

Integrated Schedule Planning for Regional Airlines Using Column Generation

Alberto Santini^{1,2} and Vikrant Vaze³

¹Department of Economics and Business, Universitat Pompeu Fabra, Barcelona, 08005 Spain

²Data Science Centre, Barcelona School of Economics, Barcelona, 08005 Spain

³Thayer School of Engineering, Dartmouth College, Hanover (NH), 03755 United States
alberto.santini@upf.edu, vikrant.s.vaze@dartmouth.edu

16 February 2025

Abstract

Problem definition: More than one-third of US domestic flights are operated by regional airlines. This paper focuses on optimizing medium-term schedule planning decisions for a network of regional airlines through the joint optimization of frequency planning, timetable development, fleet assignment, and some limited aspects of route planning, while capturing passengers' travel decisions through a general attraction discrete choice model. **Methodology:** Unlike previous studies that focused on flight-level decision variables, our formulation uses composite variables to model all non-stop flights between a pair of airports and their complex interdependencies, giving rise to an extended formulation with an extremely tight continuous relaxation. We develop an original solution approach based on column generation and a restricted master heuristic to generate high-quality solutions. Additionally, we propose a new acceleration technique based on dynamic programming to quickly generate promising columns. When combined with implicit dual smoothing, symmetry-breaking techniques, and subproblem aging, this acceleration approach allows us to solve large-scale real-world instances to near-optimality in less than 3 hours. **Results:** Through an extensive computational analysis for some of the largest regional airline networks in the US, we demonstrate the effectiveness of our overall modeling and computational framework. Our approach can generate daily profit improvements of up to \$0.4 million compared to the actual schedule of the airline. **Implications:** We identify the main operational drivers of profit improvement. Furthermore, numerous sensitivity analyses confirm that our results are robust to relaxing key modeling assumptions. Ultimately, our detailed experiments show that the proposed approach provides near-optimal solutions to real-world problem instances within practically reasonable runtimes.

Keywords: timetable development; fleet assignment; passenger choice; column generation; regional airlines.

1 Introduction

Flight schedules have a significant effect on the profitability and competitiveness of airlines. The schedule planning process consists of a sequence of steps, including fleet planning, route selection, frequency planning, timetable design, fleet assignment, aircraft routing, crew scheduling, and tail assignment (Petropoulos et al. 2024). Fleet planning involves decisions related to buying, selling, and leasing aircraft, a multi-year process that significantly affects the airlines' future schedules. Despite its importance, it is very difficult to optimize fleet planning, given the difficulty in generating accurate long-term demand forecasts and the long lead times between aircraft orders and deliveries. Very few quantitative studies focus on optimizing airline fleet planning. In comparison, medium-term strategic decisions related to route selection, frequency planning, timetable design, and fleet assignment benefit from access to more accurate information on demand projections and market competition. Finally, aircraft routing, crew scheduling, and tail assignment decisions are made on a shorter time horizon. This paper focuses on optimizing medium-term strategic decisions related to frequency planning, timetable design, and fleet assignment, as well as limited changes to the route selection process.

Airlines follow various business models, resulting in different network planning, scheduling, and pricing decisions. Examples of such differences include hub-and-spoke versus point-to-point carriers and network legacy versus low-cost carriers. A large number of airlines are classified as regional airlines. They are responsible for operating between 30% and 40% of US domestic flights. Although regional airlines can vary considerably in their strategy and operations, they often share several similarities. They typically offer short-haul flights using smaller aircraft, leverage hub-and-spoke networks, and help passengers connect with flights operated by their mainline partner carriers. Mainline partners are usually long-haul network carriers that operate their flights using larger aircraft. Interestingly, some regional airlines can be among the largest carriers in a country or region when measured in terms of the number of flight operations. For example, in 2022, four of the ten largest US airlines by flight operations were regionals. Regional airlines often sell tickets jointly with mainline partners, and mainline carriers have a considerable say in the scheduling and operating decisions made by their regional partners. In terms of ownership, some regional airlines are wholly owned subsidiaries of the mainline carriers and serve to extend the network of only one mainline partner. Other regional airlines may serve multiple mainline carriers, but when they do so, they almost always serve a unique mainline partner out of each of its hub airports.

This research tackles the joint optimization problem of frequency planning, timetable development, and fleet assignment — as well as some limited aspects of route planning — for a regional airline network. It assumes that the mainline partner’s schedule is fixed. The contributions of this paper are as follows.

- We propose a new integrated optimization formulation for medium-term planning decisions (route selection, frequency planning, timetabling, and fleet assignment) that incorporates a realistic model of passenger choice. Unlike previous medium-term planning studies, which used flight-level decision variables, our formulation uses a single composite variable to model all non-stop flights between a pair of airports and their complex interdependencies. In contrast, the composite variables used in the existing airline scheduling literature have been used to model the itineraries of individual aircraft, crew, and passengers, primarily to solve short-term scheduling problems.
- Composite variables give rise to an extended formulation enjoying an extremely tight continuous relaxation. We exploit this formulation to obtain near-optimal solutions via column generation and a restricted master heuristic. Additionally, we propose a new acceleration approach based on Dynamic Programming (DP) to quickly generate promising columns. When combined with implicit dual smoothing, symmetry-breaking techniques, and subproblem aging, this acceleration approach allows us to solve large-scale real-world instances in up to 3 hours of runtime.
- Through an extensive computational analysis based on ablation studies, we demonstrate the effectiveness of our overall modeling and computational framework. We identify the main operational drivers of the profit improvements enabled by our approach. Furthermore, numerous sensitivity analyses confirm that our results are robust to relaxing key modeling assumptions. Ultimately, our detailed experiments show that the proposed approach provides near-optimal solutions to real-world problem instances within practically reasonable runtimes.

2 Literature Review

Given the focus of this paper on frequency planning, timetable development, and fleet assignment, we now provide an overview of the prominent studies in each of these areas.

2.1 Fleet Assignment and Passenger Flow Modeling

Fleet assignment involves assigning aircraft types to individual flights to minimize assignment costs while satisfying many constraints, including aircraft-flight compatibility, flight cover, flow balance, and aircraft availability. Fleet assignment is among the oldest and most studied problems in the literature on airline operations research. Early studies such as (Abara 1989) and (Hane et al. 1995) assumed

passenger demand to be leg-based, exogenously known, and deterministic. Subsequent studies began to relax these assumptions, progressively leading to increasingly more realistic representations of real-world passenger demand patterns. Barnhart, Kniker, et al. (2002) and Barnhart, Farahat, et al. (2009) relaxed the leg-based demand assumption and explicitly modeled passenger demand at an itinerary level. Barnhart, Kniker, et al. (2002) proposed an itinerary-based fleet assignment model (IFAM) that captures network effects using a passenger mix model (PMM). A PMM simplifies passenger flows by assuming that airlines can control them to maximize revenue, leading to an optimistic fare revenue approximation. Barnhart, Farahat, et al. (2009) developed a subnetwork decomposition approach to solve an itinerary-based fleet assignment model.

Some studies focused on problems that are extensions or complementary to solving the fleet assignment problem. Jacobs et al. (2008) combined a leg-based fleet assignment model with network revenue management considerations and solved the resulting formulation using Benders' decomposition. Ahuja et al. (2007) tackled fleet assignment with a particular emphasis on capturing the revenue premium for the *through connections* and solved it using a very large-scale neighborhood search algorithm. Soumis and Nagurney (1993) and Dumas and Soumis (2008) focused on modeling the flow of passengers in a network for a given schedule and fleet assignment while accounting for stochastic demand and passenger spill-recapture behaviors. Sherali, Bish, et al. (2005) developed and analyzed a model to adjust fleet assignment in response to improved demand forecasts available closer to the day of departure. Sherali and Zhu (2008) developed a two-stage stochastic programming model to allow itinerary-based fleet assignment to benefit from improved passenger demand forecasts by deciding fleet family assignments in the first stage and delaying intra-fleet adjustment decisions to the second stage.

Wang et al. 2014 introduced choice-based demand for more accurate capture of demand substitution effects in a fleet assignment model, replacing more traditional spill-recapture modeling methods. They also suggested potential model extensions to incorporate optional flights, optional markets, departure time selection, and block time selection.

2.2 Incremental Timetable Development

After the early successes in developing and solving fleet assignment models, some researchers started incorporating incremental flight timetabling decisions into the fleet assignment models. These Incremental Timetabling and Fleet Assignment Models (IT-FAM) focused on up to two sets of incremental timetable development ideas: retiming existing flights within departure time windows and adding/removing optional flights to/from a schedule.

Desaulniers et al. (1997) combined aircraft routing decisions into an IT-FAM based on departure time windows and solved the resulting models via an exact solution method combining column generation (to generate aircraft routes) and Dantzig-Wolfe decomposition with branch-and-bound. Lohatepanont and Barnhart (2004) combined a passenger mix model into an IT-FAM based on optional flights and solved it using a column generation (to generate passenger itineraries) and row generation method combined with branch-and-bound. S. Yan et al. (2008) developed an IT-FAM using departure time windows while accounting for stochastic passenger demand and solved it using arc-based and route-based heuristic methods. Sherali, Bae, et al. (2010) also developed an IT-FAM based on optional flights and solved it using Benders' decomposition. Sohoni et al. (2011) developed a stochastic integer program for flight retiming to improve the robustness of an existing schedule and for a given fleet assignment and solved it via an exact cut generation scheme. Jiang and Barnhart (2013) also developed a robust scheduling model that utilizes IT-FAM based on departure time windows and solved it using a column generation approach to generate passenger itineraries. Sherali, Bae, et al. (2013a) and Sherali, Bae, et al. (2013b) developed an IT-FAM with both optional flights and departure time windows and solved it using Benders' decomposition.

More recent IT-FAMs have started incorporating choice-based passenger demand into the incremental timetabling process. This often requires resorting to heuristics rather than exact methods. For example, Abdelghany et al. (2017) designed a bilevel optimization model using an IT-FAM based on departure time windows while incorporating passenger itinerary choice models and developed a

metaheuristic to solve this model. Xu, Wandelt, et al. (2021) also developed an IT-FAM for robust scheduling with optional flights and departure time windows and solved it using a variable neighborhood search heuristic. Finally, C. Yan et al. (2022) recently proposed an IT-FAM based on optional flights under choice-based passenger demand. They extended the subnetworks-based decomposition method by Barnhart, Farahat, et al. (2009) to become the first study to develop an exact solution approach for an IT-FAM under choice-based demand.

2.3 Frequency Planning and Comprehensive Timetabling

Several studies in the last few years have proposed optimization models and solution approaches that make more substantial changes to flight schedules than those made in the incremental timetabling and fleet assignment studies presented in Section 2.2. These more recent studies develop comprehensive timetabling and fleet assignment models (CT-FAMs) as well as frequency planning models. These more substantial scheduling decisions often require researchers to consider the implications of choice-based demand and the effects of the services offered by competing transportation providers.

Vaze and Barnhart (2012) proposed a frequency planning model under airport slot constraints and solved it with an exact approach using dynamic programming. Pita et al. (2013) and Cadarso et al. (2017) both approximated choice-based demand using piecewise linear functions for each period of the day (for example, morning, afternoon, or evening) for a slot-constrained version of CT-FAM. The resulting problem sizes were small enough to be solved directly using an off-the-shelf mixed-integer optimization solver to near-optimality. Wei et al. (2020) embedded a discrete choice passenger demand model at the itinerary level, through sales-based linear programming, into a full-scale CT-FAM applied to a major US carrier. They developed a multiphase heuristic to obtain good solutions within two hours of runtime, substantially outperforming off-the-shelf solvers that produced inferior solutions even after 48 hours of runtime. Kiarashrad et al. (2021) combined CT-FAM with ticket pricing decisions and solved the resulting nonlinear optimization model using a tabu-search metaheuristic. Xu, Adler, et al. (2023) also combined CT-FAM with ticket pricing and solved it using a large neighborhood search heuristic, which was shown to outperform exact methods based on branch-and-bound and branch-and-price. Santini 2025 proposed a joint optimization model for route planning, frequency planning, timetabling, and fleet assignment that focuses on regional airlines' operations at slot-constrained airports, but ignores passengers' travel choices. They specifically underscored the need to take passenger travel decisions into account, using a discrete choice model, to plan medium-term regional airline schedules. In this paper, we address this need.

In summary, all CT-FAM studies either simplify the choice-based demand models or rely on heuristic solution approaches. In contrast, we develop an exact method to solve CT-FAM while incorporating frequency planning and some limited aspects of route selection. Existing scheduling studies that use column generation (CG) do so to generate either aircraft routes or passenger itineraries (e.g., Desaulniers et al. 1997; Kenan et al. 2018; Lohatepanont and Barnhart 2004; Xu, Adler, et al. 2023). In contrast, we use CG to generate segment schedules, which is a new application of CG to airline scheduling. Finally, we accelerate the solution approach by developing and solving a new passenger mix model (PMM) using dynamic programming to speed up the pricing subproblem.

