Airline Disruption Management with Delay Ledgers

Christopher Chin, Hamsa Balakrishnan
Massachusetts Institute of Technology
Cambridge, MA, USA
{chychin, hamsa}@mit.edu

Karthik Gopalakrishnan
Stanford University
Stanford, CA, USA
gkarthik@stanford.edu

Max Z. Li
University of Michigan
Ann Arbor, MI, USA
maxzli@umich.edu

Abstract—The impact of disruptions may result in reduced capacities at airports, forcing airlines to revise schedules and delay flights. However, due to myriad factors (e.g., passengers who may miss their connections, remaining flights to be performed by an aircraft, high-valued passengers with elite statuses), a delayed flight may be more or less costly to an airline, even when compared to another similarly delayed flight. Currently, identifying optimal slot swaps between airlines requires sharing the airline-specific delay cost of each flight. However, this is not amenable as sharing these private delay costs could reveal sensitive business practices. We propose the use of a procedure called the Delay Ledger (DᴇLᴇᴅ) which enables airlines to identify a set of beneficial slot swaps across a network of airports which guarantees improvements in terms of private delay costs while ensuring that no private flight-specific valuations are shared. DᴇLᴇᴅ is guaranteed to lower airline delay costs, incentivizes truthful airline participation, and supports flexible airline privacy preferences. We evaluate DᴇLᴇᴅ across 30 days with 8 major US airlines, resulting in average reductions in private delay costs of 8-22% per day compared to current approaches.

Keywords: Traffic flow management; Airport ground holding; Flight delays; Airline disruption management

I. Introduction

Poor weather often leads to a decrease in the number of aircraft arrivals and departures that can be safely accommodated at an airport. As a consequence, airlines operating at affected airports are unable to execute their desired flight schedules, and flights incur delays. The appropriate flight delays are identified by air traffic managers to ensure that the revised schedule can satisfy reduced arrival and departure capacities. Several optimization-based approaches have been proposed to identify the most efficient reassignments of flights (e.g., single and multi-airport airport ground holding problems [1, 2], air traffic flow management problem [3]).

From an airline's perspective, prior to the disruption, the airline had a set of flights with scheduled arrival and departure times. Now, in the revised schedule, these flights may have been delayed. One of the major objectives of an airline during such disruptions is to minimize the impact of these delays. In particular, the airline needs to manage several consequences due to the revised schedule (e.g., crew time-outs, passengers missing their connections, delays in subsequent legs of the same aircraft) [4]. Naturally, these negative impacts are a function of the delay severity (e.g., delay minutes), but are

also not the same for all aircraft [5]. For example, a delayed flight that has few connecting passengers could be less costly for an airline than another delayed flight which has a significant number of passengers with tight connections. Thus, the objective of an airline is to minimize the delay costs, i.e., the impacts of these delays, rather than the delay itself. To achieve this goal, airlines often swap flights through a process called *Intra-Airline Substitution* [6] to reduce their delay costs. Fig. 1 shows an example where airlines can use Intra-Airline Substitution to minimize delay costs.
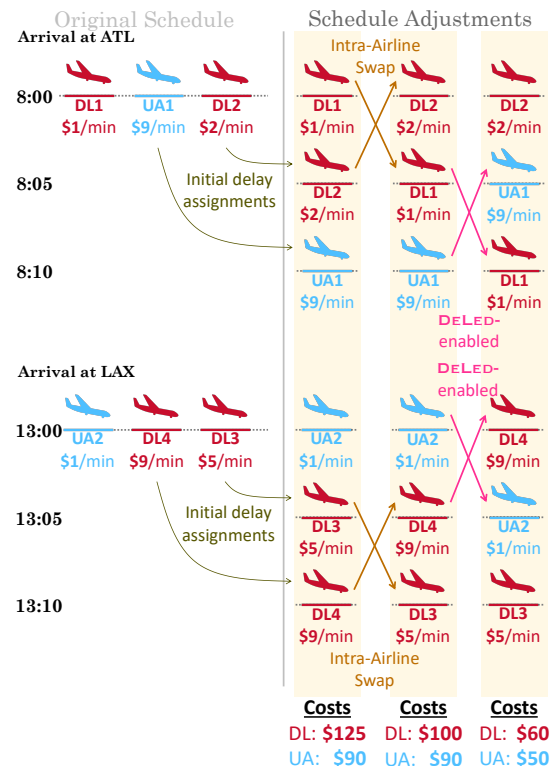

Fig. 1. Intra-Airline Substitution and the swaps between airlines that DᴇLᴇᴅ will enable. Notional per-delay minute flight-specific valuations are given beneath flight number.

At Atlanta (ATL), Delta Air Lines (DL) has two flights: DL1 and DL2 with the same scheduled arrival time of 8:00. DL2 has a higher valuation than DL1. After the first schedule adjustment, DL1 keeps its 8:00 arrival, whereas DL2 is delayed

until 8:05. Delta could then swap these flights so that the less valuable DL1 incurs the 5-minute delay, rather than DL2. At Los Angeles (LAX), Delta swaps DL3 and DL4 to reduce costs. Intra-airline substitution allows airlines to reduce their delay costs, and is widely adopted in practice [7]. However, we hypothesize that the slot swapping benefits could be much higher if airlines could swap slots with each other. In Fig. 1, after intra-airline swapping, Delta and United (UA) could reduce their costs by swapping with each other: UA1 could move to DL1's slot at LAX, and DL4 could move to UA2's slot at ATL. Note that the total delay of each airline remains unchanged, but the delay *costs* of each airline have decreased.

