

Privacy and Stability in Airport Ground Delay Programs

Karthik Gopalakrishnan and Hamsa Balakrishnan

Abstract—Adverse weather can reduce airport capacity. When the number of arriving aircraft exceeds this reduced capacity, flights can get delayed. A Ground Delay Program (GDP) is a strategy by which aircraft landing slots can be redistributed so that the flights are delayed on the ground itself and not in the air. This increases safety, reduces the fuel consumption and hence the operating costs for airlines. In this paper, we present six algorithms that perform this slot reassignment. These algorithms differ in the extent to which slots can get re-arranged in real-time and the amount of information that the airlines must reveal to the central planner during the implementation. These two features of the algorithm are called as stability and privacy respectively. The efficiency of these six algorithms, measured in terms of the expected delay cost for flights, is compared using operational data for La Guardia Airport in New York. A two-step Receding Horizon Static (2-step RHS) model is shown to be a good compromise based on present expectations of privacy and stability in the system.

I. INTRODUCTION

Airports are capacitated resources that are constrained by the rate at which aircraft can take off or land [1], [2]. The capacity of an airport can vary significantly depending on factors such as wind, runway configuration, and the visibility [3]. Adverse weather can decrease the rate at which aircraft can land (also called as the airport arrival capacity). If such capacity decrease is not accounted for, aircraft will continue to arrive at the vicinity of the affected airport, but will be placed in an airborne queue until they are able to land. The airborne holding increases the operating costs for airlines (due to fuel burn), increase Air Traffic Controller (ATC) workload and decreases safety. If a decrease in airport capacity is predicted in advance, the demand-capacity imbalance can be better managed by initiating a Ground Delay Program (GDP). This traffic management strategy is based on the premise that it is safer and less expensive to delay an aircraft prior to departure, on the ground, than in an airborne queue near the destination airport. Once a GDP is initiated, flights that are departing to the impacted airport get revised departure times, ensuring that they will be able to land without significant airborne delays [4], [5], [6], [7].

Optimally assigning revised departure times has two main practical challenges. Firstly, since the airport capacity forecast is probabilistic (which may be represented using scenario trees), the optimal reassignment is a stochastic optimization problem. The solution can therefore be highly

dynamic, meaning that the optimal departure times for flights may change as we get more certain information about the capacity. Secondly, a GDP trades off airborne delay for ground delay. Hence, the optimization requires a knowledge of flight-specific costs for airborne and ground delay. For example, if the cost for an hour of ground delay for a flight is one-third that of an hour of airborne delay, that flight would prefer up to 3 hours of ground delay, to an hour of airborne delay. However, the competitive nature of the airline industry makes such flight-specific delay costs proprietary information, since they may be used to deconstruct scheduling strategies.

Privacy in the context of a GDP denotes the amount of information that needs to be shared by a participating airline. We refer to a GDP that requires the airlines to share less or no information about their delay costs as one which respects airline privacy concerns. If the airlines have to report their delay costs to the central planner, then the GDP does not permit airlines to maintain a high level of privacy. If a GDP does not change the flight schedules frequently, and sticks to a prescribed reassignment, then we refer to the GDP allocation as more stable. On the other hand, if the GDP schedule got updated every hour, it is less stable. In this paper, we present different approaches to implementing GDPs that provide varying degrees of flexibility in terms of airline privacy and stability. We show that a new approach, called the 2-step Receding Horizon Static formulation (2-step RHS) is a good compromise between the conflicting requirements of more stability, privacy and optimality. Further, we define two important metrics (*Price of Stability* and *Price of Privacy*) that quantify the loss in system efficiency due to higher privacy and stability requirements.

A. Current structure of a GDP

We first describe how GDPs are currently implemented (Fig. 1) in the US by the Federal Aviation Administration (FAA). When a reduced airport capacity is anticipated, the FAA initiates a GDP by issuing a start time, an end time, and a prescribed arrival rate for the duration. For example, the FAA may issue a GDP for Los Angeles airport between 3-6 pm, with arrival rates of 30, 30 and 40 landings/hour, for each hour for the 3-hour period. Once the GDP is initiated, the flights are reassigned slots so that they match the prescribed arrival rate. This is done on a first-scheduled first-served basis (also known as Ration by Schedule or RBS). All flights that are arriving during the GDP period are now assigned a revised departure time. The RBS allocations are not made considering the flight-specific delay costs of each aircraft since the airlines do not share this information.

This work was partially supported by NASA under cooperative agreement number NNX16AF61A, and by NSF under CPS:Frontiers:FORCES, grant number 1239054.