3 Problem Statement and Assumptions

In this section, we describe the research problem and specify our key assumptions.

Regional Airline Perspective: We take the perspective of a regional airline (or a group of regional airlines) supporting a single mainline carrier. For some regional airlines, different subnetworks may support different mainlines. For example, Republic Airlines (YX) has contracts with American Airlines (AA), Delta Airways (DL), and United Airlines (UA). However, the schedules of the regional flights that support each mainline are developed independently. On the other hand, a mainline carrier can coordinate the scheduling of multiple regionals supporting its own network. Therefore, this paper focuses on scheduling regional airline flights (operated by one or more regionals) that support a single

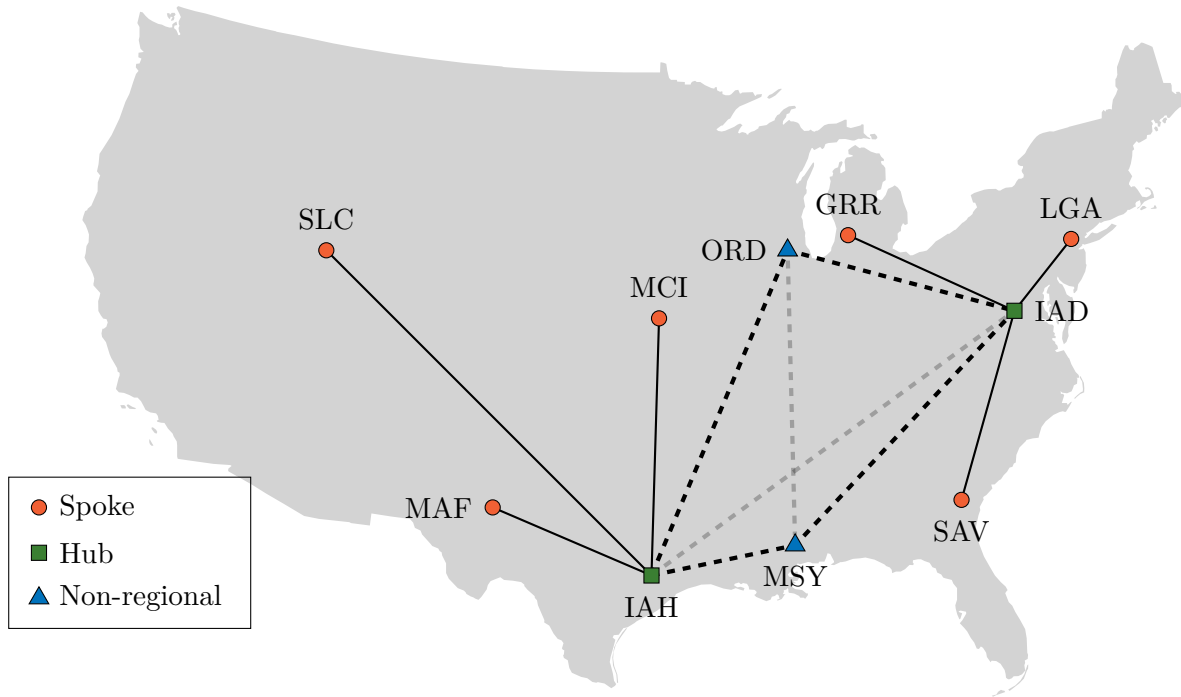


Figure 1: A part of the network for regional airline Mesa (solid) and mainline carrier United Airlines (dashed).

Segments →	Hub-to-hub			Hub-to-NR		Hub-to-spoke			NR-to-hub		NR-to-spoke		Spoke-to-hub			Spoke-to-NR		Spoke-to-spoke			MktSz
	1	2	>2	2	>2	1	2	>2	2	>2	2	>2	1	2	>2	2	>2	1	2	>2	
9E DL	1.2	0.3	0.1	0.4	0.4	19.7	1.4	0.1	0.3	0.4	20.4	1.1	19.5	1.5	0.1	20.4	1.2	7.1	4.2	0.1	2.8M
C5 UA	0.0	0.2	0.1	0.2	0.4	17.3	1.7	0.1	0.1	0.3	27.4	1.9	17.4	1.8	0.1	27.2	2.1	0.0	1.7	0.1	642.7k
G7 UA	0.1	0.6	0.0	0.6	1.4	23.2	0.4	0.0	0.4	1.0	22.9	1.1	23.4	0.4	0.1	21.8	1.4	0.0	1.4	0.0	389.7k
MQ AA	0.0	0.2	0.0	0.3	0.3	14.9	2.0	0.1	0.3	0.3	24.2	1.4	14.6	2.0	0.1	24.3	1.6	8.5	4.7	0.1	2.8M
OH AA	0.1	0.3	0.0	0.4	0.4	17.5	1.5	0.2	0.3	0.4	24.6	1.2	17.3	1.5	0.2	24.7	1.4	0.2	7.9	0.0	2.9M
OO AA	0.0	0.1	0.0	0.3	0.2	12.2	1.9	0.1	0.3	0.2	25.6	1.9	12.1	2.1	0.1	26.6	2.1	10.9	3.2	0.1	1.7M
OO AS	6.0	0.5	0.0	0.7	0.1	22.2	0.7	0.0	0.5	0.1	3.9	0.3	21.9	1.3	0.0	4.8	0.7	35.4	0.9	0.0	971.7k
OO DL	0.2	0.2	0.1	0.5	0.4	19.9	3.6	0.1	0.1	0.4	19.5	2.1	20.0	3.6	0.2	19.8	2.4	0.1	6.1	0.5	3.2M
OO UA	0.1	0.4	0.1	0.2	0.2	20.9	3.4	0.1	0.1	0.2	20.4	1.0	20.4	3.6	0.2	21.1	1.5	0.0	5.9	0.3	3M
PT AA	0.0	0.2	0.0	0.2	0.5	8.5	0.7	0.1	0.1	0.5	33.9	2.6	8.6	0.7	0.1	33.4	2.7	5.0	2.5	0.0	815.2k
YV AA	0.0	0.1	0.0	0.3	0.3	16.0	0.6	0.1	0.2	0.3	30.1	1.4	15.9	0.6	0.1	29.6	1.7	0.0	2.7	0.0	903.3k
YV UA	0.1	0.3	0.0	0.6	0.4	28.1	0.3	0.0	0.3	0.2	18.3	0.7	28.0	0.4	0.0	18.0	1.0	0.5	2.7	0.0	752.4k
YX AA	0.8	0.2	0.0	1.3	0.4	26.1	1.2	0.0	1.1	0.4	11.7	1.1	26.0	1.4	0.1	12.0	1.0	11.2	4.0	0.0	1.9M
YX DL	6.3	0.0	0.0	2.3	0.5	28.9	0.6	0.0	2.5	0.4	8.0	0.6	28.8	0.8	0.0	7.9	0.6	10.6	1.2	0.0	1.2M
YX UA	0.0	0.2	0.0	0.5	0.5	30.5	0.2	0.0	0.2	0.3	14.5	0.7	29.5	0.2	0.0	13.9	0.9	5.2	2.7	0.0	1.1M
ZW UA	0.0	0.0	0.0	0.2	0.2	12.8	0.5	0.0	0.1	0.2	29.8	1.8	12.5	0.5	0.0	29.7	2.2	6.1	3.2	0.0	574.1k

Table 1: Market sizes for the main regional airlines operating in the United States.

mainline. For example, both YX and Mesa Airlines (YV) have flights that support UA. So, UA is interested in coordinating the schedules of YX and YV.

Flow of Flights and Passengers: Each passenger intends to travel from an origin airport to a destination airport. Every ordered pair of origin and destination airports defines a market. An itinerary is a sequence of flights that a passenger takes during a trip. A path is the sequence of airports visited by a passenger during a trip. A segment is any ordered pair of airports between which non-stop flights could be operated.

Airports: The regional airline operates at two types of airports: hubs and spokes. Hubs are airports where the regional airline has a base (for aircraft and crew) and sees a high traffic volume. The hubs also allow passengers to transfer between regional and mainline flights. The spokes are all other airports where a regional airline flies. They are typically smaller and have less traffic. The mainline carrier operates at the hubs of the regional airline but usually not at its spokes. Moreover, mainline carriers also operate at other airports where the regional does not operate; we call these “non-regional airports”. Figure 1 shows an example portion of a network for the regional airline Mesa Airlines and its mainline partner carrier, United Airlines. Circles represent spokes, squares are hubs, and triangles are non-regional (NR) airports. Solid lines represent segments operated by the regional airline, while dashed lines correspond to segments operated by the mainline. Of these, the black dashed lines connect hubs and NR airports, while the lightly colored dashed lines correspond to hub-to-hub or NR-to-NR segments.

Markets: We consider eight market types that are relevant for regional airlines: hub-to-hub, spoke-to-spoke, hub-to-NR (and vice versa), hub-to-spoke (and vice versa), and spoke-to-NR (and vice versa). We do not consider NR-to-NR markets because these are very rarely served by regional flights. Table 1 shows the relative sizes of each type of market for the major regional airlines in the US, using data from the Bureau of Transportation Statistics (2023c). The first column is the regional airline’s IATA code, while the second column is the code of the corresponding mainline. The last column, MktSz, reports the total number of passengers carried in the reference period (quarter 2 of 2022). All other columns report the percentage of passengers belonging to a given market type for each regional airline. For each market, the data are further decomposed by the number of segments that make up the passenger itinerary. For example, in the network depicted in Figure 1, a non-stop passenger from MAF to IAH would count toward the spoke-to-hub markets and one segment. A passenger flying from MAF to ORD with a connection at IAH would count toward the spoke-to-NR markets and two segments. The highlighted rows indicate the regional-mainline combinations that we use in our computational experiments (see Sections 6 and 7).

Paths: We assume that for each market, all passengers who use at least one regional flight follow the same path. For example, the MAF-ORD market can be served via the MAF-IAH-ORD and MAF-IAH-MSY-IAD-ORD paths, but the former is much more popular. In fact, typically, only one path for each market captures almost all of the demand served by the regional airline. For example, when considering mainline UA and regionals YV and YX, 91.29% of all passengers take the most popular path; when aggregating on a market-by-market basis, the median percentage of passengers taking the most popular path is 100%. In our model, we assume that all passengers use the most popular path for each market.

Passenger Connection Times: We assume that all passengers in the same market follow the same itinerary as long as they meet the following two conditions: a) their itinerary includes one regional flight and one mainline flight, and b) they take the same regional flight. For example, passengers on the MAF-IAH-ORD path who take a regional MAF-IAH flight with Mesa Airlines that departs at 8 am and arrives at 10 am will all continue their journey on the same mainline flight. If there are mainline flights departing at 10:50 am, 11:50 am, and 12:50 pm on segment IAH-ORD, then all MAF-IAH-ORD passengers arriving on the Mesa flight to IAH at 10 am will take the 10:50 am departure to ORD, assuming that the 50-minute connection time is more desirable for passengers than the 110- or 170-minute connection times.

Number of Connections: We assume that all itineraries have at most two segments. Indeed, Table 1 shows that only a tiny fraction of passengers use itineraries with more than two segments. When considering the mainline UA and the regionals YV and YX, only 2.36% passengers use itineraries of length > 2 . The two-segment itineraries can be of two types: both segments served by regionals or one segment served by a regional airline and the other by the mainline. In our model, we directly include all markets served by two-segment itineraries that involve one regional segment and one mainline segment. On the other hand, we split every market served by a two-segment itinerary involving two regional segments into two separate markets whose demand is the same as that of the original market. Each of these two markets is modeled to be served by a single-segment itinerary. For mainline UA and regionals YV and YX, only 1.33% of the passengers use an itinerary with more than one segment operated by regional airlines. In Section 7.4, we show that this simplification, which is helpful for modeling purposes, only minimally affects the computational results.

4 Models

In this section, we develop an optimization algorithm for the combined problem of Frequency planning, Timetable development, and Fleet Assignment (FTFA) for a network of one or more regional airlines that support a single mainline. Given the interdependence of the decisions made on different segments of the network, the overall optimization problem must be modeled and solved jointly rather than solving a separate optimization problem for each segment. However, to motivate our mathematical and algorithmic developments, in Section 4.1, we start by formulating the FTFA on a single segment served by a regional airline. In Section 4.2, we formulate the network-wide FTFA. Then, in Section 5, we present a column generation algorithm to solve the network-wide FTFA using the single-segment FTFA.