One approach to implement these swaps is for a third-party central entity to solve a centralized optimization problem using each airline's internal, private delay costs [8, 9]. However, airlines may be wary of sharing the exact delay costs of their flights to a central entity, as it may reveal sensitive business practices. Another approach would be for airlines to propose candidate slot swaps to each other. If a proposal is beneficial to both airlines, then they could agree to such a swap without revealing any private information. However, a systematic process for airlines to propose acceptable swaps — without explicitly knowing the internal delay costs of the other airlines' flights — remains unknown. In fact, a naive approach may require a combinatorially large number of proposals to be evaluated, leading to computational tractability concerns. Thus, the fundamental challenge in slot swaps among airlines is to identify swaps that are feasible to all airlines involved *without* explicitly using flight-specific private delay costs.

A. Our approach: DELED

We propose a new slot swapping procedure, DELED, based on a *Delay Ledger* that is able to overcome aforementioned limitations. DELED enables a large number of slot swaps among airlines, and by design incentivizes truthful participation while ensuring privacy in flight-specific delay costs. DELED centers on three key ideas:

- Slot swaps between airlines can occur across different airports (as shown in Fig. 1). Since disruptions occur on a day-by-day, week-by-week basis, slot swaps could also occur across scheduling *rounds*. Thus, an airline may accept incurring slightly more delay in this round, in exchange for advantageous slot swaps in future rounds. We can leverage this flexibility to unlock more slot swap opportunities that reduce delay costs even when complete information on delay costs are not revealed. We use a *Delay Ledger* to track the cumulative additional delay incurred by each airline; this value could be negative for airlines that benefited from advantageous swaps.
- In every round, one *airline* is assigned by the Delay Ledger to coordinate slot allocations for all airlines. By doing so, the *coordinating airline* can use all of its own

private information to identify slot swaps that minimize its delay impact. The Delay Ledger ensures that the coordinator role is appropriately passed around.
- DELED constrains the coordinating airline's actions such that non-coordinating *participating airlines* do not suffer an increase in delay costs. Participating airlines share a coarse version of flight delay costs with the coordinating airline by designating some flights as high-priority (i.e., can only be moved earlier) and others as more flexible (could be delayed further).

B. Our contributions

The main contribution of our work is the design of DELED, a Delay Ledger-based procedure to enable slot swaps among airlines. DELED has several desirable properties:

- *Individual rationality*, meaning that every airline is guaranteed to do no worse in terms of private delay costs;
- *Privacy preserving* by enabling airlines to share only a coarse version of flight-specific delay costs.

Our experiment with historical airline data combining traditional Intra-Airline Substitutions with DELED shows significant benefits for all airlines, with average reductions in private delay costs of 8.3-22.3% per round. Thus, our work suggests that it may be possible to enact mutually beneficial slot swaps among airlines without compromising privacy.

II. RELATED WORKS

Inter-airline slot exchanges have been modeled before with at-most, at-least (AMAL) trades [10, 11]. With an AMAL offer, airlines indicate willingness to move flight f_1 to a time period no later than t_1, as long as flight f_2 is moved to a time period no later than t_2. Slot credit substitution (SCS) has been adopted in practice with Ground Delay Programs (GDPs) [12]. Under SCS, airlines offer to relinquish a slot (by canceling a flight) in exchange for obtaining a desired replacement slot. In contrast to both AMAL trades and SCS, we do not require airlines to propose trades but instead allow them to assign a privatized priority level to flights. We also track the delay impacts to each airline across arbitrarily long time frames with the delay ledger and leverage that information within our mechanism.

Perspectives that closely align with our ledger-based mechanism include (i) delay credit models, where airlines pay or are compensated based on delay decreases or increases, respectively [13, 14], (ii) modeling delays as "goods" that airlines can buy from a central coordinator (e.g., the FAA) on a per-flight basis [15], and (iii) an iterative airspace pricing and allocation scheme [16]. We examine day-of-operation adjustments, and do not require airlines to share private information to a central entity. We also consider scheduling benefits across rounds.

Our model also shares some similarities to the User Driven Prioritization Process (UDPP) concept proposed as a part

of SESAR [5]. UDPP involves individual airlines assigning a priority value to a subset of flights (Fleet Delay Apportionment/Fleet Delay Reordering) and protects high-priority flights (via the *selective flight protection* feature) by finding opportunities to assign additional delays to low-priority flights. Pre-operational trials with airlines have shown positive results [17]. Our work is distinct from UDPP, as we consider the multi-airline setting case where coarse flight priorities are used to enable privacy-preserving slot swaps among airlines.