Karthik Gopalakrishnan and Hamsa Balakrishnan are with the Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA. {karthikg, hamsa}@mit.edu

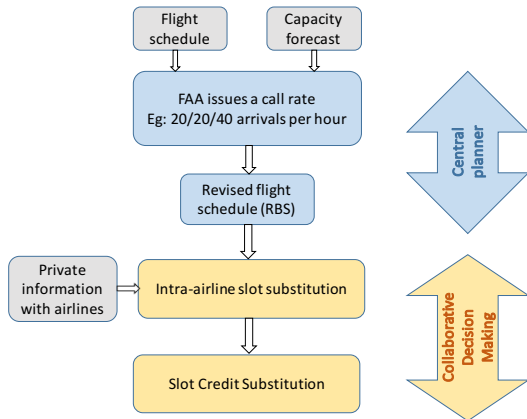


Fig. 1. Steps in the implementation of a GDP.

Once the RBS allocations are made, the next step, called Collaborative Decision Making (CDM) allows the airlines to use their private information to improve the efficiency of the allocation. The first component of CDM is intra-airline swaps. Delay-sensitive flights can be swapped within an airline and assigned to low delay slots, based on the private flight-specific costs. The second component of CDM is Slot Credit Substitution. If an airline cannot utilize an assigned slot or decides to cancel its flight, it can release the slot and request a new slot at any preferred time in the future.

This current GDP implementation has two important features:

- 1) Uncertainties in the capacity estimate (which may be due to weather) are abstracted out and a deterministic departure time is given to the aircraft. Consequently, there is no need for any real time communication during the GDP between the airlines and the FAA and the revised takeoff time does not change. The issue occurs when the capacity does not pan out as expected, and then the FAA will have to revise the GDP rates (or even cancel it).
- 2) The FAA does not have access to airline specific information, like the delay cost of each flight. The FAA introduced the second step of intra-airline substitution to incorporate more flexibility in decision making while respecting airlines privacy concerns.

B. Prior literature

Existing literature on GDPs focused on two aspects. One relates to obtaining an optimal arrival rate from scenario trees (which is the role of the central planner in Fig 1) and the other focuses on privacy and equitable allocation of slots.

The earliest models computed an optimal arrival rate for a given flight schedule based on a probabilistic capacity scenario tree [6]. The advantages of this approach was that no private cost information was required from airlines and once the allocation is made, it is followed irrespective of any future updates in the capacity forecast. A further refinement, where the departure slot is dynamically allocated based on

evolving capacity information is described in [7] and [4]. An important point about these models is that they assume no further intra-airline swaps or slot credit substitution procedures which are presently employed. Hence they are not reflective of the actual practical costs of implementing such an algorithm. Recently, formulations that obtain a landing slot allocation while accounting for a future CDM step have been developed [5].

Fairness and equity in resource allocation has been studied in the air traffic flow management setting [8], [9], [10]. In the context of GDPs, one could explicitly enforce equality among flights of different duration [11], reduce exemption bias [12] or incorporate a Ration By Distance formulation [13]. The mediated bartering model [14], [15] and Top Trading Cycles Algorithm [16] address the intra-airline slot exchange process. Cox and Kocenderfer [17], [18] proposed a CDM compliant Markov Decision Model. However the computational complexity of such an approach poses a challenge, even when simple cost functions are chosen.

There are two questions not addressed in previous literature. One is related to privacy- If airlines did not display such a strong sense of privacy and if they were willing to share information to the central planner (and not to other competing airlines), what is the increase in efficiency? The second is- How do algorithms that partially use updated capacity information perform? We address these two questions and put them into a formal perspective using the notions of stability and privacy.

C. Potential GDP architectures

A GDP architecture where airlines have privacy constraints involves two steps. The first step is where the central planner (the FAA in USA) decides on an optimal arrival rate based on a probabilistic capacity scenario tree. In the second step, the airlines use their private flight-specific delay costs to swap flights and run slot substitutions in case of cancellations. Thus, a GDP structure that incorporates privacy concerns is run in two stages and we refer to it as a *2-step algorithm*. If privacy is not a concern, the airlines would submit their flight specific cost to the central planner and the optimal reassignment can be obtained in a single step. A GDP architecture where privacy is not a concern is called a *1-step algorithm*.

With regard to various degrees of stability, we consider three models. The *static* model is the most stable, and there are no updates or changes to the revised landing slots once they are made. On the other end, we have a *dynamic* formulation, where flight revisions are made as and when new capacity information is available. The third model is a *Receding Horizon Static (RHS)* formulation. It involves running the static formulation repeatedly and is more stable than the dynamic model but less stable than the static formulation.

In this paper, we present six formulations (1 and 2-step variants of static, dynamic and RHS) and compare their efficiency. We will discuss the tradeoffs between efficiency, privacy and stability. Further, we will illustrate why a 2-step




Fig. 2. Scenario tree representation of uncertain airport capacity. Each branch(scenario) q has an associated probability π_q [19].