4.1 Single-Segment Scheduling

Consider a pair of airports, i and j , connected by regional flights, i.e., $(i, j) \in L$ in the notation of Table 2. We start by tackling the problem of devising an optimal schedule between these two airports, given the available fleet. The output of this process will be a segment schedule S , i.e., a sequence of pairs (t, a) in which t is the departure time of a flight from i to j using an aircraft of type a . This segment schedule can be used to determine the cost of operating flights between i and j . However, to evaluate the revenue earned by selling tickets to passengers that use segment (i, j) as part of their itinerary, we need a model that calculates passenger flow given the segment schedule and the fares of available tickets. Therefore, a scheduling model aiming to maximize profits, i.e., revenue minus costs, must integrate a passenger allocation model.

Determining how many passengers will take each offered flight is challenging because this number depends on passenger preferences, the entire offering of flights between i and j , and on other segments that are parts of itineraries containing segment (i, j) , the available capacity on each flight, and the available outside options. Following Wei et al. (2020), we use a General Attraction Model (GAM) to incorporate endogenous passenger choices and determine the resulting passenger flows. In the absence of a tight seating capacity constraint, under the GAM, the number of passengers of market $q \in Q_{ij}$ taking a flight from i to j with a flight departure time s is:

$$M_q \cdot \frac{\tilde{\alpha}_{ijqt}}{\beta_q + \sum_{t \in T_{ij}^*} \alpha_{ijqt}}, \quad (1)$$

where T_{ij}^* is the set of offered departure times on segment (i, j) and M_q is the size of market q . $\tilde{\alpha}_{ijqt}$, ω_{ijqt} and α_{ijqt} are, respectively, the attractiveness, the shadow attractiveness, and the adjusted attractiveness of the itinerary containing a flight in segment (i, j) with flight departure time in period t to passengers in the market q . Note that $\alpha_{ijqt} = \tilde{\alpha}_{ijqt} - \omega_{ijqt}$. $\tilde{\beta}_q$ is the attractiveness of the outside option. The outside option represents the aggregation of all alternatives available to a passenger that do not include flying with the regional airline under consideration. This may include traveling with flights from other airlines, traveling using other modes of transportation, or not traveling. Finally,

Set	Description
D	Set of airports served by the regional airline.
L	Set of segments served by the regional airline.
T	Set of periods in the planning horizon.
$T_{ij} \subseteq T$	Valid departure periods for flights in segment $(i, j) \in L$.
$T' = T \cup \{ T + 1\}$	Set of time periods in the extended time horizon; it includes the planning horizon and the period after that.
Q	Set of all markets included in the model.
$Q_{ij} \subseteq Q$	Set of markets whose itineraries include segment $(i, j) \in L$.
A	Set of aircraft types.
$A_{ij} \subseteq A$	Aircraft types that can operate segment $(i, j) \in L$.
$T_{ij}^a \subseteq T_{ij}$	Valid departure periods for flights in $(i, j) \in L$ operated by an aircraft of type $a \in A_{ij}$.
$D^a \subseteq D$	Set of airports where an aircraft of type $a \in A$ can operate.
$L^a \subseteq L$	Set of segments that can be operated by an aircraft of type $a \in A$.
$A_{ijt} \subseteq A_{ij}$	Aircraft types that can operate a flight in segment $(i, j) \in L$ departing in period $t \in T_{ij}$.
$A_i \subseteq A$	Set of aircraft types that can operate at airport $i \in D$.
\mathcal{S}	Set of all feasible schedules.
$\mathcal{S}_{ij} \subseteq \mathcal{S}$	Set of all feasible schedules for segment $(i, j) \in L$.
$\mathcal{S}_{ijt}^a \subseteq \mathcal{S}_{ij}$	Schedules for segment $(i, j) \in L$, containing a flight departing in period $t \in T_{ij}$ and operated by an aircraft of type $a \in A_{ijt}$.
$\hat{\mathcal{S}}_{ijt}^a \subseteq \mathcal{S}_{ijt}^a$	Schedules for segment $(i, j) \in L$, containing a flight operated by an aircraft of type $a \in A_{ijt}$ whose sum of arrival time and turn-around time corresponds to period $t \in T_{ij}^a$.
Parameter	Description
$M_q \in \mathbb{Z}^+$	Size of market $q \in Q$.
$\check{\alpha}_{ijqt} \geq 0$	Attractiveness of an itinerary containing a flight in segment $(i, j) \in L$ departing at time $t \in T_{ij}$ for passengers of market $q \in Q_{ij}$.
$\omega_{ijqt} \geq 0$	Shadow attractiveness of an itinerary containing a flight in segment $(i, j) \in L$ departing at time $t \in T_{ij}$ for passengers of market $q \in Q_{ij}$.
$\check{\beta}_q \geq 0$	Attractiveness of the outside option for passengers of market $q \in Q$.
$\alpha_{ijqt} \geq 0$	Adjusted attractiveness of an itinerary containing a flight in segment $(i, j) \in L$ departing at time $t \in T_{ij}$ for passengers of market $q \in Q_{ij}$.
$\beta_q \geq 0$	Adjusted attractiveness of the outside option for passengers of market $q \in Q$.
$f_{ijqt} \geq 0$	Fare for a passenger in market $q \in Q_{ij}$ flying on segment $(i, j) \in L$ at time $t \in T_{ij}$.
$c_{ijt}^a \geq 0$	Cost of operating a flight in segment $(i, j) \in L$ departing at time $t \in T_{ij}$ using an aircraft of type $a \in A_{ijt}$.
$t_{ij}^a \in \mathbb{Z}^+$	Number of periods needed for an aircraft of type $a \in A_{ij}$ to be ready for the next flight, after it flies on segment $(i, j) \in L$; it includes the flight time from i to j and the turn-around time at j .
$N^a \in \mathbb{Z}^+$	Fleet size for aircraft type $a \in A$.
$P^a \geq 0$	Seating capacity of aircraft type $a \in A$.
$\sigma_{ij} \in \mathbb{Z}^+$	Minimum number of periods between two consecutive departures for segment $(i, j) \in L$.
$\phi_{ij} \in \mathbb{Z}^+$	Minimum frequency for segment $(i, j) \in L$.
$\psi_{ij} \in \mathbb{Z}^+$	Maximum frequency for segment $(i, j) \in L$.
$\delta_{ijt}^a \in \{0, 1\}$	Takes value 1 iff an aircraft of type $a \in A_{ijt}$ can operate a flight at time $t \in T_{ij}^a$ on segment $(i, j) \in L$ when performing incremental retiming.
$C_{ij} \in [0, 1]$	Maximum load factor for segment $(i, j) \in L$.
$p_S \in \mathbb{R}$	Profit earned when operating the flights in schedule $S \in \mathcal{S}$.
Variable	Description
$z_{ijt}^a \in \{0, 1\}$	1 iff a flight departing during period $t \in T_{ij}^a$ is operated by an aircraft of type $a \in A_{ijt}$.
$x_{ijqt} \geq 0$	Number of passengers of market $q \in Q_{ij}$ taking a flight at time $t \in T_{ij}$.
$x_{q0} \geq 0$	Number of passengers of market $q \in Q_{ij}$ taking the outside option.
$y_S \in \{0, 1\}$	1 iff we operate the segment schedule $S \in \mathcal{S}$.
$w_{it}^a \in \mathbb{Z}^+$	Number of aircraft of type $a \in A$ available at airport $i \in D^a$ at the beginning of period $t \in T'$.

Table 2: Sets and parameters used to formulate the single-segment (Section 4.1) and network-wide (Section 4.2) scheduling problems.

$\beta_q = \check{\beta}_q + \sum_{t \in T_{ij}^*} \omega_{ijqt}$ is the adjusted attractiveness of the outside option. The concept of shadow attractiveness was introduced by Gallego et al. (2015) to overcome the Independence of Irrelevant Alternatives property of the multinomial logit model. In the GAM, the attractiveness of the itinerary $\check{\alpha}_{ijqt}$ is the exponential of a linear function of the itinerary attributes, such as the fare, time of day, travel time, connection time, inconvenience, etc. (Barnhart, Fearing, et al. 2014). We refer the reader to Gallego et al. (2015) for more details on the GAM and to Wei et al. (2020) on the adjusted attractiveness of an airline itinerary. When there is enough seating capacity, expression (1) equals the number of passengers carried, and when the seating capacity constraint is tight, expression (1) serves as an upper limit to the number of passengers carried.

A mathematical formulation for the single-segment scheduling problem under a GAM passenger allocation model for segment (i, j) uses variables $z_{ijt}^a \in \{0, 1\}$ ($a \in A_{ij}$, $t \in T_{ij}^a$), $x_{ijqt} \geq 0$ ($q \in Q_{ij}$, $t \in T_{ij}$) and $x_{q0} \geq 0$ ($q \in Q_{ij}$). Variable z_{ijt}^a equals one if and only if the schedule includes a flight operated by an aircraft of type a departing during period t . Variables x_{ijqt} and x_{q0} are the number of passengers in market q taking, respectively, a regional airline flight at time t or the outside option. The formulation reads as follows.

$$\begin{aligned} \max \quad & \sum_{a \in A_{ij}} \sum_{t \in T_{ij}^a} \left(\left(\sum_{q \in Q_{ij}} f_{ijqt} x_{ijqt} \right) - c_{ijt}^a z_{ijt}^a \right) & (2a) \\ \text{s.t.} \quad & \sum_{\substack{t \in T_{ij}^a \\ s - t_{ij}^a - t_{ji}^a < t \leq s}} z_{ijt}^a \leq N^a & \forall a \in A_{ij}, \forall s \in T & (2b) \\ & \beta_q \cdot x_{ijqt} \leq \alpha_{ijqt} \cdot x_{q0} & \forall t \in T_{ij}, \forall q \in Q_{ij} & (2c) \\ & x_{q0} + \sum_{t \in T_{ij}} x_{ijqt} = M_q & \forall q \in Q_{ij} & (2d) \\ & \sum_{q \in Q_{ij}} x_{ijqt} \leq \sum_{a \in A_{ijt}} P^a z_{ijt}^a & \forall t \in T_{ij} & (2e) \\ & z_{ijt}^a \in \{0, 1\} & \forall a \in A_{ij}, \forall t \in T_{ij}^a & (2f) \\ & x_{ijqt} \geq 0 & \forall q \in Q_{ij}, \forall t \in T_{ij} & (2g) \\ & x_{q0} \geq 0 & \forall q \in Q_{ij}. & (2h) \end{aligned}$$

The objective function (2a) maximizes the schedule's profit. Constraints (2b) ensure that the segment schedule does not use more aircraft than available in any time window of size $t_{ij}^a + t_{ji}^a$. The time window's size reflects that the quickest way for an aircraft flying from i to j to be ready to fly again from i is, upon arriving at j , to turn around and fly back to i immediately. Constraints (2c) define the demand for each flight to be proportional to its attractiveness. Constraints (2d) make sure that the sum of the number of passengers taking the various alternatives (including the outside option) equals the total passenger demand in the market. Constraints (2e) ensure that the offered number of seats is enough to carry all passengers taking each flight.

Constraints (2c)–(2e), first introduced by Gallego et al. (2015), allow incorporating the nonlinear GAM through a small family of linear inequalities. This is commonly known as the Sales-Based Linear Program (SBLP) reformulation of the discrete choice models that belong to the GAM family. This approach, first applied to airline scheduling by Wang et al. (2014), allows us to model the single-segment scheduling problem as a compact mixed-integer program (MIP), compared to techniques that used an exponential number of constraints (Méndez-Díaz et al. 2014) or conic integer programming (Sen et al. 2017).

4.1.1 Single-Segment Model Extensions.

In the following, we describe possible modeling extensions that capture further real-life constraints in the single-segment scheduling model. The corresponding new parameters are also reported in Table 2.

- **Padding between consecutive flights.** Marketing and operational considerations often suggest that consecutive flights on the same segment should not be scheduled to depart too close to each other. Our model can take this requirement into account through the following constraint.

$$\sum_{a \in A_{ijt}} z_{ijt}^a + \sum_{a \in A_{ij}} \sum_{\substack{t' \in T_{ij}^a \\ t \leq t' < t + \sigma_{ij}}} z_{ijt'}^a \leq 1 \quad \forall t \in T_{ij}. \quad (3)$$

- **Minimum and maximum number of flights.** A planner can impose hard limits on the minimum and maximum number of flights on a segment, adding the following constraint.

$$\phi_{ij} \leq \sum_{a \in A_{ij}} \sum_{t \in T_{ij}^a} z_{ijt}^a \leq \psi_{ij} \quad (4)$$

If $\phi_{ij} = \psi_{ij}$ and both values are equal to the number of flights in segment (i, j) in the current airline schedule, the planner forbids frequency changes and only makes fleet assignment and retiming decisions.