III. Methodology

A. Problem setup and assumptions

Consider a network of airports served by flights from several airlines. We assume that each airline's *original schedule* (or nominal schedule) is public, as in reality, such schedules are regularly shared as part of the collaborative decision-making process [18], or through sources such as OAG [19]. Reductions in airport capacity can cause flights to be delayed. Since our work focuses on the design of DeLed and its deployment on a realistic set of airline schedules, we assume that airport capacities are deterministic, and known for the entire round. For future work, we could remove this assumption in a straightforward manner, e.g., by partitioning each round into sub-time intervals then using a rolling horizon implementation of DeLed, or by using chance-constrained optimization approaches [20]. We assume that revised slots and associated delays are assigned by a centralized multi-airport ground holding problem (MAGHP) [2]; we call this the *Baseline Solution*. We note that other commonly-used baselines include first-come-first serve and ration-by-schedule procedures [13, 16].

When flights are delayed, we can use the original schedule to calculate the *public schedule delays* of each airline. We assume that each airline has a different *valuation* for each of its individual flights, i.e., one minute of delay on a flight can be more costly than one minute of delay on another flight. In addition to the reasons for a high valuation discussed in Section I, another key example is that when a flight is delayed, it can propagate that delay to connecting flights operated by the same aircraft (i.e., late-arriving aircraft delay). Thus, an airline may put a high valuation on Flight A flown by an aircraft that is scheduled to immediately turn around and operate as Flight B. We refer to Flight B as a *connecting flight* of Flight A. We assume that airlines know and publicly share the sequence of flights that each aircraft will operate; this assumption is standard given information-sharing standards to enable, e.g., collaborative decision-making [18]. We do not consider aircraft reassignment by airlines, which could alter connecting flights. For future work, we could incorporate aircraft reassignments in a rolling horizon, giving airlines more flexibility mid-round.

Since flight valuations are directly related to potentially sensitive business practices, we assume that individual flight valuations are private to the airlines and thus will not be shared. However, we assume that airlines could share *flight priorities*, which are coarse versions of individual flight valuations, as long as they were guaranteed to benefit from it. Thus, the metric we focus on is *private delay costs*, i.e., the sum of public schedule delays scaled by the private valuation for each flight. We assume that airlines want to minimize their private delay costs; thus, after airlines receive their new set of slots from the Baseline Solution, they perform an intra-airline swap to reduce their delay costs (as shown in Fig. 1). We call the modified schedule after all airlines perform Intra-Airline Substitution an *Intra-Airline Solution*.

B. DeLed overview

To further reduce private delay costs, airlines could engage in a collaborative process where slots are traded. DeLed presents a fundamentally different approach to collaborative airline disruption management, wherein airlines can reduce their private delay costs relative to the DeLed *Input Solution*, which we assume to be the Intra-Airline Solution. Fig. 2 outlines the overall process. After the Input Solution is determined, DeLed assigns a *coordinating airline* (coordinator) to perform schedule adjustments on behalf of itself and other *participating airlines* (participants). The coordinator keeps its role for an entire round, which is a time interval of reasonable length, e.g., from a few hours to an entire day. Within a round, the coordinator can adjust participants' schedules as well as its own for disruption management and schedule recovery. Participants provide flight priorities which ensure the coordinator does not increase their *private* delay costs. Coordinators can use their own private delay costs to inform schedule adjustments.

We note that DeLed schedule adjustments may come at the cost of *public* delay debt owed to some participants, which is recorded in a cumulative *delay ledger*. The delay ledger tracks the change in public schedule delay for each airline relative to its delay in the input schedule, cumulative across all rounds. To account for differences in airline size, the delay ledger tracks *average* delays. After each round, the coordinator role is reassigned to the airline with the highest delay ledger value, i.e., the airline that experienced the largest cumulative average increase in public delay relative to input schedules. The delay ledger distributes the coordinator role, ensuring overall balance *across airlines* and *across time*: An airline that has been a participant for multiple rounds and has accumulated excess delay debts will eventually be assigned the coordinator role. Critically, this delay ledger enables the cooperative trading of slots between airlines *across multiple rounds* and *across multiple resources in a network of airports*.

In each round, the coordinator receives a reduction in its private delay costs relative to the Input Solution. However, DeLed also ensures that the private delay costs of participants

do not increase relative to the input. That is, by participating in DELED, participants will never be worse off in terms of private delay costs. DELED guarantees this by enabling participants to share a *coarse* version of their private valuations with the coordinator. We refer to *flight priorities* as a coarse version of private valuations because it maps the actual private valuations of a flight into broad *priorities*: high, medium, or low. DELED constrains the coordinator such that it *cannot* increase the private delay costs of participants, even though it *can* increase the public schedule delays of participants, subject to a maximum-allowable increase. Note that airlines could decide to *not* participate in DELED, in which case their schedule would be left untouched by DELED. However, we expect that airlines will want to participate, given the possibility that their private delay cost will decrease. In the next subsection, we discuss the technical details of how this constraint and other DELED features are implemented.

C. DELED model details

Recall from above that broadly, DELED consists of a coordination step (assignment of coordinator and participant roles, and coordinator-mediated schedule adjustments constrained by participants' submitted flight priorities) followed by a delay ledger update step. We now detail the formulation of the linear program that the coordinator runs, as well as the specifications unique to DELED which ensure that flight priorities of the participants are taken into account.