RHS formulation is the best compromise by simulating a GDP at New York City's LaGuardia airport.

II. GDP FORMULATIONS

In this section, we describe all the 1-step formulations (airlines do not have privacy concerns) followed by the 2-step formulations (airlines have privacy concerns).

In these formulations, the uncertain capacity is represented using a scenario tree. Fig. 2 shows a structured scenario tree. It is easy to see that any scenario tree that does not have more than two branches at any time instant can be represented using this structure. The numbers in each of the nodes is the airport arrival capacity for that time interval. T is the length of the GDP and Q as the set of all possible scenarios. In this paper, we use a scenario tree structure as in Fig. 2, and for this structure, $|Q| = T$. A scenario tree is the set $(M_{q,t}, \pi_q)$, where $M_{q,t}$ is the arrival capacity of the airport at the discrete time $t \in \{1, \dots, T\}$ in scenario $q \in Q$ and π_q is the probability of scenario q occurring. For example, in Fig. 2, in the third time interval ($t=3$) the capacity is 20 with probability 0.7 (if either S_3 , S_4 or S_5 occurs) and 40 with probability 0.3 (if either S_1 or S_2 occurs).

A. No privacy concerns

The central decision maker (FAA) has a list of flights f along with their schedule information (the arrival time arr_f , departure time dep_f , and flight duration dur_f), the scenario tree $(M_{q,t}, \pi_q)$ and the flight-specific delay cost. $C_{f,n}^g$ is the cost of ground holding flight f for n time steps. C^a is the cost of holding one aircraft in the air for 1 unit of time. For n time steps, the air hold cost of an aircraft is nC^a . The air-hold cost is assumed to be constant for all flights. When airlines release information about their flight specific cost, there is no need for the CDM step and the flight assignment is done through a single integrated optimization (a *1-step* algorithm).

Truthful reporting is not guaranteed in *1-step* algorithms. If an airline A1 was indeed more delay sensitive than airline A2, then flights of A1 will be delayed lesser than flights of A2. This would also create an incentive for airlines to lie about the true cost of their flight. We assume that when airlines do not have privacy concerns, they would report

costs truthfully. Developing algorithms that induce truthful revelation of preferences is an open problem.

Once a 1-step slot assignment is given to an airline, the airline would not obtain a lower cost even if it did any internal substitution (because if such a swap was possible, then the central allocation is not optimal). Also, if all airlines reported their true costs to the FAA, then the total delay costs will be lower than the 2-step process.

1) *1-step static*: X_t^f is a 0-1 variable that is 1 when flight f reaches the destination airport at time t . This is the time when it joins the landing airborne queue and the actual landing time may be at t or later depending on the state of the airborne queue. $W_{q,t}$ is the airborne queue at time t in scenario q . The optimal solution to a GDP of length T with Q scenarios is given by

$$\min \sum_f \sum_{t=arr_f}^{T+1} C_{f,t-arr_f}^g X_t^f + C^a \sum_q \pi_q \sum_{t=1}^T W_{q,t} \quad (1)$$

such that

$$\sum_{j=arr_f}^{T+1} X_j^f = 1, \quad \sum_{j=1}^{arr_f-1} X_j^f = 0 \quad \forall f \quad (2)$$

$$W_{q,t} \geq W_{q,t-1} - M_{q,t} + \sum_f X_t^f \quad \forall q, t \quad (3)$$

$$X_t^f \in \{0, 1\} \quad \forall f, t \quad (4)$$

$$W_{q,t} \geq 0 \quad \forall q, t \quad (5)$$

The objective function (1) minimizes the sum of the flight-specific ground holding and expected air holding costs. Constraint (2) ensures that all flights do not take off before their schedule and (3) describes the evolution of the airborne queue. The 1-step static formulation is ideal when airlines report true cost of delay and prefer to obtain a deterministic take-off time for their flights (i.e. a *stable* solution).