- **Incremental flight retiming.** The planner can decide to allow only *incremental retiming* of the currently operated schedule. For example, they might allow flight departure times to change by no more than 30 minutes compared to this schedule. In our model, this can be achieved by forcing variables z_{ijt}^a to take the value zero if departure times t fall outside the allowed ranges.

$$z_{ijt}^a \leq \delta_{ijt}^a \quad \forall a \in A_{ij}, \forall t \in T_{ij}^a.$$

- **Maximum load factor.** Flight load factors are influenced by many decisions, especially related to revenue management, that are outside the scope of our model. The decision maker can specify a maximum load factor for each segment by adding the following constraint to ensure that our model is not too optimistic in assigning passengers to flights.

$$\sum_{t \in T_{ij}} \sum_{q \in Q_{ij}} x_{ijqt} \leq C_{ij} \sum_{t \in T_{ij}} \sum_{a \in A_{ijt}} P^a z_{ijt}^a. \quad (5)$$

- **Aircraft-Type-Dependent Attractiveness and Fares.** The attractiveness of an airline itinerary to a passenger can sometimes be influenced by the types of aircraft used to operate its flights. Moreover, an airline might consider charging different fares for the same itinerary as a function of the aircraft type. Such situations can be modeled by defining new fare parameters f_{ijqt}^a to replace f_{ijqt} , new attractiveness parameters α_{ijqt}^a to replace α_{ijqt} , and new passenger allocation decision variables x_{ijqt}^a to replace x_{ijqt} , each with an additional index $a \in A_{ij}$. Moreover, the objective function (2a) becomes:

$$\max \sum_{a \in A_{ij}} \sum_{t \in T_{ij}^a} \left(\left(\sum_{q \in Q_{ij}} f_{ijqt}^a x_{ijqt}^a \right) - c_{ijt}^a z_{ijt}^a \right).$$

Finally, the following constraints replace (2c)–(2e):

$$\begin{aligned} \beta_q x_{ijqt}^a &\leq \alpha_{ijqt}^a x_{q0} && \forall a \in A_{ij}, \forall t \in T_{ij}^a, \forall q \in Q_{ij} \\ x_{q0} + \sum_{a \in A_{ij}} \sum_{t \in T_{ij}^a} x_{ijqt}^a &= M_q && \forall q \in Q_{ij} \\ \sum_{q \in Q_{ij}} x_{ijqt}^a &\leq P^a z_{ijt}^a && \forall a \in A_{ij}, \forall t \in T_{ij}^a. \end{aligned}$$

4.2 Network-Wide Scheduling

The network-wide schedule optimization problem for a regional airline combines individual single-segment problems (such as those defined in Section 4.1) and links them through fleet size and aircraft balance constraints. In this section, we propose an extended MIP formulation for network-wide schedule optimization.

Denote by \mathcal{S}_{ij} the set of all feasible segment schedules for a given segment $(i, j) \in L$, and by \mathcal{S} the set of all feasible segment schedules. We introduce two sets of decision variables using the notation presented in Table 2. First, the variables $y_S \in \{0, 1\}$ take the value one if and only if the regional airline decides to operate the segment schedule $S \in \mathcal{S}$. Second, the variables $w_{it}^a \in \mathbb{Z}^+$ represent the number of aircraft of type $a \in A$ that are available at the airport $i \in D^a$ at the beginning of the period $t \in T'$. The formulation reads as follows.

$$\max \sum_{S \in \mathcal{S}} p_S y_S \tag{6a}$$

$$\text{s.t.} \quad \sum_{S \in \mathcal{S}_{ij}} y_S = 1 \quad \forall (i, j) \in L \tag{6b}$$

$$\sum_{i \in D^a} w_{i1}^a \leq N^a \quad \forall a \in A \tag{6c}$$

$$w_{i,t+1}^a = w_{it}^a - \sum_{j: (i,j) \in L^a} \sum_{S \in \mathcal{S}_{ij}^a} y_S + \sum_{j: (j,i) \in L^a} \sum_{S \in \mathcal{S}_{ji}^a} y_S \quad \forall i \in D, \forall a \in A_i, \forall t \in T \tag{6d}$$

$$w_{i1}^a = w_{i,|T|+1}^a \quad \forall i \in D, \forall a \in A_i \tag{6e}$$

$$y_S \in \{0, 1\} \quad \forall S \in \mathcal{S} \tag{6f}$$

$$w_{it}^a \in \mathbb{Z}^+ \quad \forall a \in A, \forall i \in D^a, \forall t \in T'. \tag{6g}$$

The objective function (6a) maximizes the total profit from the chosen segment schedules. Constraints (6b) ensure that exactly one segment schedule is selected for each segment operated by the regional airline. Constraints (6c) ensure that no more aircraft of each type are used than are available. Constraints (6d) link the variables y and w that impose flow conservation at each airport in each period and for each aircraft type. Constraints (6e) ensure that the aircraft are correctly positioned at the end of the planning horizon to ensure that the same schedule can be repeated. Remark that for a given segment (i, j) , the number of variables $y_S \in \mathcal{S}_{ij}$ is exponential in the number of periods and aircraft types. Furthermore, in the worst case, the number of feasible solutions of (6a)–(6g) is exponential in the number of segments.

The schedule for a segment (i, j) that appears in an optimal network-wide solution is not necessarily optimal for the single-segment problem and vice versa. Therefore, one cannot directly use the single-segment model to build the set \mathcal{S} required to optimally solve the network-level model. Enumerating all feasible schedules would be prohibitively expensive, due to the size of the set. Moreover, most feasible schedules are not attractive from a commercial point of view. These observations suggest that an appropriate solution method for the network-wide problem should use the single-segment model to identify promising schedules. At the same time, information from the network-wide model should be used at the single-segment level to ensure that the resulting segment schedules work well when used jointly. To achieve this goal, in Section 5, we propose a column generation algorithm that uses dual information from the continuous relaxation of (6a)–(6g) to generate promising schedules using a variation of the single-segment model (2a)–(2h).

In summary, formulation (6a)–(6g) is challenging to solve due to the exponentially many variables y in each segment. However, as demonstrated in Section 6, this formulation enjoys an especially tight continuous relaxation. Furthermore, we develop an exact column generation procedure and a novel acceleration technique, presented in Section 5, that efficiently solve the continuous relaxation. Ultimately, a tight continuous relaxation and an efficient column generation procedure together yield high-quality solutions in short runtime budgets (of less than three hours) for planning problems at the scale of real-world airline networks.

4.2.1 Network-Wide Model Extensions.

In Section 4.1.1, we presented extensions of the single-segment scheduling model that take into account additional real-world complexities and planning requirements. We now present extensions to the network-wide model that can be used to capture additional network-level requirements. See Appendix A for a detailed overview of each extension, where we describe the eventual new constraint(s) needed and how they affect the dual structure of formulation (6a)–(6g).

- **Route selection.** Relaxing the equality constraint (6b) into a less-than-or-equal-to (\leq) constraint, one can incorporate route selection decisions. The resulting model can decide not to assign any schedule to some segments, which corresponds to not operating the segments.
- **Airport slot restrictions.** Many airports limit the number of flight operations that an airline can have at that airport during each period to the number of slots assigned to that airline. Our model allows limiting the number of flights departing from or arriving at each airport during some or all periods to account for such slot constraints.
- **Mainline capacity.** Passengers on two-segment itineraries can take the regional flight only if they can get a seat on the mainline flight. Given the list of mainline flights with the respective capacities that are expected to be available to regional passengers, we can impose the requirement that the chosen regional flights must allow for a feasible continuation of the itinerary on the mainline carrier's network.
- **Average aircraft utilization.** Several factors not captured endogenously by our model (e.g., planned maintenance of the aircraft) can affect aircraft utilization. To ensure that the results correspond to real-life conditions, we allow the average utilization of each fleet type to be capped.

5 Column Generation Algorithm

In this section, we develop a column generation algorithm to solve the network-wide schedule optimization problem presented in Section 4.2. The variables y_S are the columns that we generate in our column generation algorithm. In particular, our algorithm uses dual information from the continuous relaxation of (6a)–(6g) to generate promising segment schedules using a variation of the single-segment model (2a)–(2h). Denote the continuous relaxation of formulation (6a)–(6g) as the *Master Problem* [MP]. Denote with $\lambda_{ij}^{(6b)} \in \mathbb{R}$ and $\lambda_{iat}^{(6d)} \in \mathbb{R}$ respectively the dual variables associated with constraints (6b) and (6d) in [MP]. Given a segment (i, j) , the reduced cost of a schedule $S \in \mathcal{S}_{ij}$ is

$$\bar{p}_{ijS} := p_S - \lambda_{ij}^{(6b)} + \sum_{(t,a) \in S} \left(\lambda_{j,a,t+t_{ij}^a}^{(6d)} - \lambda_{iat}^{(6d)} \right). \quad (7)$$

Let \bar{x}_{ijqtS} be the number of passengers of market q taking the flight in segment schedule $S \in \mathcal{S}_{ij}$ departing in period t . Using the notation of the single-segment problem, the profit associated with schedule S is

$$p_S = \sum_{(t,a) \in S} \left(\left(\sum_{q \in Q_{ij}} f_{ijqt} \bar{x}_{ijqtS} \right) - c_{ijt}^a \right). \quad (8)$$

We denote the following quantities with \bar{f}_{ijqt} and \bar{c}_{ijt}^a .

$$\bar{f}_{ijqt} = f_{ijqt}, \quad \bar{c}_{ijt}^a = c_{ijt}^a - \left(\lambda_{j,a,t+t_{ij}^a}^{(6d)} - \lambda_{iat}^{(6d)} \right). \quad (9)$$

Note that f and \bar{f} are identical as defined in (9), but this is not necessarily the case when using the model extensions presented in Section 4.2.1. Appendix A explains how to amend the definitions of \bar{f}_{ijqt} and \bar{c}_{ijt}^a for each considered extension. With this notation, we can write the reduced cost of S in the following way:

$$\bar{p}_{ijS} = \sum_{(t,a) \in S} \left(\left(\sum_{q \in Q_{ij}} \bar{f}_{ijqt} \bar{x}_{ijqtS} \right) - \bar{c}_{ijt}^a \right) - \lambda_{ij}^{(6b)}. \quad (10)$$

Comparing (10) with the objective function (2a) of the single-segment problem, we note that we can obtain the schedule $S \in \mathcal{S}_{ij}$ with the highest reduced cost by solving problem (2a)–(2h) with the following changes: replace f_{ijqt} with \bar{f}_{ijqt} , replace c_{ijt}^a with \bar{c}_{ijt}^a , and add the constant term $-\lambda_{ij}^{(6b)}$ to the objective function. After these changes, the objective function becomes

$$\max \sum_{a \in A_{ij}} \sum_{t \in T_{ij}^a} \left(\left(\sum_{q \in Q_{ij}} \bar{f}_{ijqt} x_{ijqt} \right) - \bar{c}_{ijt}^a z_{ijt}^a \right) - \lambda_{ij}^{(6b)}. \quad (2a')$$

The single-segment problem, after changing the objection function to (2a'), becomes the pricing subproblem for a column generation algorithm based on [MP]. The algorithm is initialized by solving a restricted [MP] with a single dummy column y_{S_0} . This dummy column has coefficient one in the inequalities (6b), does not appear in (6d), and has a large negative objective coefficient. Adding the dummy column ensures that [MP] is always feasible. If $y_{S_0} > 0$ in the optimal solution to [MP], then the problem is infeasible.

A new segment schedule can improve the objective function of [MP] if its reduced cost is strictly positive. Therefore, in each iteration of the algorithm, we solve one pricing subproblem for each segment and add to [MP] all the y variables associated with positive-reduced-cost segment schedules. We note that each pricing subproblem can produce multiple positive-reduced-cost columns; to this end, we instruct the MIP solver used to optimize the subproblem to keep a list of all visited feasible solutions, and we add all schedules with positive reduced cost to [MP]. When all subproblems prove that no positive-reduced-cost segment schedule exists (i.e., all optimal solutions to the single-segment problems have a non-positive objective value), the incumbent solution to the restricted [MP] is optimal for the [MP]. At that point, we re-solve the restricted [MP] but with integrality constraints (6f) and (6g) reinstated. In other words, we solve (6a)–(6g) only using the generated columns. The procedure presented in this section has been called a *reduced master heuristic* or a *price-and-branch* algorithm (Desrosiers and Lübbecke 2011; Joncour et al. 2010).