1) Variables and objective: We first utilize several constraints from the MAGHP: airport departure and arrival capacities, minimum flight times, minimum connection times between flights operated by the same aircraft, and feasible departure and arrival times [2]. For our case study, we describe in Section V how these constraints are populated by relevant data. The set of all flights is denoted as F. The set F_c contains all flights that belong to the coordinating airline. P is the set of all participants in a round. For every round, the set F_p contains the flights belonging to the participating airline $p \in P$. The coordinator knows its own flight valuations v_f, but does not know the flight valuations of participants. Instead, the coordinator knows the priority of participants' flights. $F_{p,h}$, $F_{p,m}$, $F_{p,l}$ contain high, medium, and low priority flights of p, respectively. Flights in $F_{p,h}$ have $v_f \geq b$, flights in $F_{p,m}$ have $a < v_f < b$, and flights in $F_{p,l}$ have $0 \leq v_f \leq a$. The coordinator also knows several public delay values from the Input Solution: Total delay D_{total}, participant delay D_p, and individual flight delay D_f.

For each round, the coordinator solves an optimization problem to assign a new delay d_f to each flight, subject to several constraints. The objective of the coordinator is to minimize its own private delay costs, as shown in (1). The objective does not contain any terms for the participants; the priorities of the participants are addressed in the constraints.

$$\min \sum_{f \in F_c} v_f d_f \tag{1}$$

subject to

$$\sum_{f \in F} d_f \leq D_{\text{total}} \tag{2}$$

$$\sum_{f \in F_p} d_f \leq c_{\text{incr}} D_p \qquad \forall p \in P \tag{3}$$

$$d_f \leq D_f \qquad \forall f \in F_{p,h} \cup F_{p,m} \quad \forall p \in P \tag{4}$$

$$b \sum_{f \in F_{p,h}} (D_f - d_f) + a \sum_{f \in F_{p,m}} (D_f - d_f) \geq$$
$$a \sum_{f \in F_{p,l}} \max (d_f - D_f, 0) \quad \forall p \in P \tag{5}$$

2) Constraints: Constraint (2) stipulates that the total delay in the resulting DELED Solution must not exceed the total delay of the Input Solution. This is to avoid under-utilizing system capacity (e.g., runway slots). Constraint (3) ensures that the total public delay assigned to each participant is no more than some constant (c_{incr}) times the participant's total delay in the Input Solution; this is to upper bound the amount of excess public schedule delay penalties that the coordinator can levy on a participant. Note that in order for Constraint (2) to hold, if the coordinator increases public delay for one participant, it must reduce that amount of public delay elsewhere. Constraint (4) requires that high and medium-priority flights can only shift earlier in the schedule, not later. There is no such requirement for low priority flights, i.e., they can move earlier or later. Critically, Constraint (5) ensures that each participant does not experience an increase in private delay cost. The left-hand side of Constraint (5) records the *minimum decrease* in private delay costs, given that high priority flights have $v_f \geq b$ and medium priority flights have $v_f \geq a$. The right-hand side of Constraint (5) records the *maximum increase* in private delay costs stemming from moving low priority flights later in the schedule, as low priority flights have $v_f \leq a$. Note that this constraint is structured such that participants will not experience an increase in private delay cost. Constraint (5) can be easily modified to accommodate airlines submitting different prioritizations (e.g., more priority levels or different ranges for each priority).

We emphasize that the coordinator can optimize (1) subject to (2)-(5) using their true private delay costs $v_f, f \in F_c$; thus, the coordinator will likely see a large reduction in private delay costs. In return for submitting flight priorities, participants are guaranteed that high and medium priority flights will not be delayed further, and their private delay costs will not increase. The delay ledger ensures that the most cumulatively penalized participant in terms of public schedule delay will be assigned the coordinator role in the subsequent round.
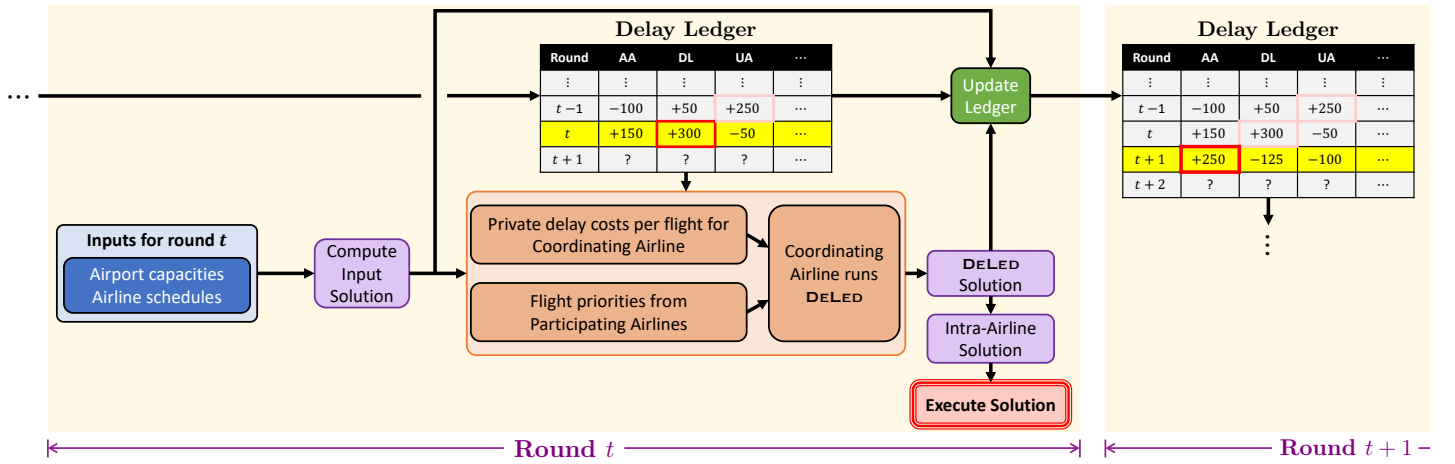
Fig. 2. Process diagram for DeLed.

