and compare the total cost results of several different possible resource allocation schemes within a highly idealized CTOP framework, where preference inputs are incorporated into the resource allocation decision process. In particular, the results of the system-optimal schemes are compared against the First-Submitted, First-Assigned (FSFA) algorithm. In FSFA, we assume that a flight is assigned the best resource available at the time of their input submission, which would consist of (2) like FISO. Therefore, we conjecture that flight operators have no incentive to submit inputs that are implicitly competitive in value, as the competition for resources is embodied in the submission time rather than the submission itself. We also assume that this submission order is random and independent of other flight and flight operator characteristics (this assumption was further explored in [18] under a game theoretic setup). The total cost results of all allocation schemes are represented as a ratio of the FISO results; FISO always yields the most efficient solution in any situation and therefore can be used as a benchmark. However, given that FISO may not be achievable and rather, the user equilibrium equivalent is more likely, we investigate how it compares against FSFA.

We consider an application of the FISO scheme that results in a non-competitive user equilibrium and corresponding revised competitive system-optimum, in a highly simplified illustrative example. Suppose 50 flights are to be reassigned routes and departure times due to an en route constraint. The nominal route remains open, but at a greatly reduced capacity. In this illustration, the CTOP will have a total of four route options, one of which includes the nominal route under a decreased capacity (route 4). For routes $r = 1, \dots, 4$, the constant departure headways (g, minutes) at Fix A and additional en route time (ρ, in minutes) are $(g, \rho)_r = \{(3,30)_1, (6,20)_2, (5,15)_3, (8,0)_4\}$.

Now we introduce the flight demand characteristics. We continue with the assumption that all flights are identical in their air-to-ground cost ratios, such that $\alpha_n = \alpha \; \forall n$. We use a value of 2 for this ratio, as it has been cited in the literature that one unit of en route delay is equal in cost to roughly two units of ground delay [5]. We retain the assumption that the $\Delta_{n,r}$ submission order in FSFA is random and independent.

Figure 4 presents the results. The x-axis represents increasing values of the standard deviation of the random term in the flight cost model. Specifically, each point on the x-axis represents the value of σ as a proportion of the average flight cost in FISO under perfect information conditions, $\hat{c}_{FISO}$. For instance, the point "0.10" means that σ is 10% of $\hat{c}_{FISO}$. The y-axis represents the total cost result of each scheme's results as a ratio of the full information system-optimal (FISO) total cost, or $y = C_s / C_{FISO} = C'_s$; $s \in \{\text{FSFA}, \text{FISO}_{\text{SUE}}\}$. We compare the performances of FSFA and FISO_{SUE} against the non-competitive system-optimal FISO solution (C_{FISO}) because it represents the most efficient solution possible in any situation. Each point represents 500 simulation runs of FISO and FSFA and 30 of the FISO_{SUE} results (which were obtained using the method of successive averages).

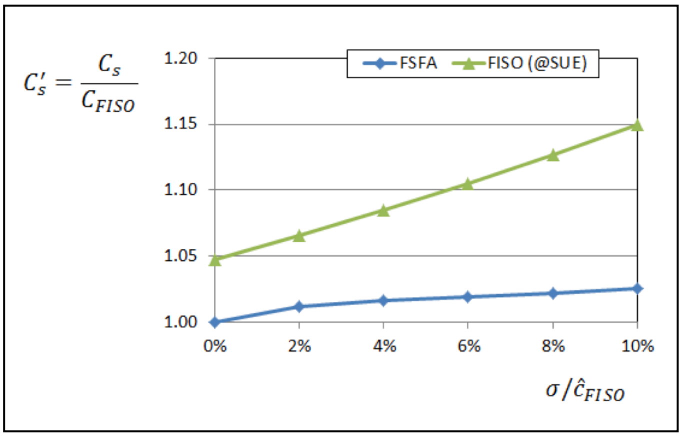

Figure 4. Example total cost results

As all flights are homogeneous in their deterministic costs (i.e. $\alpha_n = \alpha \; \forall n$), when $\sigma = 0$, FSFA=FISO. The main feature of Figure 4 is the fact that the FISO_{SUE} solutions are inferior to FSFA for all values of σ shown. If application of FISO in fact yields a non-competitive user equilibrium while appearing to be system-optimum with the submitted values, FISO is always more inefficient than FSFA, in the example of Figure 4.