5.1 Heuristic Column Generation

In principle, there is no need to solve the restricted [MP] to optimality, and any feasible solution with positive reduced cost, when added to the column pool, can improve the objective function of the restricted [MP]. However, special care is needed when the solution to the subproblem is not optimal. For any given segment $(i, j) \in L$ and period $t \in T_{ij}$, for the SBLP reformulation to work correctly, either (2e) is tight, or all constraints (2c) are tight, or both. This condition always holds for optimal solutions to the pricing subproblem. However, in a suboptimal solution, this condition might be violated for some $t \in T_{ij}$. Consider such a suboptimal solution where, for the given $(i, j) \in L$ and $t \in T_{ij}$, the variables take the values $x_{ijqt} = \tilde{x}_{ijqt}$ and $x_{q0} = \tilde{x}_{q0} \forall q \in Q_{ij}$. Then it is possible to set some of the variables x_{ijqt} to a value higher than \tilde{x}_{ijqt} and, correspondingly, set some of the variables x_{q0} to a value lower than \tilde{x}_{q0} while still satisfying constraints (2c)–(2e). The resulting new solution would have a better objective value. Therefore, if we used the suboptimal solution to build a segment schedule $S \in \mathcal{S}$ as a new column in the restricted [MP], its objective coefficient p_S could be wrong because it is calculated using a passenger number that does not follow the GAM.

To avoid such erroneous objective coefficients in the restricted [MP], we must ensure that for any segment schedule defined by variables z_{ijt}^a , the remaining variables (namely, x_{q0} and especially x_{ijqt}) maximize the objective function of the pricing subproblem across all solutions with that segment schedule. To compute p_S with GAM-consistent passenger allocation, one can fix the variables z_{ijt}^a and re-solve (2a'),(2b)–(2h) to optimality (we call this procedure “ z -fixing”). This transition from a feasible solution that does not follow the GAM to one that does, by definition, cannot decrease the reduced cost.

In terms of the overall implementation of column generation, this issue is linked to the classical question of whether to add single or multiple columns at each pricing iteration. In most cases, adding multiple columns improves the convergence of the column generation algorithm and the quality of the

reduced master heuristic solution. However, adding more than one column does not usually involve any additional computation. However, in our case, each potentially suboptimal column requires running the z -fixing procedure.

This led to three different pricing strategies that we experimented with. As the first strategy, we solve the pricing subproblem, and we instruct the MIP solver to stop as soon as it finds the first positive-reduced-cost solution. We add that single column to the restricted [MP], potentially after running the z -fixing procedure in case the pricing solution is suboptimal. The second strategy involves solving the pricing subproblem to optimality. In this case, the corresponding solution is guaranteed to follow the GAM and no z -fixing is required. As the third strategy, we solve the pricing subproblem to optimality and inspect all the integer solutions encountered by the MIP solver. We run a z -fixing procedure for each suboptimal integer solution and add all positive-reduced-cost integer solutions to the column pool. We conducted a preliminary experiment to determine the most effective strategy and the results showed that the third strategy outperforms the others. This is because the z -fixing procedure is quick and does not significantly impact the algorithm’s overall time while still enriching the column pool. Therefore, all the results in this paper follow the third strategy.

5.2 Approximating an Upper Bound on the Reduced Cost

We describe how to derive an upper bound on the reduced cost of a column by solving a relaxation of the pricing subproblem using Dynamic Programming (DP). This procedure can result in two cases. On the one hand, if the upper bound is negative, then no positive-reduced-cost columns exist. On the other hand, while producing the bound, the procedure can build a segment schedule that can be added to the column pool of the reduced [MP]. Computational experiments show that the DP usually produces high-quality columns and helps to improve the overall convergence of the column generation algorithm.

To apply DP, we first aggregate all the markets whose itineraries include segment $(i, j) \in L$ and, therefore, the aggregate market size is given by $M_{ij} = \sum_{q \in Q_{ij}} M_q$. To overestimate the reduced cost, we assume that in any given period $t \in T_{ij}$, all passengers have the highest fare across all markets $q \in Q_{ij}$, given by $\bar{f}_{ijt} = \max_{q \in Q_{ij}} \{\bar{f}_{ijqt}\}$. We also relax the assumption of GAM-following market shares. Instead, we assume that the passengers are allocated across different flights in a segment schedule in a way that maximizes its reduced cost. In other words, we assume that the airline can decide how many passengers to carry on each flight in the segment schedule while still satisfying the aircraft capacity constraints for all individual flights and respecting the aggregate market size limit. Finally, we disregard the constraints (2b) and, if applicable, the constraints (4). This is consistent with obtaining an overestimation of the reduced cost of a segment schedule. However, we note that constraints (2b) are unlikely to be active anyway because the fleet size is usually much larger than the number of aircraft needed to operate a profitable single-day schedule on a single segment. We define a “*relaxed* segment schedule” (RSS) as a segment schedule that satisfies all constraints other than (2b) and (4). We define the “*relaxed* reduced cost” (RRC) of an RSS as its reduced cost calculated under these modifications.

Denote by $\theta(t, M)$ the highest RRC across all RSS whose first flight does not depart before $t \in T$ when the aggregate market size is $M \in \mathbb{Z}^+$. Then $\theta(1, M_{ij})$ — the highest RRC across all RSS — constitutes an upper bound on the maximum reduced cost across all segment schedules. We compute $\theta(1, M_{ij})$ using DP.

Recall that T_{ij} is the set of valid flight departure times for flights in segment (i, j) . We denote with $v(t) \in T_{ij}$ the first valid departure period that is not before $t \in T$:

$$v(t) = \min\{s \in T_{ij} \mid s \geq t\}.$$

Then, the Bellman equation of our DP algorithm is as follows:

$$\theta(t, M) = \max \left\{ \theta(v(t+1), M), \max_{a \in A_{ijt}} \max_{\xi \in \{1, \dots, \min\{M, P^a\}\}} \{ \bar{f}_{ijt}\xi - \bar{c}_{ijt}^a + \theta(v(t + \sigma_{ij}), M - \xi) \} \right\}. \quad (11)$$

The top line of (11) refers to the case where no flight is scheduled at time t . The bottom line defines the maximum RRC achievable when operating a flight departing at time t using the best possible aircraft type a and carrying the best possible number of passengers ξ . This RRC is the sum of two components: the first one, from the flight departing at time t , is $\bar{f}_{ijt}\xi - \bar{c}_{ijt}^a$; the second one, $\theta(v(t + \sigma_{ij}), M - \xi)$, is the best possible RRC achieved by flights departing at time $t + \sigma_{ij}$ or later and carrying at most $M - \xi$ passengers. We can set $\sigma_{ij} = 1$ if the constraints (3) are not used.

We note that, in recursion (11), the first argument to function θ always strictly increases. Therefore, to initialize the algorithm, we provide θ values for all t values greater than the last valid departure period:

$$\theta(t, M) = -\lambda_{ij}^{(6b)} \quad \forall t \in T, t > \max\{s \mid s \in T_{ij}\}, \quad \forall M \in \{0, \dots, M_{ij}\}.$$

Value $-\lambda_{ij}^{(6b)}$ appears as a constant in the pricing subproblem's objective function (2a').

The table for the DP defined above has a size proportional to $|T| \cdot M_{ij}$. However, changing the number of passengers carried by one does not significantly impact the total reduced cost of a segment schedule. As a result, a one-passenger resolution is often unnecessary to obtain an accurate approximation of the reduced cost of a segment schedule. We propose to reduce the size of the DP table by dividing the interval $\{0, \dots, M_{ij}\}$ into buckets of size B ; therefore, we only consider values M if they are multiples of B or if $M = P^a$. However, we remark that the RRC computed using this approximation is no longer guaranteed to be an upper bound on the maximum reduced cost of a segment schedule for all $B > 1$.

The output of the DP algorithm is an RSS S and its approximate RRC $\theta(1, M_{ij})$. To determine whether S is feasible and, if so, what its exact reduced cost is, we solve the model (2a'),(2b)–(2h) fixing variables z_t^a to value one if $(t, a) \in S$ or to zero otherwise. When the resulting segment schedule is feasible and has a positive actual reduced cost, we add the corresponding column to the pool. Otherwise, we must solve (2a)–(2h) to determine whether any positive-reduced-cost columns exist.

5.3 Acceleration Strategies

We present three additional strategies to speed up the column generation algorithm.

5.3.1 Dual Stabilization.

We counteract the tendency of column generation algorithms towards dual variable oscillations and tailing-off (see, e.g., Vanderbeck and Wolsey 1996) by using the dual barrier algorithm to solve [MP]. When using this algorithm, the optimal dual solution is more likely to lie midface rather than at a vertex. This property smooths the dual values used in the subproblems (Vanderbeck 2005). To obtain a midface dual optimum when one exists, we disable the crossover procedure (Bixby and Saltzman 1994) of the barrier algorithm, which would otherwise postprocess the midface solution and return a dual vertex. We also tried to use techniques that perform explicit dual smoothing — in particular, we implemented the method of Pessoa et al. (2018) — but we did not get any further computational gains.

5.3.2 Selecting Which Subproblems to Solve (Subproblem Aging).

After a few dozen column generation iterations, most subproblems do not return any positive-reduced-cost solutions for many iterations and, eventually, many of them never will. We define *cold* subproblems as those that do not generate positive-reduced-cost columns in a given iteration. Other subproblems are defined as *hot*. When, in some iteration, more than half of the subproblems are cold, we enter a *shortcut* mode. In shortcut mode, we solve the subproblems in decreasing order of the number of columns they produced during the previous iteration. We stop solving subproblems as soon as one fails to produce positive-reduced-cost columns after at least one such column was produced during the current iteration. At that point, we solve [MP] again and move to the next column generation iteration. The main idea behind this strategy is to attempt to solve hot subproblems only and resort to cold ones at the end of the algorithm to prove that no more positive-reduced-cost columns exist.

Instance	Hubs	Spokes	Segments	Flights	Markets	Aircraft types	Aircraft
YV	2	28	64	90	230	2	38
YX	2	30	70	128	206	2	48
OO	5	124	416	784	2810	3	260
YV + YX	4	44	134	218	454	3	86
YV + OO	6	145	518	874	2819	4	298
YX + OO	6	146	488	932	2721	4	308
YV + YX + OO	7	149	534	1002	3312	5	346

Table 3: Key features of the considered instances.

Instance	Base		No Barrier		No SP Aging		No Noise		No DP	
	Gap %	Time (s)	Gap %	Time (s)	Gap %	Time (s)	Gap %	Time (s)	Gap %	Time (s)
YV	0.11	52.80	0.09	4602.45	0.11	277.76	0.10	1042.75	0.10	703.75
YX	0.20	884.05	0.20	1506.19	0.20	1696.42	0.20	2541.12	0.20	1883.55
OO ¹	0.00	10800.00	0.00	10800.00	0.01	10800.00	—	10800.00	0.00	10800.00
YV + YX	0.17	1115.75	0.14	1742.86	0.19	1666.04	0.14	8481.55	0.13	6051.81
YV + OO ²	0.08	10800.00	—	10800.00	0.11	10800.00	—	10800.00	—	10800.00
YX + OO ³	0.12	10800.00	—	10800.00	—	10800.00	—	10800.00	—	10800.00
YV + YX + OO ⁴	0.09	10800.00	—	10800.00	—	10800.00	—	10800.00	—	10800.00

Table 4: Ablation study results. Column Generation runtimes: ¹5h7m; ²15h39m; ³5h38m; ⁴11h9m.

5.3.3 Breaking Symmetry in the Subproblems.

When solving the pricing subproblem for a segment $(i, j) \in L$, the objective function coefficients of variables x and z depend on passenger fares f_{ijqt} , operating costs c_{ijt}^a , and the dual values. Although the fares and operating costs could vary between periods t , in practice, some of them could be identical. For example, operating a flight at 10 am might cost the same as at 10:15 am. The duals are also likely to take only a few different values, and most of them take value zero because they correspond to non-binding primal constraints. Indeed, in our experiments, we observed that the pricing subproblem is affected by symmetry. For example, operating a flight at 8:00 am might yield the same reduced cost as operating it at 8:15 am, 8:30 am, etc. To break this symmetry, we add a negligible amount of noise to all fares f_{ijqt} ; more precisely, we add a Gaussian noise with mean zero and standard deviation $10^{-3} \cdot f_{ijqt}$. Computational experiments confirmed that instances that only differ in whether or not the noise was added to the fares produce primal and dual bounds that are indistinguishable for practical purposes.