IV. Integrating DeLed with Intra-Airline Substitution

In this section, we detail how DeLed fits in with existing intra-airline schedule adjustment procedures. Recall from Section III-B that we assume that the input to DeLed is the Intra-Airline Solution after a baseline MAGHP is run. It is possible to skip this first intra-airline step and directly input the Baseline Solution into DeLed. However, the Baseline Solution typically has much higher private delay costs than the Intra-Airline Solution. Thus, participants will generally receive lower private delay costs from DeLed when the Intra-Airline Solution—not the Baseline Solution—is used as the input. This is because DeLed guarantees that participants will not see an increase in private delay costs relative to the input. Hence, a Intra-Airline Substitution step should be done before DeLed.

We have motivated the utility of inputting the Intra-Airline Solution into DeLed. After the DeLed Solution, we assume that airlines will want to further reduce private delay costs by performing another round of intra-airline substitutions. Because participants shared flight priorities with the coordinator, many of their high priority flights could have been moved earlier; this final intra-airline step allows airlines to perform refinements to the schedule using their internal private valuations. For example, airlines can swap two high priority flights with different valuations. After all airlines resubmit their adjusted schedules, these schedules are aggregated into the final, public schedule to be executed in this round. Recall that the delay ledger is updated by comparing the DeLed Solution to the Input Solution; thus, the final intra-airline step does not affect the delay ledger.

V. Data and experimental setup

To evaluate the performance of DeLed, we first require flight schedules and airport capacities to generate the Input Solution. We start with the Bureau of Transportation Statistics' (BTS) Marketing Carrier On-Time Performance database for May 2019, which includes flight-level information on origin, destination, airline, aircraft tail number, along with scheduled and actual departure and arrival times [21]. We include flights that depart from or arrive at one of the Federal Aviation Administration (FAA) Core 30 airports and omit flights between non-Core 30 airports. We also omit flights from two airlines that had a small market share in the FAA Core 30: Allegiant Air (less than 0.3% of flights) and Hawaiian Airlines (less than 1.1%). We select flights that depart or arrive within 06:00 and 16:00 in Eastern Time (UTC-5:00). Each round spans this time period, and we use one round per day, for each day in May 2019. After this filtering, the input consists of 4,515 flights on average per round. Note that the choice of the time period is arbitrary and would likely be strategically chosen.

For each flight, we use the computerized reservation systems (CRS) times as the scheduled departure and arrival times. We discretize time into 15-minute bins to simplify airport capacity constraints. When adjusting schedules, we also need to know the earliest and latest feasible departure and arrival times. We estimate the earliest departure time as the maximum of (i) the scheduled departure time or (ii) the actual departure time minus *transferable* delays. We consider "National Aviation System (NAS)" and "Late-arriving aircraft" (airlines report the cause of delays to BTS) as transferable delays. We assume that transferable delays are preventable with schedule adjustments. For example, NAS delays could be reduced by swapping flights, and late-arriving aircraft delays could be mitigated if the delay of the preceding flight was reduced. This requires knowledge of connecting flights, which we identify by using the aircraft tail number of each flight. Each pair of connecting flights has a minimum-required connecting time, which we estimate as the actual connecting time minus any transferable

delay of the succeeding flight. For the latest feasible departure time, we assume that each flight can be delayed up to 5 hours. Then, to get the earliest and latest feasible arrival times, we add the actual travel time to the earliest and latest feasible departure time, respectively. Our choice for the maximum flight-specific departure delay of 5 hours is arbitrary; in practice, each airline could set its own maximum value. Airlines that set very low maximum delays would essentially be opting out of DeLed, and thus would be unlikely to become coordinator.

Besides estimating the demand for each round, we estimate airport capacities as the actual number of departures and arrivals per 15-minute period within our filtered set of flights in the BTS data, at each of the 30 airports. To avoid underestimating capacity during periods of low demand in the data, we set a minimum capacity of two flights per 15-minute period. Note that this estimate of capacities for our subset of flights do not reflect actual airport capacities, which are strictly higher since we do not include all flights and airlines.