2) *1-step dynamic*: $X_{q,t}^f$ is a 0-1 decision variable that is 1 when flight f reaches the destination airport at time t in scenario q . $W_{q,t}$ is the same as Section II-A.1.

$$\min \sum_q \pi_q \left[\sum_f \sum_{t=arr_f}^{T+1} C_{f,t-arr_f}^g X_{q,t}^f + C^a \sum_{t=1}^T W_{q,t} \right] \quad (6)$$

such that

$$\sum_{j=arr_f}^{T+1} X_{q,j}^f = 1, \quad \sum_{j=1}^{arr_f-1} X_{q,j}^f = 0 \quad \forall f, q \quad (7)$$

$$W_{q,t} \geq W_{q,t-1} - M_{q,t} + \sum_f X_{q,t}^f \quad \forall q, t \quad (8)$$

$$X_{q_1,t}^f = X_{q_2,t}^f \quad \forall t \text{ and } q_1, q_2 \in G_{t-dur_f} \quad (9)$$

$$X_{q,t}^f \in \{0, 1\} \quad \forall f, q, t \quad (10)$$

$$W_{q,t} \geq 0 \quad \forall q, t \quad (11)$$

The dynamic solution allocates a take-off time to each aircraft in real time. Hence, the decision variable $X_{q,t}^f$ is a function of the scenario q . However, a decision at the take-off time t for any flight f can only incorporate all

the information available till t . The decision to ground hold or take-off can not depend on scenarios that cannot be distinguished at t . This is enforced by constraint (9) and G_t is the set of all scenarios that are indistinguishable at time t . For example, in Fig. 2, at $t = 3$, we cannot distinguish between S_3, S_4 and S_5 . Therefore, $G_3 = \{S_3, S_4, S_5\}$

3) *1-step RHS*: This architecture involves planning the GDP in two stages. The first stage is from time $t = 1$ to $t = \tilde{t}$ and the second stage is from $t = \tilde{t}$ to $t = T$. The time \tilde{t} , with $1 < \tilde{t} < T$, is called the update time. Two coupled static problems are solved corresponding to the two stages. The first stage corresponds to the first static solution and is applicable until time \tilde{t} . In the second stage, which runs from $t = \tilde{t}$ to T , a different static solution is implemented. At the update time \tilde{t} , because certain capacity scenarios are realized, either of the two cases occur- a scenario gets realized and can be uniquely identified (the resolved scenario), or the scenarios are still ambiguous (the unresolved scenarios). The sets Q_{res} represents scenarios that could possibly get resolved at the update time and Q_{unres} represents the scenarios that cannot be resolved at the update time. In Fig. 2, if the update time is $\tilde{t} = 3$, then the scenarios S_1 and S_2 are resolved and scenarios $S_3 - S_5$ are unresolved.

For update time $t = \tilde{t}$, $Q_{res} = \{1, \dots, \tilde{t} - 1\}$ is the set of resolved scenarios and $Q_{unres} = \{\tilde{t}, \dots, T\}$ is the set of unresolved scenarios when the scenario tree structure is as Fig. 2. $X_{q,t}^f$ and $W_{q,t}$ are defined previously in II-A.2. For a given update time \tilde{t} , the optimal slot allocation is

$$\min \sum_f \sum_q \pi_q \sum_{t=arr_f}^{T+1} C_{f,t-arr_f}^g X_{q,t}^f + C^a \sum_q \pi_q \sum_{t=1}^T W_{q,t} \quad (12)$$

such that

$$\sum_{j=arr_f}^{T+1} X_{q,j}^f = 1, \quad \sum_{j=1}^{arr_f-1} X_{q,j}^f = 0 \quad \forall f, q \in Q \quad (13)$$

$$W_{q,t} \geq W_{q,t-1} - M_{q,t} + \sum_f X_{q,t}^f \quad \forall t, q \in Q \quad (14)$$

$$X_{q,t}^f \in \{0, 1\} \quad \forall f, t, q \in Q \quad (15)$$

$$W_{q,t} \geq 0 \quad \forall t, q \in Q \quad (16)$$

$$X_{q_1,t}^f = X_{q_2,t}^f \quad \forall f, t < \tilde{t} + dur_f \text{ and } q_1, q_2 \in Q \quad (17)$$

$$X_{q_1,t}^f = X_{q_2,t}^f \quad \forall f, t \geq \tilde{t} + dur_f \text{ and } q_1, q_2 \in Q_{unres} \quad (18)$$

Constraint (17) ensures that stage 1 solution is independent of q (i.e it is a static solution) and Constraint (18) ensures that the stage 2 solution for unresolved cases is independent of Q_{unres} . Each of the resolved scenarios to have their own solution.

Once this algorithm is run, each flight gets a revised take-off time in stage 1 or stage 2. Stage 1 flights have a scenario independent takeoff time. Stage 2 flights will get their take-off times assigned to them at $t = \tilde{t}$, which depends on the scenario information at that time. It is important to note that there is no computation to be done in real time at \tilde{t} . The optimal update time \tilde{t} can be computed by enumerating the

total cost for all values of \tilde{t} . For more details on the RHS formulation, we refer the readers to [20].