6 Computational Results

In this section, we present the results of our computational experiments to evaluate the impact of the main algorithmic design choices. We consider seven instances obtained using real-world data for the three largest regionals supporting United Airlines: Mesa Airlines (YV), Republic Airways (YX), and SkyWest Airlines (OO). This combination of mainline carrier and regionals gives rise to one of the largest instances that can be built from publicly available data. We build one instance for each possible combination of one, two, and three regional airlines. Table 3 reports the main network characteristics of each instance: the number of hub and spoke airports, the number of segments operated, the number of flights flown, the number of markets served, and the number of individual aircraft and aircraft types used. This data was obtained from the Bureau of Transportation Statistics (2023a), Bureau of Transportation Statistics (2023b), and Bureau of Transportation Statistics (2023c)

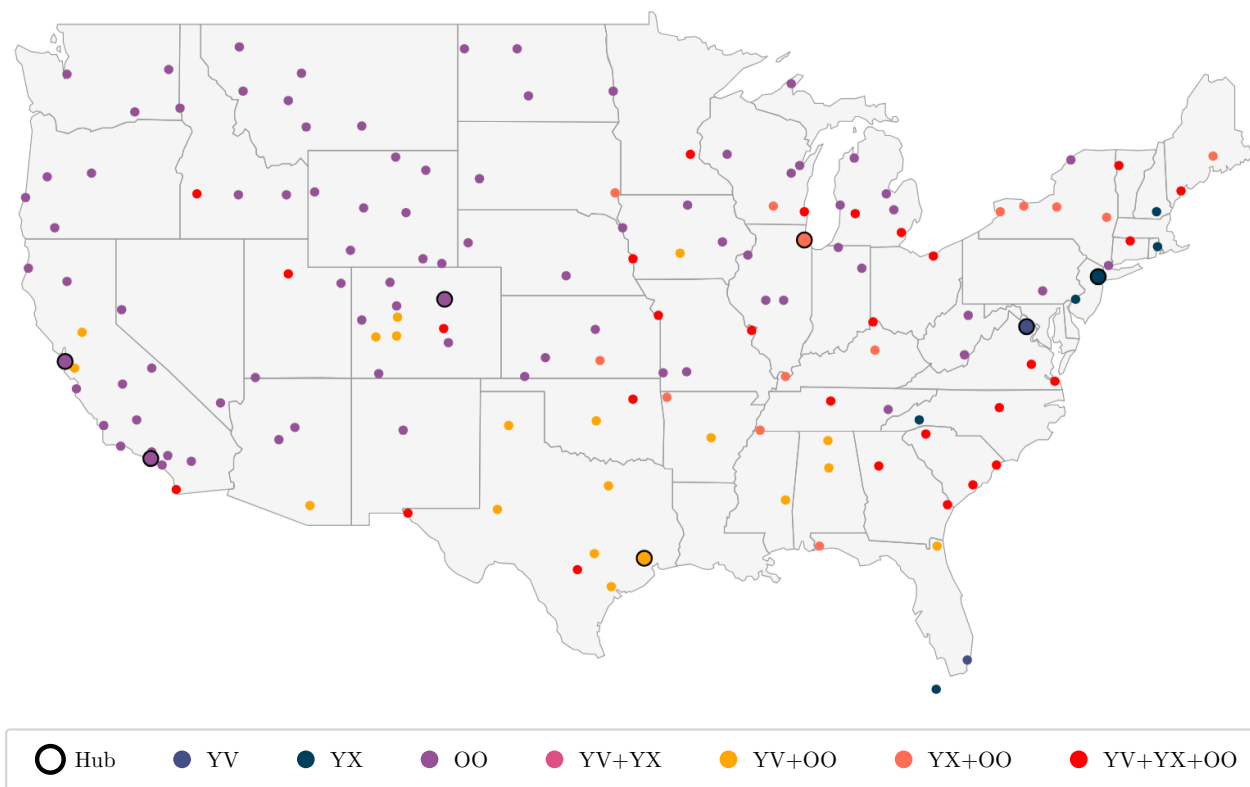


Figure 2: Hubs and Spokes included in the instances.

for June 29, 2022 (Wednesday). The largest joint instance ($YV + YX + OO$) involves operations at 156 airports throughout the continental US, with more than 3,000 markets (that is, more than 3,000 unique origin-destination pairs between which passengers want to travel), and optimizes a fleet of almost 350 aircraft. When creating instances combining multiple regionals, the listed numbers of flights and aircraft in the combined instance are always simple sums of the corresponding values for the individual regionals. However, other characteristics, namely hubs, spokes, segments, markets, and aircraft types, sometimes overlap, resulting in the combined characteristics being different from the sum of the parts. Figure 2 depicts the geographical distribution of hubs and spokes included in the instances. It shows that our instances cover almost all parts of the continental US. Appendix D provides more details on our setup of the case study and data preprocessing steps.

The tested algorithmic design choices are:

- **Barrier.** Using the dual barrier algorithm when solving [MP] to increase the likelihood that the optimal dual solution returned by the solver lies midface, as explained in Section 5.3.1.
- **SP Aging.** Cutting short some column generation iterations when the likelihood of producing new positive-reduced-cost columns later in the same iteration is low, as explained in Section 5.3.2.
- **Noise.** Adding a tiny amount of noise to fares f_{ijqs} to break the symmetry in the subproblem, as explained in Section 5.3.3.
- **DP.** Using the Dynamic Programming algorithm introduced in Section 5.2 and solving MIP (2a)–(2h) only if the DP algorithm does not produce any positive-reduced-cost column.

Table 4 shows the results of an ablation study comparing five configurations. In the “Base” configuration, all algorithmic improvements are active. In the other four configurations, we disable one improvement at a time to evaluate how much the corresponding algorithm’s performance differs from the Base configuration. We ran all experiments on a cluster with Intel Xeon CPUs running at 2.6GHz. The algorithms were coded in C++, and we used Gurobi 9.0.0 as the Linear and Integer Programming solver with a time limit of three hours (unless explicitly stated otherwise). For the DP algorithm, we

used a bucket size of $B = 10$.

In the four instances that involved OO, the column generation algorithm did not terminate within the time limit. In these cases, we ran the final MIP (which solves almost instantaneously) with the columns generated before the three-hour timeout. However, we could not always obtain a feasible MIP solution. The dashes in Table 4 represent cases in which the columns generated in the three-hour limit produced an infeasible MIP.

All gaps are computed relative to the optimal solution of [MP]. When such a solution was not available within the three-hour time limit, we ran the column generation algorithm until [MP] was solved optimally. We remark that, under a hard time budget constraint, an upper bound on the cost of the optimal [MP] solution can be obtained as explained in Appendix B.

The impact of our algorithmic improvements is twofold. In smaller instances (which do not involve OO), they cause a large reduction in computational times, up to $87\times$ compared to “No Barrier”, $5\times$ compared to “No SP Aging”, $20\times$ compared to “No Noise”, and $13\times$ compared to “No DP”. The greatest improvements are for YV, while the improvements are moderate ($1.5\times$ – $7.6\times$) for YX and YV + YX. All configurations show extremely small gaps ($\leq 0.20\%$) that demonstrate the strength of our extended formulation.

In larger instances, the column generation algorithm without the algorithmic improvements is often unable to obtain even a feasible solution within the three-hour time limit. For example, the “Base” configuration is the only one producing a feasible integer solution in the YX + OO and YV + YX + OO instances. The reason is that its more efficient column generation procedure allows us to complete more pricing rounds, thus generating more columns and enabling the MIP solver to build a feasible solution. The “Base” configuration not only obtains feasible solutions to all four hard instances within the three-hour time limit, but the solutions are also not more than 0.12% away from the optimum in all cases. In summary, our algorithmic improvements provide a $1.5\times$ to $87\times$ runtime speedup for smaller instances and produce near-optimal solutions for larger instances for which the algorithm without the improvements is often unable to find even a feasible solution.

7 Practical Insights

In this section, we quantify the impact of our modeling and algorithmic approach on key business metrics of an airline network (Section 7.1). We also identify the main operational drivers of improvement in Section 7.2. In Section 7.3, we demonstrate the value of comprehensive schedule planning compared to more incremental approaches. Sections 7.4 and 7.5 assess the robustness of our numerical results and our main qualitative findings to our two key assumptions: those about regional-regional split markets and minimum connection times. All experiments in this section and in Appendix E are performed using the YV, YX, and YV + YX instances. Furthermore, we indirectly capture the effects of demand uncertainty by imposing a hard upper limit of 93% on the maximum number of seats sold in each segment, setting the appropriate value of C_{ij} in constraints (5). This is a reasonable value because, across the 772 combinations of segments and carriers, for the segments operated by the YV and YX carriers in 2022, with at least one flight per day, only 46 had an average load factor greater than 93%. Still, in Appendix E, we further assess the impact of demand uncertainty and revenue management practices on our key financial results and the continued relevance of our main operational insights.

7.1 Value of Optimization

Table 5 presents the effects of our modeling and algorithmic approach on key financial metrics (that is, total revenue, total cost, and total profit) relative to the real-world schedules of the respective airlines. The first row, named “Status Quo”, shows these financial metrics obtained by running our single-segment scheduling model to optimality for each segment using the actual schedules of Mesa (YV), Republic (YX), or both (YV + YX). Specifically, we fix variables $z_{ijt}^a = \bar{z}_{ijt}^a$, where \bar{z}_{ijt}^a refers to the actual schedules of the respective airlines, solve model (2a)–(2h) to optimality, and then sum the

objective function (2a) across all segments in the network to get the total profit. Moreover, the cross-segment sums of each of the two individual terms in (2a), namely, $\sum_{a \in A_{ij}} \sum_{t \in T_{ij}^a} \sum_{q \in Q_{ij}} f_{ijqt} x_{ijqt}$ and $\sum_{a \in A_{ij}} \sum_{t \in T_{ij}^a} c_{ijt}^a z_{ijt}^a$, give the total revenue and total cost, respectively, as presented in Table 5. Here, we need not run the network-wide model because fixing the schedule allows us to decouple segment-specific models from each other. The remaining rows (Freq $\pm n$) of the table show the results obtained by our column generation algorithm while allowing the segment frequency to be different from that of the actual real-world schedules by at most n on each segment while ensuring that the flight frequency never goes below a threshold m on any segment. Using the notation introduced in Section 4.1.1, we set parameters ϕ and ψ as follows:

$$\phi_{ij} = \left(\sum_{a \in A_{ij}} \sum_{t \in T_{ij}^a} \bar{z}_{ijt}^a \right) + n \quad \text{and} \quad \psi_{ij} = \max \left\{ \left(\sum_{a \in A_{ij}} \sum_{t \in T_{ij}^a} \bar{z}_{ijt}^a \right) - n, m \right\}$$

where $n = 0, 1, 2$ and $m = 0, 1$. We note that the case with $n = 0$ fixes frequencies and thus only allows retiming and fleet assignment decisions while setting $n = 1, 2$ allows frequency planning decisions to be captured in the model. Using $m = 1$ ensures that the frequencies remain positive, while using $m = 0$ allows us to remove some segments from the schedule if the corresponding frequency is zero. Thus, the last two groups of rows of Table 5 correspond to capturing some limited amount of route selection decisions.

In Table 5, all the numbers in parentheses are the differences relative to “Status Quo”. The table provides several insights. First, the costs and revenues computed by our approach are both higher than those under the status quo in all cases. Furthermore, in each case, a profit higher than the status quo is obtained because the absolute revenue growth is greater than the absolute cost increase. As expected, in the “Freq ± 0 ” case, with the segment frequencies not allowed to change, the cost growth is driven only by the reassignment of aircraft types corresponding to different operating costs. Consequently, the corresponding cost changes are small (in the 0.02%-2.83% range). On the other hand, revenue increases by 2.04%-2.49%, primarily driven by prioritizing high-value passengers (see Section 7.2), leading to a profit increase of 3.50%-4.39%. In the “Freq ± 1 ” and “Freq ± 2 ” cases (with $m = 1$), the costs and revenues increase across the board by 6.65%-12.23%, due to the increase in the number of flights operated and passengers carried, respectively. In all cases, the absolute increase in revenue is greater than the cost increase, leading to larger profits by 6.96%-13.32%.

Setting $m = 0$ rather than $m = 1$ allows the model to consider solutions with even lower frequencies on some segments (zero instead of one), which can lead to fewer scheduled flights, lower costs, and lower revenues on those segments. However, this flight reduction also creates opportunities to use the corresponding spare capacity in a potentially more profitable way in other segments. Indeed, the last two groups of rows of Table 5 support this conclusion. Specifically, comparing $m = 0$ with the corresponding $m = 1$ rows (for both $n = 1$ and $n = 2$), we remark that costs remain almost constant (with increases between 0.14 and 1.15 percentage points), hinting at the fact that the aircraft utilisation constraints were already binding for $m = 1$. At the same time, replacing less profitable routes with more profitable ones leads to consistent revenue increase (between 2.96 and 4.93 percentage points) and, ultimately, increased profits (with increases between 4.74 and 8.56 percentage points). In terms of absolute daily profit increase for the YV + YX instance, we remark that only performing fleet assignment and retiming (Freq ± 0) yields an improvement of a little over \$80k; adding a limited frequency optimization component raises this figure to nearly \$250k (Freq ± 2 and $m = 1$); further allowing route selection brings the figure to almost \$400k (Freq ± 2 and $m = 0$).

In the remainder of this section, we focus our attention on the $m = 1$ case.