To demonstrate the reduction of private delay costs via DeLed, each flight must be assigned a private valuation. Since we clearly do not have access to such airline-internal data, we assume nothing regarding the private valuations, and assign a random valuation to each flight drawn from a uniform distribution between 1 and 9, inclusive. It is straightforward to substitute these randomly-generated valuations with actual valuations, e.g., dollar amounts per delay minutes per flight. To convert these valuations v to priorities, we set $a = 3$ and $b = 7$, such that high-priority flights have $v \in [7, 9]$; medium priority flights have $v \in [4, 6]$; and low priority flights have $v \in [1, 3]$. Finally, we set $c_{\text{incr}} = 1.2$ such that the public delay of an individual airline can increase by up to 20% within each round, but the total public delay must not increase relative to the DeLed input. Note that priority thresholds a and b need not be constant for each airline, and could be submitted to the coordinator as additional information, resulting in airline-specific versions of constraint (5).

VI. Results and discussion

We evaluate the performance of DeLed by computing private and public delays for the Baseline, Intra-Airline only, Intra-DeLed-Intra, and Minimum Private Cost Solution. The Minimum Private Cost Solution is the hypothetical best-case performance if all airlines shared exact private flight valuations in a centralized MAGHP. The top-half of Figure 3 compares the performance of the models in terms of total public delay, calculated relative to the original schedule. As expected, the Minimum Public Delay Solution has the lowest total public delay. The Intra-Airline Solution has higher public delay than the Minimum Public Delay Solution. This is because when airlines perform Intra-Airline Substitution, it may increase delays for a low priority flight with subsequent flights operated by the same aircraft, so that a high priority flight can be

moved earlier. Because of the higher valuation for the high priority flight, even though public delays have increased, the airline's private delay costs have decreased. The Intra-DeLed-Intra Solution has slightly higher public delays than the Intra-Airline Solution; recall that DeLed does not increase public delay relative to its input, so this increase is due to the final Intra-Airline swapping step in Intra-DeLed-Intra. The Minimum Private Cost Solution has the highest public delay, as it optimizes for private delay cost, rather than public delay.

The bottom-half of Fig. 3 shows private delay costs. The trends across the models are opposite of the trends for public delay: The Minimum Private Cost Solution naturally has the lowest private delay cost, while the Minimum Public Delay Solution has the highest. The Intra-Airline Solution results in a 24.2% reduction in private delay cost on average, representing what airlines can achieve on their own through internal swaps only. However, the Intra-DeLed-Intra Solution yields an additional 12.2% reduction in private delay costs on average, demonstrating the benefits of DeLed. We also note that the Minimum Private Cost Solution will never be achieved in reality, as it requires all airlines to publicly share exact private flight valuations.
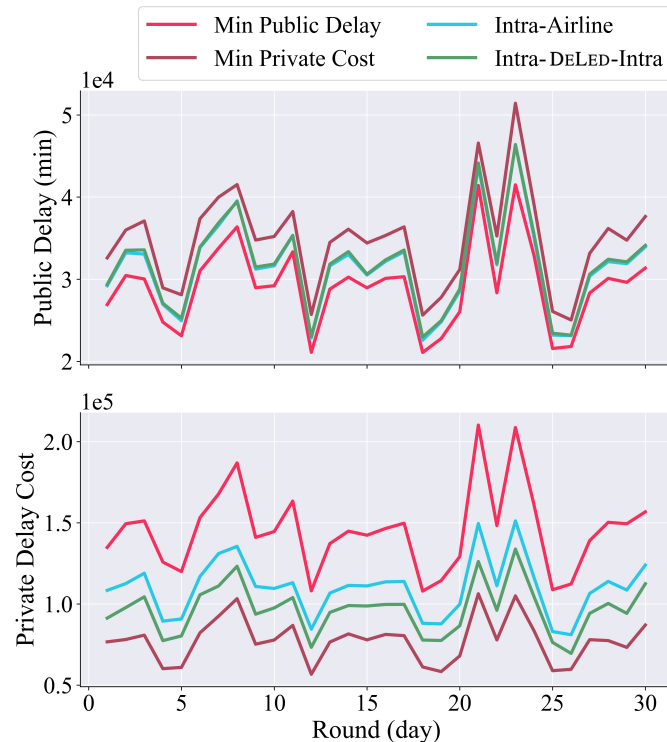

Fig. 3.   Public Delay and Private Delay Cost Comparison

While we know that Intra-DeLed-Intra performs well on an aggregate basis, we now evaluate the impacts on coordinators and participants. Fig. 4 shows the absolute (top) and percent

(bottom) change in private delay cost relative to the input to DeLed (Intra-Airline Solution). Private delay cost changes for the coordinator in each round is color-coded by airline and shown with a filled circle, while private delay cost changes for participants are given by boxplots for each round. We first note that across the 30 rounds (recalling that for this experiment, each round is an operational day), every airline was assigned coordinator at least once. Across all rounds, coordinators always see significant reductions in private delay costs. Additionally, participants generally receive a reduction in private delay costs as well, and at worst experience no change in private delay cost. In absolute terms, the larger airlines (e.g., American, United, and Delta) see a greater reduction in private delay cost. Some smaller airlines such as Spirit and Frontier have modest reductions in private delay cost in absolute terms, but receive much larger percent reductions.