B. Privacy concerns

Privacy concerns of airlines are addressed by having a 2-step GDP process. An initial allocation is done using an average ground and air delay cost for all flights. This allocation is communicated to the airlines, and they have the flexibility to make any internal slot substitutions using private delay cost information. An important feature of 2-step GDP architectures is that they are more fair than their corresponding 1-step solutions. This is because airlines are not distinguished based on their operating cost structures. An airline that leaves ample buffer between flights and is not very sensitive to additional delay is not penalized by getting more delay.

1) *2-step static*: Since flight specific cost information is not available to the central planner, the optimization described in Section II-A.1 is solved with $C_{f,n}^g = C_n^g$. For example, the cost of ground holding for each time period can be 1 unit and the cost for air holding can be 2.5 units. This gives an optimal solution X_t^f . The result, which includes the ground delay for each flight as well as its expected air hold time is communicated to the airlines.

Using flight-specific costs, each airline would like to swap slots internally such that delay-sensitive flights occupy low-delay slots and delay-insensitive flights occupy the high-delay slots. The airline thus solves a minimum-cost allocation problem for the CDM step.

Let $a \in \mathcal{A}$ denote the set of airlines and \mathcal{F}_a denote the set of aircraft operated by airline a . The initial flight corresponding to slot s_k is f_k . The cost of assigning any flight $f_i \in \mathcal{F}_a$ to slot s_k associated with flight f_k (slot s_k) is given by

$$C_{f_i, s_k} = C_{f_i, grd(s_k)}^g + C_{f_i, air(s_k)}^a \quad f_i, f_k \in \mathcal{F}_a \quad (19)$$

$grd(s_k)$ and is the ground delay assigned to slot s_k and $air(s_k)$ is the average air delay of the slot. $X_{i,k}$ is a binary variable that is 1 when flight f_i is swapped and allotted to slot s_k . The optimization problem finds the minimum cost, scenario-independent assignment of flights to slots.

$$\min \sum_{f \in \mathcal{F}_a} \sum_{k: f_k \in \mathcal{F}_a} X_{f,k} C_{f,s_k} \quad (20)$$

such that

$$\sum_{f \in \mathcal{F}_a} X_{f,k} = 1 \quad \forall k: f_k \in \mathcal{F}_a \quad (21)$$

$$\sum_{k: f_k \in \mathcal{F}_a} X_{f,k} = 1 \quad \forall f \in \mathcal{F}_a \quad (22)$$

$$X_{f,k} \leq feas(f, k) \quad \forall f \in \mathcal{F}_a, k: f_k \in \mathcal{F}_a \quad (23)$$

$$X_{f,k} \in \{0, 1\} \quad \forall f \in \mathcal{F}_a, k: f_k \in \mathcal{F}_a \quad (24)$$

$feas(f, k)$ is 1 if the flight f can be allocated to slot s_k . The condition for feasibility is that the flight cannot take off before its original published time, i.e. $feas(f, k) = 1$ if and only if $f, f_k \in \mathcal{F}_a$ and $dep(f) \geq dep(f_k) + grd(s_k)$.

2) *2-step dynamic*: The first step is a dynamic slot allocation similar to the one described in Section II-A.2 and an average ground hold cost of $C_{f,n}^g = C_n^g$ is used. The optimal solution $X_{q,t}^f$ prescribes the takeoff-time for each aircraft for all scenarios. Communicating the ground hold for each aircraft in all scenarios to the airlines is impractical. Instead, we suggest that the central planner communicates the following information to the airlines regarding each of their flights:

- 1) Minimum ground holding time: The minimum ground hold that will be assigned to a slot s across all scenarios is $grd_{min}(s)$.
- 2) Expected ground holding time: This is the average ground holding time that will be assigned to the slot. For a slot s , the expected time is $grd(s)$.
- 3) Expected air holding time: The airborne queue that a slot will experience is dependent on the scenario, and the mean wait time for a slot s is denoted as $air(s)$.

Finally, depending on which scenario materializes, the actual departure time for each flight is updated through the course of the GDP. While the full scenario specific schedule is not known to the airline, it can make swaps based on $grd_{min}(s)$, $grd(s)$ and $air(s)$ for all slots s it owns. This optimization is similar to the one presented in Section II-B.1. The feasibility space for swaps is more restricted in the dynamic model [5] as the slots are dependent on the flight duration. Consequently, $feas(f,k) = 1$ if and only if $f, f_k \in \mathcal{F}_a$ and $dep(f) \geq dep(f_k) + grd(s_k)$ and $dur(f) = dur(f_k)$. We emphasize that the slot substitution with a dynamic formulation, is more restrictive than the static formulation. Only flights with the same flight duration can be swapped in the dynamic model and this is not a requirement for the static model slot substitution [5].