7.2 Operational Drivers of Improvement

To understand the main reasons for the improvement obtained by our approach, in Table 6, we present key operational metrics associated with the various solutions corresponding to the joint YV + YX instance. Table 6 presents the number of flights, average flight distance in nautical miles (nm)

	m	YV			YX			YV+YX		
		Revenue	Cost	Profit	Revenue	Cost	Profit	Revenue	Cost	Profit
Status Quo		1,664,272	696,246	968,026	1,723,135	839,696	883,439	3,387,407	1,535,941	1,851,465
Freq ± 0	1	1,698,251 (+2.04%)	696,348 (+0.02%)	1,001,903 (+3.50%)	1,762,283 (+2.27%)	842,073 (+2.83%)	920,211 (+4.16%)	3,471,732 (+2.49%)	1,538,942 (+0.20%)	1,932,790 (+4.39%)
Freq ± 1	1	1,805,247 (+8.47%)	742,565 (+6.65%)	1,062,682 (+9.78%)	1,869,700 (+8.51%)	924,730 (+10.13%)	944,970 (+6.96%)	3,713,894 (+9.64%)	1,670,419 (+8.76%)	2,043,474 (+10.37%)
Freq ± 2	1	1,841,328 (+10.64%)	751,723 (+7.97%)	1,089,605 (+12.56%)	1,933,830 (+12.23%)	936,473 (+11.53%)	997,357 (+12.89%)	3,790,470 (+11.90%)	1,692,362 (+10.18%)	2,098,108 (+13.32%)
Freq ± 1	0	1,884,902 (+13.26%)	744,803 (+6.97%)	1,140,099 (+17.78%)	1,953,937 (+13.39%)	933,405 (+11.16%)	1,020,532 (+15.52%)	3,814,298 (+12.60%)	1,683,081 (+9.58%)	2,131,217 (+15.11%)
Freq ± 2	0	1,903,727 (+14.39%)	752,728 (+8.11%)	1,150,999 (+18.90%)	1,994,866 (+15.77%)	938,120 (+11.72%)	1,056,746 (+19.62%)	3,957,599 (+16.83%)	1,709,924 (+11.33%)	2,247,675 (+21.40%)

Table 5: Key financial metrics of the solutions of the Mesa (YV), Republic (YX), and joint (YV + YX) instances.

	Status Quo	Freq ± 0	Freq ± 1	Freq ± 2
Number of flights	218	218	240	246
Avg flight distance (nm)	571.0	571.0	584.7	587.4
RASM	36.5	37.3	35.5	35.2
CASM	16.6	16.5	16.0	15.7
Passengers carried	14908	14920	16211	16666
Non-stop	66%	66%	66%	66%
one-stop	34%	34%	34%	34%
Avg fare	\$227	\$233	\$229	\$227
Non-stop	\$223	\$229	\$227	\$225
one-stop	\$242	\$246	\$238	\$236
Light retiming	—	63%	58%	60%

Table 6: Key operational metrics of the solutions of the joint network instance (YV + YX).



Figure 3: Variation in the number of flights in the optimized YV + YX schedule compared to the “Status Quo”.




Figure 4: Flight distance distribution in the optimized YV + YX schedules.

(sometimes called “Stage Length”), revenue per available seat-mile (RASM), cost per available seat-mile (CASM), number of passengers carried (both non-stop and one-stop), average fare per passenger (both non-stop and one-stop), and the number of flights that underwent only a small amount of retiming (called “Light Retiming”). Note that “Light Retiming” is defined as the percentage of flights in a given schedule that have a matching flight in the status quo schedule on the same segment within ± 30 minutes of its departure time. If an optimized schedule has multiple flights within ± 30 minutes of the departure time of a flight in the status quo schedule, then we can match with that flight in the status quo schedule only the closest flight in the optimized schedule.

When comparing our “Freq ± 0 ” solutions with the status quo solutions, as expected, the number of flights and average flight distance remain unchanged since we can only change the flight times and the aircraft types assigned to each flight. Retiming aligns flights with the peak demand times of higher-profit markets, even if it may lead to some misalignment of flight times with peak demand times in lower-profit markets to meet the aircraft availability constraints. This increases average fares despite carrying a similar number of passengers. Consequently, RASM increases by 2.2%, while CASM remains flat, leading to the 4.4% profit improvement.

When we move from “Freq ± 0 ” to “Freq ± 1 ” and “Freq ± 2 ” solutions, the drivers of the further improvements are quite different from those that provide improvement when going from “Status Quo” to the “Freq ± 0 ” solution. First, “Freq ± 1 ” and “Freq ± 2 ” progressively increase the number of flights relative to “Freq ± 0 ”, by 10% and 13%, respectively. This leads to 9% and 12% growth in the number of passengers but causes a reduction in the average fares because the “Freq ± 0 ” solution already captures most of the high-fare passengers requiring the “Freq ± 1 ” and “Freq ± 2 ” solutions to dig deeper in search of additional revenue. Moreover, Table 6 shows that compared to the “Freq ± 0 ” solution, the “Freq ± 1 ” and “Freq ± 2 ” solutions use more long-distance flights, which also helps reduce CASM, and further increases profit.

The split of the one-stop versus non-stop passengers stays similar for all four schedules. Furthermore, roughly 60% of the flights in the optimized schedules correspond to light retiming. In other words, approximately 40% of the flights in each of the optimized schedules have a departure time substantially different from that of a matching flight in the status quo schedule.

Figure 3 groups all 134 segments in the YV + YX instance based on daily frequency changes (with respect to the “Status Quo”) in the “Freq ± 1 ” and “Freq ± 2 ” solutions. The number in the cell identified by value $x \in \{-2, -1, 0, 1, 2\}$ on the horizontal axis and $y \in \{-1, 0, 1\}$ on the vertical axis gives the number of segments whose frequency changed by x in the “Freq ± 2 ” solution and by y in the “Freq ± 1 ” solution. First, 102 out of the 134 segments undergo no frequency change in either the “Freq ± 1 ” or “Freq ± 2 ” solution, implying that those segments only underwent flight retimings and/or aircraft type adjustments. In contrast, 18 segments see the maximum allowable frequency increase and two segments see the maximum allowable frequency decrease, thus highlighting the segments with

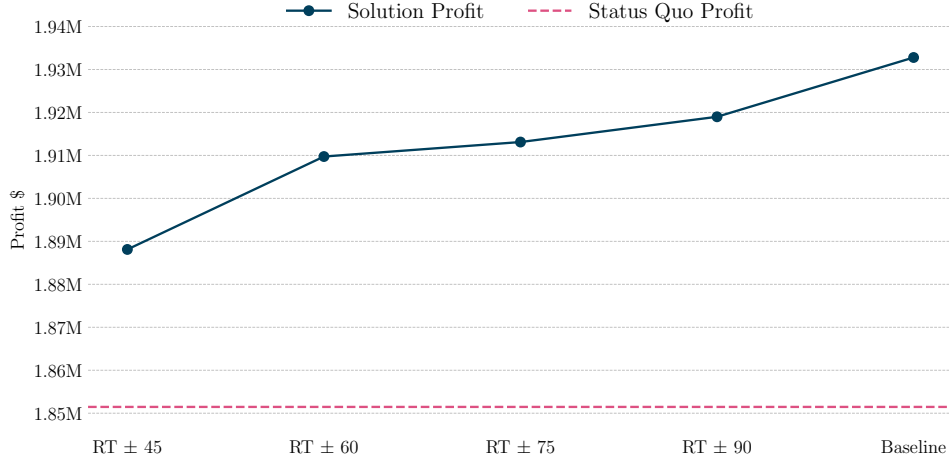


Figure 5: Effect of retiming on profits in the joint instance (YV + YX). “Baseline” case allows unlimited retiming. “RT ± t ” cases allow flights to be retimed by $< t$ minutes compared to the status quo schedule.

the greatest frequency adjustment potential. Four other segments see an increase of one flight per day, and two others see a decrease of one flight per day, regardless of the frequency change limit, indicating that the status quo frequency was off by one per day compared to the optimal schedule. More interestingly, two segments do not undergo a frequency change in the “Freq ±1” schedule, but the “Freq ±2” schedule reduces their frequency by one. This phenomenon indicates that the model sacrifices frequency on these segments to provide enough aircraft time to operate the more profitable segments that see their frequency increase by two. Finally, a more extreme version of this phenomenon is seen in the case of four other segments, which have their frequency increase by one in the “Freq ±1” schedule but undergo a frequency reduction by one in the “Freq ±2” schedule.

Figure 4 shows the total number of flights by flight distance for the YV + YX instance. The left subfigure includes all segments, the middle includes segments with more than one daily flight in the status quo schedule, and the right subfigure includes the segments with more than two daily flights in the status quo schedule. The right subfigure is of particular interest because it shows the trends for only those segments that allow the optimization model to explore the full spectrum of alternatives, ranging from reducing the daily frequency by two to increasing the daily frequency by two. As we move from the “Freq ±0” schedule to “Freq ±1” schedule to “Freq ±2” schedule, we notice that the number of flights increases for most flight distances higher than 400 nautical miles (nm) but slightly decreases for the 300-400 nm flight distance range. This indicates that a major driver of profit improvement is an increase in longer and more profitable flights with some reduction in the number of shorter-distance flights to satisfy the aircraft availability constraints. This is consistent with the fact that, as shown in Table 6, compared to the “Freq ±0” schedule, the average flight distance increases by 2.4% for the “Freq ±1” schedule and by 2.9% for the “Freq ±2” schedule.

7.3 Value of Comprehensive Schedule Planning

Our approach focuses on a comprehensive timetabling problem, which provides the optimization model complete freedom to choose the departure times and aircraft types for each flight. In comparison, several previous studies have focused on incremental timetabling decisions. The incrementality is manifest in one of two ways: some studies only allow retiming flights within a small window around their departure times in an existing schedule, while others only allow choosing flights from a predetermined set of optional flights. Thus, an interesting question relates to the additional benefits accrued by the comprehensive timetabling approach compared to the incremental approaches. Figures 5 and 6 provide one way to visualize these benefits.

Figure 5 plots the profit of the YV + YX instance against the size of the retiming window. Specifically,




Figure 6: Effect of retaining a random subset of possible departure times for each segment in the joint instance (YV + YX). “Baseline” retains all departure times. Each violin plot uses 10 random draws.

from left to right, the window size increases from ± 45 minutes (min) to ± 90 min. The rightmost point shows the profit for our “Freq ± 0 ” schedule, which allows unlimited retiming (denoted as “Baseline” in this figure). As noted before, the “Baseline” profit is 4.39% higher than that of the status quo schedule. When restricting window sizes to ± 45 min, the profit improvement relative to the status quo schedule decreases to 1.98%. Similarly, window sizes of ± 60 , ± 75 , and ± 90 min yield profit increases of 3.15%, 3.33%, and 3.65%, respectively. Thus, incremental retiming approaches can capture between 45% (for the ± 45 min window size) and 83% (for the ± 90 min window size) of the profit improvement provided by our comprehensive timetabling approach. In other words, we stand to lose more than half of the benefits of our timetable optimization if we restrict departure times to the ± 45 min retiming window.

Our comprehensive timetabling approach decides the periods used as departure times for the flights in each segment, along with the corresponding aircraft type assigned to each selected period. A common way of incremental timetabling in the literature (called the “optional flights” approach) involves pre-selecting a set of candidate periods in each segment and letting the optimization model choose from those. In essence, this optional flights approach is a restricted version of the comprehensive timetabling that involves fewer periods to choose from. To quantify the additional benefits of comprehensive timetabling compared to the optional flights approach, we characterize the relationship between profit and the number of candidate periods. To this end, we conducted a number of experiments in which we allowed flights to depart only in a subset of the periods in the planning horizon. This subset contains all the departure periods in the status quo schedule, plus other randomly selected periods, totalling $z\%$ of the available periods, with $z \in \{10, 15, 20, 25, 30, 35, 40\}$. We create ten random subsets for each value of z and demonstrate the results in Figure 6 as a series of violin plots of the profits against the percentage of candidate departure time periods retained.

Note that we recover approximately half of the benefits of comprehensive timetabling by retaining just 20% of the departure times. However, the profit curve in Figure 6 starts to flatten beyond that point, making it harder to recover the remaining profit gains through this incremental approach. Overall, this figure shows that comprehensive timetabling allows additional profit gains that are between $\approx 310\%$ (at $z = 10\%$) and $\approx 33\%$ (at $z = 40\%$) compared to those achieved by an optional flights approach.