While the coordinator cannot increase the total public delay relative to the Input Solution, the coordinator can increase the public delay of participants. In addition, the final intra-airline swapping step can modestly increase public delays. In total, the increase in public delay is relatively small, with Southwest incurring the largest average increase of 1.1% (see first row of Table I). Some airlines see a *reduction* in public delay by minimizing connecting delays when they are assigned to be the coordinator. Every airline receives an overall reduction in average private delay cost; this reduction is significantly larger than any increase in public delay. As an example, American Airlines held the coordinator role once, received a 8.3% reduction in private delay cost on average in exchange for an average increase of 0.6% in public delay. All other airlines have even greater benefits under Intra-DeLed-Intra.

Finally, we emphasize that most of the private delay cost reductions comes from DeLed, rather than the final intra-airline swap. For example, on average, Delta received a 12% reduction in private delay cost from DeLed and a 1.8% reduction from intra-airline swapping. The last two rows of Table I show the change in private delay cost when airlines are participants or coordinators. For each airline, the reduction in private delay cost is higher when they are coordinator, but each airline still receives a reduction in private delay cost when they are a participant.

VII. CONCLUDING REMARKS

When faced with airspace resource and capacity constraints, airlines simultaneously compete against and cooperate with each other to operate their respective schedules. We propose DeLed, a procedure that enables airlines to jointly optimize schedules across a network of airports. DeLed guarantees improvements in terms of private delay costs (i.e., the impact of delays experienced by an airline given different internal, private valuations of each flight), while ensuring that no private flight-specific valuations are shared. Additional schedule delay penalties are minimal in comparison to private delay cost improvements. The key features of DeLed are (i) the use of a delay ledger to keep track of additional delay penalties and to iteratively compensate airlines via the coordinator role, and (ii) the ability for airlines to share only coarse, aggregate flight priorities in lieu of sensitive, private flight-specific valuations. We evaluate DeLed through a 30-day case study across 30 airports involving 8 major US airlines, with average reductions in private delay costs ranging from 8.3% to 22.3% per round, and with average schedule delay increases of less than 1.1%.

In terms of future work, we can explore relaxing assumptions of deterministic and known airport capacities via a rolling horizon implementation within a round, along with incorporating chance-constrained or scenario-based optimization techniques. Other directions for future work include taking into consideration crew assignments, since this will reduce the set of possible schedule adjustments. We can also experiment with more realistic flight-specific valuations, mirroring how an airline might value the impact of delays on a specific flight.

REFERENCES

[1] 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.

[2] P. B. Vranas, D. J. Bertsimas, and A. R. Odoni, "The multi-airport ground-holding problem in air traffic control," *Operations Research*, vol. 42, no. 2, pp. 249–261, 1994.

[3] D. Bertsimas, G. Lulli, and A. Odoni, "An integer optimization approach to large-scale air traffic flow management," *Operations research*, vol. 59, no. 1, pp. 211–227, 2011.

[4] International Air Transport Association, "Disruption Management Manual (DMM)," 2022.

[5] SESAR, "User Driven Prioritisation Process," 2018, accessed: February 2022.

[6] 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.

[7] Federal Aviation Administration, "Flight Schedule Monitor User's Guide," 2016.

[8] T. W. M. Vossen and M. O. Ball, "Slot trading opportunities in collaborative ground delay programs," *Transportation Science*, vol. 40, no. 1, pp. 29–43, 2006.

[9] H. Erkan, N. K. Erkip, and Özge Şafak, "Collaborative decision making for air traffic management: A generic mathematical program for the rescheduling problem," *Computers & Industrial Engineering*, vol. 137, p. 106016, 2019.

[10] D. Bertsimas and S. Gupta, "Fairness and collaboration in network air traffic flow management: An optimization approach," *Transportation Science*, vol. 50, no. 1, p. 57–76, 2016.

[11] H. D. Sherali, J. M. Hill, M. V. McCrea, and A. A. Trani, "Integrating slot exchange, safety, capacity, and equity mechanisms within an airspace flow program," *Transportation Science*, vol. 45, no. 2, p. 271–284, 2011.