V. CONCLUDING REMARKS

The purpose of this paper is to gain some insight into how system-optimal goals in collaborative, user input-driven en route resource allocation programs like the CTOP may not align the objectives of the ANSP and flight operators, and therefore may result in more inefficient outcomes than aimed for or anticipated. We have used simple, classic traffic assignment principles to illustrate that if flight operators are able to provide preference inputs that are used directly in a system-optimal resource allocation, they will communicate their desire to minimize their own travel costs much like drivers on a roadway. Analogously, a system-optimal solution may not be achievable and in fact the solution may be closer to a user equilibrium. We show that in a simple two-route case with homogeneous flight costs, a user equilibrium is achieved when operators submit revised flight cost parameters (air-to-ground ratios) that are twice the value of their cost parameters in the originally non-competitive situation. This finding has two major implications. Firstly, in ATFM programs such as CTOP where user preferences are considered, system-optimal solutions may be an unrealistic goal that cannot be attained due to competition for constrained resources. Secondly, resource allocation schemes like First-Submitted, First-Assigned

(FSFA) that aim to provide greater equity by sacrificing efficiency may not be as inefficient as they first appear, when compared against schemes that aim to provide system-optimality. The results of this paper are a step towards better understanding the implications of information exchange between flight operators and an ANSP. There are many ways to extend our analyses of operator/ANSP interactions in user input-based system-optimal resource allocation schemes. It would be beneficial to consider heterogeneity in the parameters of the flight cost model (α, ε_r) in the competitive analysis of the FISO and PASO inputs schemes; the results of [12] and [14] would help in doing so. Also we can assess other allocation schemes [18] against a user equilibrium solution to the FISO allocation scheme.

REFERENCES

[1] FAA (2010). CTOP Industry Day. Retrieved April 2011 from Federal Aviation Administration: *http://cdm.fly.faa.gov/ad/R7docs/CTOP%20Industry%20Day%20presentations%2010132010.pdf*

[2] Bertsimas, D., & Stock Patterson, S. (1998). The Air Traffic Flow Management Problem with Enroute Capacities. Operations Research , 46 (3), 406-422.

[3] Hoffman, R., Burke, J., Lewis, T., Futer, A., & Ball, M. (2005). Resource Allocation Principles for Airspace Flow Control. San Francisco: AIAA Guidance, Navigation, and Control Conference and Exhibit.

[4] Jakobovits, R., Krozel, J., & Penny, S. (2005). Ground Delay Programs to Address Weather within En Route Flow Constrained Areas. San Francisco, CA: AIAA Guidance, Navigation and Control Conference and Exhibit.

[5] Mukherjee, A., & Hansen, M. (2009). A dynamic rerouting model for air traffic management. Transportation Research Part B: Methodological , 43, 159-171.

[6] Vossen, T., & Ball, M. (2005). Optimization and Mediated Bartering Models for Ground Delay Program. Naval Research Logistics , 53 (1), 75-90.

[7] Pourtaklo, N. V., & Ball, M. (2009). Equitable Allocation of Enroute Airspace Resources. Eighth USA/Europe Air Traffic Management Research and Development Seminar (ATM2009). Napa.

[8] Ball, M., Futer, A., Hoffman, R., & Sherry, J. (2002). Rationing Schemes for En-route Air Traffic Management (CDM Memorandum).

[9] Vossen, T., Ball, M., Hoffman, R., & Wambsganss, M. (2003). A General Approach to Equity in Traffic Flow Management and its Application to Mitigating Exemption Bias in Ground Delay Programs. 5th EUROCONTROL/FAA ATM R&D Seminar. Budapest, Hungary.

[10] Goodhart, J. (2000). Increasing Airline Operational Control in a Constrained Air Traffic System. University of California, Berkeley: PhD Dissertation.

[11] Hoffman, R., Lewis, T., & Jakobovits, R. (2004). Flow Constrained Area Rerouting Decision Support Tool, Phase I SBIR: Final Report. Metron Aviation.

[12] Daganzo, C. F. (1983). Stochastic Network Equilibrium with Multiple Vehicle Types and Asymmetric, Indefinite Link Cost Jacobians. Transportation Science , 17 (3), 282-300.

[13] Sheffi, Y. (1985). Urban Transportation Networks: Equilibrium Analysis with Mathematical Programming Methods. Prentice-Hall.