3) *2-step RHS*: It is possible to obtain a 2-step RHS solution by setting $C_{f,n}^g = C_n^g$ for the optimization in Section II-A.3. However, we introduce a slightly different formulation here. In the 1-step RHS, the static solution of the first stage is *aware* that a second stage static optimization will be performed. Thus, the solution jointly optimizes over both static formulations. This would result in longer duration flights getting higher delays when they are scheduled around the update time \tilde{t} . We avoid this flight-duration specific solution by running two sequential optimization. The primary motivation for this is to allow for maximum flexibility in slot substitutions (the second step in the 2-step RHS) and to maintain more fairness across flights of different duration. We describe the first step (performed by the central planner) in the 2-step RHS (more details in [20]).

- 1) Solve the static problem (with $C_{f,n}^g = C_n^g$) involving the full duration of GDP (till T)
- 2) Identify all aircraft taking off before \tilde{t} as stage 1 flights and allot them a deterministic ground holding time.
- 3) The flights not taking off in stage 1 are the stage 2 flights.
- 4) Solve the static optimization (with $C_{f,n}^g = C_n^g$) for all the stage 2 flights with the time horizon being \tilde{t} to T .

Do this for each of the scenarios that can get resolved in stage 2 as well for the set of unresolved scenarios.

- 5) These set of flights will get a scenario dependent ground holding time (which is same for all $q \in Q_{unres}$).

The key part here is that the first static is planned independently of the second static optimization and this gives us the desired flexibility to swap flights within a stage independent of their duration. After this first step, the airlines are given the following information for each of its flights $f_k \in \mathcal{F}_a$ - the stage $stage(f_k)$, $grd_{min}(s_k)$, $grd(s_k)$ and $air(s_k)$. For flights f_k in stage 1, $grd_{min}(s_k) = grd(s_k)$ because the holding is deterministic. The feasibility condition is $feas(f,k) = 1$ if and only if $f, f_k \in \mathcal{F}_a$ and $dep(f) \geq dep(f_k) + grd(s_k)$ and $stage(f) = stage(f_k)$. The rest of the optimization is same as in Section II-B.1.

III. RESULTS

From the nature of the formulations, it is clear that the 2-step models respect airline privacy more than the 1-step models. Also, for both the 1 and 2-step versions, the dynamic model is the least stable, RHS model is more stable and the static model is the most stable. In this section, we comment on the efficiency, measured in terms of total delay costs, of these six models with varying privacy and stability properties. We make the following observations.

- 1) For the 1-step models, the delay costs will always follow the order 1-step static < 1-step RHS < 1-step dynamic. This can be seen from the optimization formulations itself. All optimal static solutions are feasible solutions to the RHS and dynamic optimization. Similarly, the optimal RHS solution is feasible for the dynamic formulation. Hence the costs follow the inequality.
- 2) 2-step solutions usually have a lower delay cost than the corresponding 1-step solution. Intra airline swaps will always decrease the cost, but each airline must own a large enough number of flights, with high cost variability for such swaps to be made.
- 3) While the 1-step static always has a higher cost than the 1-step dynamic, the same trend is not true for the 2-step case, because of the different slot substitution flexibility of the static and dynamic model [5]. Depending on the intensity and duration of the GDP, either of them could have a lower cost. Theoretically, we cannot predict the cost of the RHS formulation in relation to the static or dynamic model.

Two step models incorporate privacy requirements and are expected to have a higher delay cost. The Price of Privacy is defined as the additional cost (in percentage) that a 2-step model has over the corresponding 1-step model. Similarly, a more stable model is expected to lead to poorer system performance, and hence higher cost. The Price of Stability is defined as the additional cost (in percentage) that a model has over the corresponding dynamic solution cost (which is the least stable). It is the additional price we pay for having a stable GDP allocation.

We evaluate these models using publicly-available schedule data [21] for LaGuardia Airport (New York City) on Feb 17, 2014. A 7 hour GDP from 7 am till 2 pm is simulated with the schedule arrival demand being 24, 31, 31, 36, 31, 32 and 36 for these 7 hours. The nominal capacity is 40 landings/hour and the reduced weather capacity is 20 landings/hour (similar to Fig. 2). We generate $2 * T - 1 = 13$ different probability distributions for the GDP of length $T = 7$. The branch q of the scenario tree is such that the capacity improves after time $t = q$. Each branch has a different expected duration of reduced capacity.

Denote $S = \{1, 2, \dots, 2T - 1\}$ as the set of scenario trees. Corresponding to every scenario tree $s \in S$, there is an associated probability π_q^s for every branch q in the scenario tree s . In every distribution, some of the branches are assigned a probability of 0.01 (almost no occurrence) and the remaining probability is split between the other branches.