[12] R. Sankararaman, "Impact assessment of dynamic slot exchange in air traffic management," Ph.D. dissertation, University of Maryland, College Park, 2004.
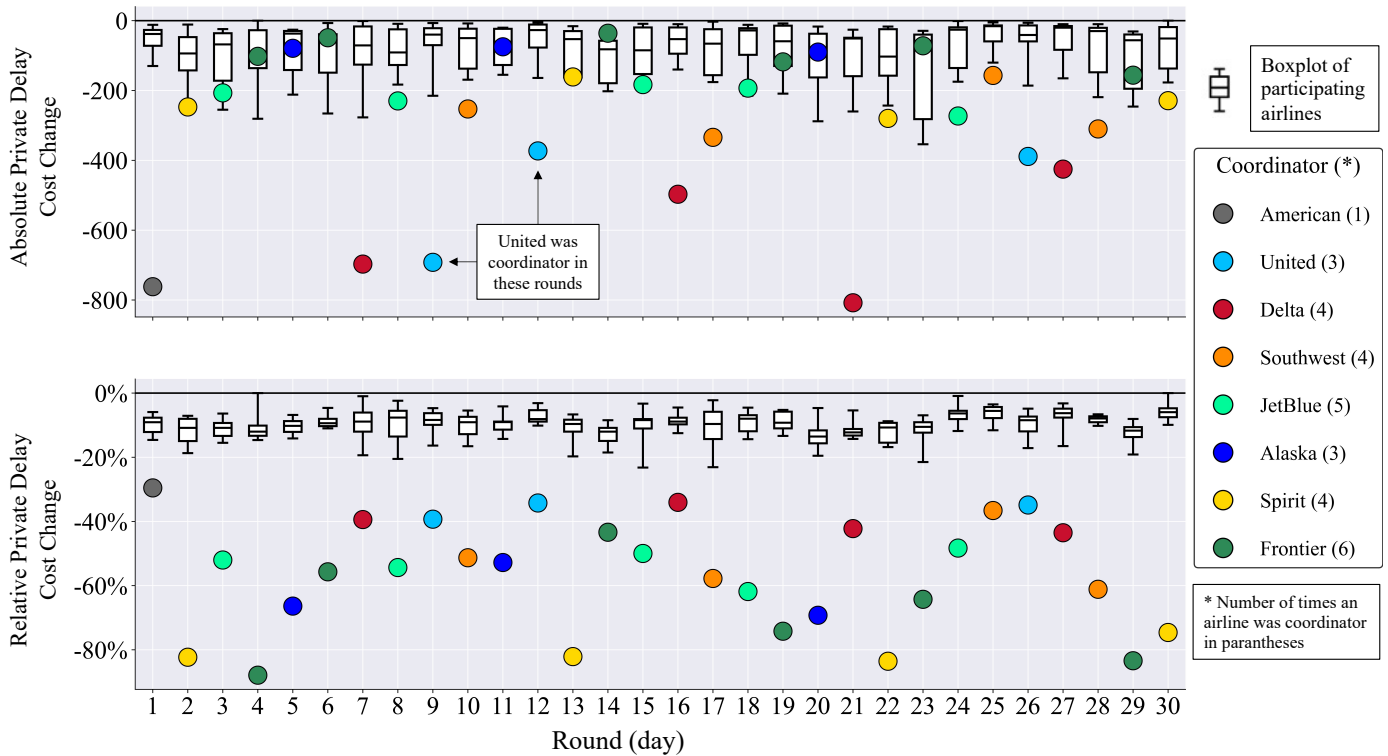
Fig. 4.   Absolute and relative change in delay costs for coordinators (ovals) and participants (boxplots). Extent of boxplot whiskers indicate minimum and maximum change in delay cost per round among participants. Results shown for Intra-DELED-Intra.

TABLE I
AIRLINE BREAKDOWN OF CHANGES IN PUBLIC DELAY INCREASES AND PRIVATE COST, RELATIVE TO INTRA-AIRLINE SOLUTION.

		American	United	Delta	Southwest	JetBlue	Alaska	Spirit	Frontier
Proportion of total flights (%)		30.3	19.3	24.5	13.5	3.6	2.8	3.1	1.9
Average change in public delay (%)		0.6	1.0	1.1	-0.8	0.5	0.4	-0.7	-4.6
Average private cost change	**Overall (%)**	**-8.3**	**-12.8**	**-13.6**	**-16.3**	**-16.8**	**-16.5**	**-18.6**	**-22.3**
	With Intra-Airline Substitution (%)	-1.8	-2.3	-1.8	-1.0	-0.3	-1.3	-0.2	-0.3
	With DELED (%)	-6.7	-10.6	-12.0	-15.4	-16.5	-15.4	-18.4	-22.0
	When pariticipant (%)	-7.6	-10.2	-9.5	-10.8	-9.5	-11.4	-9.1	-10.8
	When coordinator (%)	-29.6	-36.2	-39.8	-51.7	-53.3	-62.8	-80.7	-68.2

[13] L. Castelli, R. Pesenti, and A. Ranieri, "The design of a market mechanism to allocate air traffic flow management slots," *Transportation Research Part C: Emerging Technologies*, vol. 19, no. 5, pp. 931–943, 2011.

[14] P. Pellegrini, L. Castelli, and R. Pesenti, "Secondary trading of airport slots as a combinatorial exchange," *Transportation Research Part E: Logistics and Transportation Review*, vol. 48, no. 5, pp. 1009–1022, 2012.

[15] R. Mehta and V. V. Vazirani, "An incentive compatible, efficient market for air traffic flow management," *Theoretical Computer Science*, vol. 818, pp. 41–50, 2020.

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

[17] N. Pilon, L. Guichard, Z. Bazso, G. Murgese, and M. Carré, "User-Driven Prioritisation Process (UDPP) from advanced experimental to pre-operational validation environment," *Journal of Air Transport Management*, vol. 97, p. 102124, 2021.

[18] Federal Aviation Administration (FAA), "CDM Hub System," 2022, accessed: February 2022.

[19] Official Airline Guide (OAG), "Global Airline Schedules," 2022, accessed: February 2022.

[20] J. Chen, L. Chen, and D. Sun, "Air traffic flow management under uncertainty using chance-constrained optimization," *Transportation Research Part B: Methodological*, vol. 102, pp. 124–141, 2017.

[21] Bureau of Transportation Statistics, "Airline On-Time Statistics and Delay Causes," 2022.