[14] Konishi, H. (2004). Uniqueness of User Equilibrium in Transportation Networks with Heterogeneous Commuters. Transportation Science , 38 (3), 315–330.

[15] Haurie, A., & Marcotte, P. (1985). On the Relationship between Nash-Cournot and Wardrop Equilibria. Networks , 15 (3), 295-308.

[16] Wie, B.-W. (1993). A Differential Game Model of Nash Equilibrium on a Congested Traffic Network. Networks , 23 (6), 557-565.

[17] Wie, B.-W. (1995). A differential game approach to the dynamic mixed behavior traffic network equilibrium problem. European Journal of Operational Research , 83, 117-136.

[18] Kim, A. (2011). Collaborative Resource Allocation Strategies for Air Traffic Flow Management. University of California, Berkeley: PhD Dissertation.

[19] Vossen, T., & Ball, M. (2006). Optimization and Mediated Bartering Models for Ground Delay Program. Naval Research Logistics, 53 (1), 75-90.

[20] Newell, G. (1987). The morning commute for non-identical travelers. Transportation Science, 21 (2), 74–88.

APPENDIX

The total cost of a system-optimal allocation in a 3-route case with competitive air-to-ground cost ratio α' is:

$$C = \alpha'(X_1\rho_1 + X_2\rho_2 + X_3\rho_3) + 0.5g(X_1^2 + X_2^2 + X_3^2)$$

where $X_3 = N - X_1 - X_2$.

$$\frac{\partial C}{\partial X_1} = \alpha' \cdot (\rho_1 - \rho_3) + 2gX_1 + gX_2 - gN = 0$$

$$\frac{\partial C}{\partial X_2} = \alpha' \cdot (\rho_2 - \rho_3) + gX_1 + 2gX_2 - gN = 0$$

We solve the above to find:

$$X_1 = \frac{N}{3} + \frac{\alpha'}{3g} \cdot (\rho_3 + \rho_2 - 2\rho_1)$$

$$X_2 = \frac{N}{3} + \frac{\alpha'}{3g} \cdot (\rho_3 - 2\rho_2 + \rho_1)$$

$$X_3 = \frac{N}{3} + \frac{\alpha'}{3g} \cdot (-2\rho_3 + \rho_2 + \rho_1)$$

The expected cost of flying any route is equal at the non-competitive UE with α. The following must be satisfied:

$$E[c_1] = E[c_2] = E[c_3] \Longrightarrow \alpha\rho_1 + \frac{gX_1}{2} = \alpha\rho_2 + \frac{gX_2}{2} = \alpha\rho_3 + \frac{gX_3}{2}$$

And we find that at UE: $X_2 = \frac{2\alpha \cdot (\rho_1 - 2\rho_2 + \rho_3) + gN}{3g}$

We solve the two expressions for X_2 shown above and find that $\alpha' = 2\alpha$, which is identical to the two-route case. □

In a two-route case with slots spaced at g_1 and g_2 on routes 1 and 2 respectively:

$$X_1 = \frac{g_2 N + \alpha' \cdot (\rho_2 - \rho_1)}{g_1 + g_2}$$

Say $\varepsilon \sim \text{Gumbel}(\alpha, \nu), \nu > 0$; the probability of traffic managers assigning a flight to route 1 when they have incomplete information about flights (i.e. do not have ε) is

$$P_1 = \frac{X_1}{N} = \left(1 + exp\left(\frac{E[c_1] - E[c_2]}{\nu}\right)\right)^{-1}$$

It follows that

$$exp\left(\frac{E[c_1] - E[c_2]}{\nu}\right) = \frac{N}{X_1} = \frac{Ng_1 - \alpha' \cdot (\rho_2 - \rho_1)}{Ng_2 + \alpha' \cdot (\rho_2 - \rho_1)} \text{ and}$$

$$exp\left(\frac{E[c_1] - E[c_2]}{\nu}\right) = exp\left(\frac{1}{2\nu}(2\alpha - \alpha')(\rho_1 - \rho_2)\right)$$

Such that $exp\left(\frac{1}{2\nu}(2\alpha - \alpha')(\rho_1 - \rho_2)\right) = \frac{Ng_1 - \alpha' \cdot (\rho_2 - \rho_1)}{Ng_2 + \alpha' \cdot (\rho_2 - \rho_1)}$ □