For the first T distributions, $s = 1, \dots, T$,

$$\pi_q^s = \begin{cases} \frac{1-0.01 \times (T-s)}{s} & \forall q \in \{1, \dots, s\} \\ 0.01 & \forall q \in \{s+1, \dots, T\} \end{cases} \quad (25)$$

For the next $T - 1$ distributions, $s = T + 1, \dots, 2T - 1$,

$$\pi_q^s = \begin{cases} 0.01 & \forall q \in \{1, \dots, s\} \\ \frac{1-0.01s}{T-s} & \forall q \in \{s+1, \dots, T\} \end{cases} \quad (26)$$

To summarize, we create 13 different scenario trees, each with a different expected duration of reduced capacity. The ground hold cost for each aircraft is drawn from a normal distribution with mean 1 and standard deviation 0.25. The cost of air holding a flight is set at 2.5. The flights are randomly distributed between 10 airlines (to simulate the intra-airline substitutions in a 2-step model).

Fig. 3 shows the cost for all the six algorithms across different probability distributions (equivalently, the expected duration of reduced capacity). For all the 1-step models, the costs follow the expected order of dynamic \leq RHS \leq static. Also, as the expected duration of the GDP increases, the cost incurred is higher because more flights would be delayed if low capacity persists for longer. For the 2-step solutions, lower duration GDPs result in lower cost for dynamic model. As the expected GDP duration increases, the dynamic solution becomes more expensive. This is because more flights are affected and a greater swap flexibility of the static model becomes a more dominant factor than the benefits of an initial dynamic allocation. The 1-step dynamic is the socially optimal allocation, meaning that it gives the lowest total delay cost. The 2-step static is at the other end of the spectrum, and give the individual agents, airlines in this case, the most fair and equal allocation of resources while respecting their privacy.

Fig. 4 shows the impact of privacy and stability requirements. The static model is the least sensitive to privacy requirements. The 2-step static cost increases by less than 5% for most of the scenario trees. The 2-step dynamic model has the highest cost increase due to privacy, i.e. the highest increase in cost in comparison to the corresponding 1-step




Fig. 3. Delay cost comparison for all the 1-step (left) and 2-step (right) models. The x-axis is the expected reduced-capacity duration for the probability distribution.

dynamic model. This is mainly due to the low slot-exchange flexibility in the 2-step dynamic model which makes it significantly more expensive than than the 1-step solution (which doesn't involve any slot exchange).

There are two parts to the Price of Stability (PoS) plot (Fig 4, right). For the 1-step models, the PoS is always positive since the 1-step dynamic is always the lowest-cost model. The PoS for 1-step static (maximum of 68%) is higher than the 1-step RHS (maximum of 21%). For the 2-step models, the 2-step dynamic, which is the reference for all PoS computations does not have the lowest cost for scenarios where the low capacity exists for a longer duration (Fig. 3 (right)). So in 2-step models, low GDP duration requires a cost penalty if more stability is desired (upto 57% for static and 23% for RHS) whereas it effectively comes for 'free' for longer duration GDPs.




Fig. 4. Price of Privacy (left) and Price of Stability (right) for the GDP algorithms.

Fairness in the slot allocation is an interesting aspect of GDP planning. Two-step models are fair to airlines which have different delay costs. They are all treated equally by the central planner since the cost information is private and not available. However, there is a bias due to flight duration that is present in the dynamic model. Shorter duration flights tend to get more delay as they can be used more flexibly to adapt to evolving capacity scenarios. The 2-step RHS model is fair to flights of all duration whereas the 1-step RHS model introduces a small bias for flights that are planned to take off near the update time (delaying these flights will put them in the second stage and they can use the updated capacity information).

IV. CONCLUSIONS AND FUTURE DIRECTIONS

We considered GDP planning from a new perspective of asymmetric information exchange and stability of the slot

allocation. We then developed algorithms that accommodated varying privacy and stability requirements, and evaluated them using simulations of GDPs at LaGuardia Airport in New York.

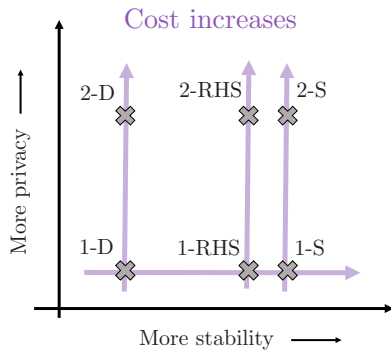


Fig. 5. Summary of the six GDP algorithms, in terms of efficiency (cost), privacy and stability. The arrows indicate the directions of cost increases. The vertical arrows correspond to the Price of Privacy, while the horizontal ones correspond to the Price of Stability.

Fig. 5 presents a summary of the models and their efficiency (in terms of total system delay costs). Two-step models consider airline privacy concerns, but are less efficient than the corresponding one-step models. A static solution is the most stable but it is not always the most efficient. The 2-step RHS model presents a good compromise between privacy and stability. It is also easy to implement, since it only requires information exchange at two time-steps, at the start of the GDP and at the update time.

Incorporating fairness across airlines with different cost structures (for example, a low-cost regional carrier and a network carrier) poses a challenge in one-step models. In addition, the proposed one-step solutions do not have incentives for truthful cost reporting. The development of equitable and truthful mechanisms is a focus of ongoing research.

REFERENCES

- [1] E. P. Gilbo, "Airport capacity: Representation, estimation, optimization," *Control Systems Technology, IEEE Transactions on*, vol. 1, no. 3, pp. 144–154, 1993.
- [2] V. Ramanujam and H. Balakrishnan, "Estimation of arrival-departure capacity tradeoffs in multi-airport systems," in *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*. IEEE, 2009, pp. 2534–2540.
- [3] J. Avery and H. Balakrishnan, "Data-driven modeling and prediction of the process for selecting runway configurations," *Transportation Research Record: Journal of the Transportation Research Board*, no. 2600, pp. 1–11, 2016.
- [4] A. Mukherjee and M. Hansen, "A dynamic stochastic model for the single airport ground holding problem," *Transportation Science*, vol. 41, no. 4, pp. 444–456, 2007.
- [5] V. Ramanujam and H. Balakrishnan, "Increasing efficiency though flexibility in airport resource allocation," in *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*. IEEE, 2014, pp. 1759–1766.
- [6] O. Richetta and A. R. Odoni, "Solving optimally the static ground-holding policy problem in air traffic control," *Transportation Science*, vol. 27, no. 3, pp. 228–238, 1993.
- [7] —, "Dynamic solution to the ground-holding problem in air traffic control," *Transportation Research Part A: Policy and practice*, vol. 28, no. 3, pp. 167–185, 1994.
- [8] C. Barnhart, D. Bertsimas, C. Caramanis, and D. Fearing, "Equitable and efficient coordination in traffic flow management," *Transportation Science*, vol. 46, no. 2, pp. 262–280, 2012.
- [9] D. Bertsimas and S. Gupta, "Fairness in air traffic flow management," in *INFORMS Meeting, San Diego, CA, USA*, vol. 32, no. 4, 2009, p. 7.
- [10] D. Bertsimas, V. F. Farias, and N. Trichakis, "On the efficiency-fairness trade-off," *Management Science*, vol. 58, no. 12, pp. 2234–2250, 2012.
- [11] C. N. Glover and M. O. Ball, "Stochastic optimization models for ground delay program planning with equity–efficiency tradeoffs," *Transportation Research Part C: Emerging Technologies*, vol. 33, pp. 196–202, 2013.
- [12] T. Vossen, M. Ball, R. Hoffman, and M. Wambganss, "A general approach to equity in traffic flow management and its application to mitigating exemption bias in ground delay programs," *Air Traffic Control Quarterly*, vol. 11, no. 4, pp. 277–292, 2003.
- [13] M. O. Ball, R. Hoffman, and A. Mukherjee, "Ground delay program planning under uncertainty based on the ration-by-distance principle," *Transportation Science*, vol. 44, no. 1, pp. 1–14, 2010.
- [14] T. Vossen and M. Ball, "Optimization and mediated bartering models for ground delay programs," *Naval Research Logistics (NRL)*, vol. 53, no. 1, pp. 75–90, 2006.
- [15] T. W. Vossen and M. O. Ball, "Slot trading opportunities in collaborative ground delay programs," *Transportation Science*, vol. 40, no. 1, pp. 29–43, 2006.
- [16] H. Balakrishnan, "Techniques for reallocating airport resources during adverse weather," in *Decision and Control, 2007 46th IEEE Conference on*. IEEE, 2007, pp. 2949–2956.
- [17] J. Cox and M. J. Kochenderfer, "Optimization approaches to the single airport ground hold problem," in *Journal of Guidance, Control, and Dynamics*, Kissimmee, FL, 2015.
- [18] —, "Ground delay program planning using markov decision processes," *Journal of Aerospace Information Science*, 2016.
- [19] V. Ramanujam, "Estimation and tactical allocation of airport capacity in the presence of uncertainty," Ph.D. dissertation, Massachusetts Institute of Technology, 2011.
- [20] K. Kavassery Gopalakrishnan, "Analysis of efficiency and fairness in stochastic ground holding models," Master's thesis, Massachusetts Institute of Technology, 2016.
- [21] Federal Aviation Administration, "ASPM database," 2014, <https://aspm.faa.gov/apm/sys/main.asp>.