7.4 Sensitivity to the Regional-Regional Split Market Assumption

Recall that we split markets that are served with an itinerary composed of two regional flights (Spoke A–Hub–Spoke B) into two separate markets (Spoke A–Hub and Hub–Spoke B). These passengers represent a small fraction of the total market size (see Table 1). For example, as mentioned in

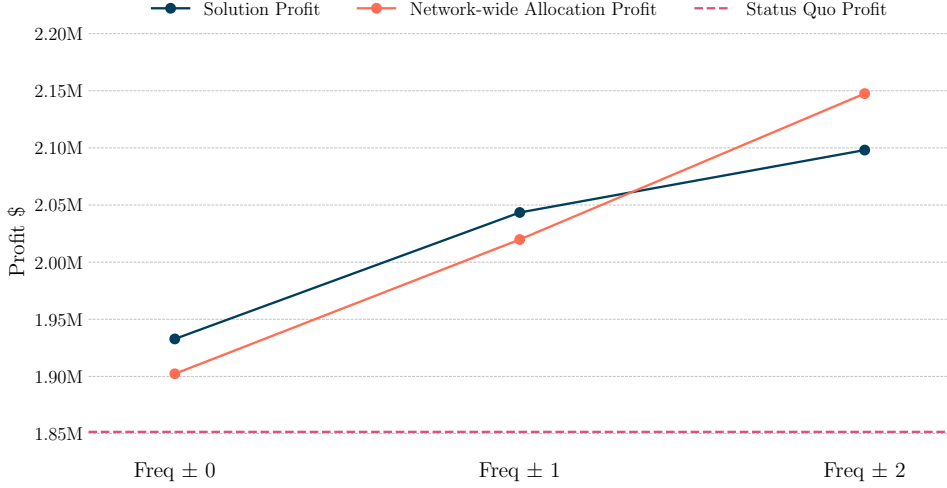


Figure 7: Profit comparison on the joint instance (YV + YX) for the schedules optimized under the split market assumption and evaluated with (in blue) and without (in orange) the assumption.

Section 3, when considering mainline UA and regionals YV and YX, only 1.33% of the passengers use an itinerary with more than one segment operated by regional airlines. This introduces some inaccuracy in our final results. However, the direction and magnitude of this inaccuracy are unclear because this assumption causes two opposite effects.

On the one hand, rescheduling some regional flights might break a previously feasible connection at the hub. Under the split market assumption, we still optimistically allocate passengers affected by this disruption to our Spoke A–Hub and Hub–Spoke B flights. In reality, these passengers might not be able to complete their itinerary using the new schedule. On the other hand, the new schedule might now enable new connections that were not possible in the status quo schedule, creating new markets that were previously not served. The first effect suggests that our results presented in Section 7.1 might overestimate the true profit improvement, whereas the second effect suggests that they might underestimate it.

In this section, we estimate the net impact of these two opposite effects by focusing on the YV + YX instance. We first solve the three scenarios (“Freq±0”, “Freq±1”, and “Freq±2”) using our column generation algorithm while retaining the split market assumption. Then, we fix the resulting schedules and replace the split markets (Spoke A–Hub and Hub–Spoke B) with the original Spoke A–Hub–Spoke B markets. At the same time, we introduce new markets not present in the original instance because they were currently unserved and, therefore, their passengers did not appear in the databases we used. For this reason, we do not know the exact market sizes and fares for these new markets. We estimate the sizes of the new markets using a gravity model (see, e.g., Morley et al. 2014) and estimate their fares by choosing, for each new market, the existing market with the most similar size and assuming the same fare per mile. In this way, we avoid using the fare of a large market for a small one and vice versa. Finally, we run a network-wide passenger allocation model (described in Appendix C) using these market sizes and attractiveness values calculated using the estimated fares for the possible itineraries in the new schedule.

Figure 7 shows the resulting profits with (in blue) and without (in orange) the split market assumption. We find that the net impact of these two opposite effects can be positive or negative. In case of the “Freq±0” and “Freq±1” instances, relaxing the assumption leads to a slight reduction in the profit improvement (between 1.3 and 1.6 percentage points) obtained by the optimized schedules, whereas the “Freq±2” instance shows a higher profit improvement (by 2.7 percentage points) when we relax the assumption. In fact, the average profit for the three instances remains identical (within $\approx 0.08\%$) with and without the split market assumption.




Figure 8: Profit comparison for different percentages of passengers having a minimum connection time of 1 hour instead of the default 30 minutes.

7.5 Impact of Connection Times

Recall our assumption that all passengers in the same market follow the same itinerary as long as their itinerary includes one regional flight and one mainline flight, and they take the same regional flight. For example, all passengers on the MAF-IAH-ORD path who take a regional MAF-IAH flight with Mesa Airlines that departs at 8 am and arrives at 10 am will continue their itinerary on the same mainline flight. If there are mainline flights departing at 10:50 am, 11:50 am, and 12:50 pm on segment IAH-ORD, then all MAF-IAH-ORD passengers arriving on the Mesa flight to IAH at 10 am will take the 10:50 am departure to ORD, assuming that the 50-min connection time is more desirable for passengers than the 110- or 170-min connection times. Specifically, we assume that all passengers will choose the connection with the shortest connection time above a minimum value of 30 min. In this section, we evaluate the impact of relaxing this assumption by allowing passengers on the same path and the same regional flight to choose from multiple mainline flights. Specifically, Figure 8 shows the effect of a varying percentage of passengers choosing a connection with 30- versus 60-min of minimum layover duration. Figure 8 demonstrates that minor violations of our assumption (for example, changing the percentage of passengers with a 30-min minimum layover duration from 100% to 80% or 90%) results in only slight changes in total profit (by 0.12%-0.36%). However, larger violations (for example, changing the percentage from 100% to 50%, 60% or 70%) can dilute the profit improvements generated by our approach. This emphasizes that in order to maximize the benefits of our approach, it is critical to capture passenger preferences with at least a moderate level of accuracy.

8 Discussion

This paper presented a new column-generation-based modeling and algorithmic approach to jointly optimize medium-term schedule planning decisions of a regional airline for a given schedule of the mainline partner carrier. We developed an original mathematical formulation for the integrated optimization of flight timetabling, fleet assignment, frequency planning, and some limited aspects of route planning. Our formulation adopted a composite variable approach with a very tight continuous relaxation, where each variable indicates a particular segment schedule. Moreover, our formulation endogenously captured passengers' booking decisions through a sales-based linear programming reformulation of a general attraction discrete choice model. We accelerated the column generation approach by developing a new acceleration strategy based on dynamic programming to quickly identify promising columns. Combined with implicit dual smoothing, symmetry breaking, and subproblem aging, this acceleration approach allowed us to solve large-scale real-world instances to near-optimality in

less than 3 hours. Extensive computational experiments based on some of the largest real-world regional airline networks that can be constructed using public datasets demonstrated that our approach can yield hundreds of thousands of dollars in daily profit improvements to the airline. In addition, these findings remained stable under a variety of sensitivity and robustness checks.

From a practical point of view, a promising direction for future research includes extending this modeling and algorithmic framework to airlines that have network structures different from regional airlines. There is also an opportunity to integrate additional modeling considerations related to aircraft routing and crew scheduling into the formulation to enable a more seamless integration with these downstream schedule planning steps. From a methodological perspective, our restricted master heuristic provided solutions with optimality gaps of up to 0.2%. Future research studies could explore the possibility of implementing a full branch-and-price strategy to potentially eliminate this small remaining degree of suboptimality.

Acknowledgements

Alberto Santini’s research was funded through grant RYC2022-035269-I financed by MCIN AEI 10.13039/501100011033 (Spain) and FSE+ (European Union).

References

- Abara, Jeph (1989). “Applying integer linear programming to the fleet assignment problem”. In: *Interfaces* 19.4, pp. 20–28. DOI: 10.1287/inte.19.4.20.
- Abdelghany, Ahmed, Khaled Abdelghany, and Farshid Azadian (2017). “Airline flight schedule planning under competition”. In: *Computers & Operations Research* 87, pp. 20–39. DOI: 10.1016/j.cor.2017.05.013.
- Ahuja, Ravindra K, Jon Goodstein, Amit Mukherjee, James B Orlin, and Dushyant Sharma (2007). “A very large-scale neighborhood search algorithm for the combined through-fleet-assignment model”. In: *INFORMS Journal on Computing* 19.3, pp. 416–428. DOI: 10.1287/ijoc.1060.0193.
- Barnhart, Cynthia, Amr Farahat, and Manoj Lohatepanont (2009). “Airline fleet assignment with enhanced revenue modeling”. In: *Operations research* 57.1, pp. 231–244. DOI: 10.1287/opre.1070.0503.
- Barnhart, Cynthia, Douglas Fearing, and Vikrant Vaze (2014). “Modeling passenger travel and delays in the national air transportation system”. In: *Operations Research* 62 (3), pp. 580–601. DOI: 10.1827/opre.2014.1268.
- Barnhart, Cynthia, Timothy S Kniker, and Manoj Lohatepanont (2002). “Itinerary-based airline fleet assignment”. In: *Transportation science* 36.2, pp. 199–217. DOI: 10.1287/trsc.36.2.199.566.
- Belobaba, Peter P. and András Farkas (1999). “Yield Management Impacts on Airline Spill Estimation”. In: *Transportation Science* 33 (2), pp. 217–232. DOI: 10.1287/trsc.33.2.217.
- Bixby, Robert and Matthew Saltzman (1994). “Recovering an optimal LP basis from an interior point solution”. In: *Operations Research Letters* 15 (4), pp. 169–178. DOI: 10.1016/0167-6377(94)90074-4.
- Bureau of Transportation Statistics (2023a). *Air Carrier Financial Reports and Statistics (Form 41, T100)*. Data for Q2 2022. Accessed on 2023-10-01. URL: https://www.transtats.bts.gov/Tables.asp?Q0_VQ=EGI.
- Bureau of Transportation Statistics (2023b). *Airline On-Time Performance (AOTP)*. Data for June 2022. Accessed on 2023-10-01. URL: https://www.transtats.bts.gov/Tables.asp?Q0_VQ=EFD.
- Bureau of Transportation Statistics (2023c). *Airline Origin and Destination Survey (DB1B)*. Data for Q2 2022. Accessed on 2023-10-01. URL: https://www.transtats.bts.gov/Tables.asp?Q0_VQ=EFI.
- Cadarso, Luis, Vikrant Vaze, Cynthia Barnhart, and Ángel Marín (2017). “Integrated airline scheduling: Considering competition effects and the entry of the high speed rail”. In: *Transportation Science* 51.1, pp. 132–154. DOI: 10.1287/trsc.2015.0617.

- Desaulniers, Guy, Jacques Desrosiers, Yvan Dumas, Marius M Solomon, and François Soumis (1997). “Daily aircraft routing and scheduling”. In: *Management Science* 43.6, pp. 841–855. DOI: 10.1287/mnsc.43.6.841.
- Desrosiers, Jacques and Marco Lübbecke (2011). “Branch-Price-and-Cut Algorithms”. In: *Wiley Encyclopedia of Operations Research and Management Science*. Ed. by James Cochran, Louis A. Jr. Cox, Pinar Keskinocak, Jeffrey Kharoufeh, and Cole Smith. John Wiley & Sons, Ltd. ISBN: 9780470400531. DOI: 10.1002/9780470400531.eorms0118.
- Dumas, Jonathan and François Soumis (2008). “Passenger flow model for airline networks”. In: *Transportation Science* 42.2, pp. 197–207. DOI: 10.1287/trsc.1070.0206.
- Federal Aviation Administration (2023). *Aircraft Registration Database*. Accessed on 2023-10-01. URL: https://www.faa.gov/licenses_certificates/aircraft_certification/aircraft_registry.
- Gallego, Guillermo, Richard Ratliff, and Sergey Shebalov (2015). “A general attraction model and sales-based linear program for network revenue management under customer choice”. In: *Operations Research* 63.1, pp. 212–232. DOI: 10.1287/opre.2014.1328.
- Hane, Christopher A, Cynthia Barnhart, Ellis L Johnson, Roy E Marsten, George L Nemhauser, and Gabriele Sigismondi (1995). “The fleet assignment problem: Solving a large-scale integer program”. In: *Mathematical Programming* 70, pp. 211–232. DOI: 10.1007/BF01585938.
- Jacobs, Timothy L, Barry C Smith, and Ellis L Johnson (2008). “Incorporating network flow effects into the airline fleet assignment process”. In: *Transportation Science* 42.4, pp. 514–529. DOI: 10.1287/trsc.1080.0242.
- Jiang, Hai and Cynthia Barnhart (2013). “Robust airline schedule design in a dynamic scheduling environment”. In: *Computers & Operations Research* 40 (3), pp. 831–840. DOI: 10.1016/j.cor.2011.06.018.
- Joncour, Cédric, Sophie Michel, Ruslan Sadykov, Dimitri Sverdlov, and François Vanderbeck (2010). “Column Generation based Primal Heuristics”. In: *Electronic Notes in Discrete Mathematics* 36, pp. 695–702. DOI: 10.1016/j.endm.2010.05.088.
- Kenan, Nabil, Ali Diabat, and Aida Jebali (2018). “Codeshare agreements in the integrated aircraft routing problem”. In: *Transportation Research Part B: Methodological* 117, pp. 272–295.
- Kiarashrad, Masood, Seyed Hamid Reza Pasandideh, and Mohammad Mohammadi (2021). “A mixed-integer nonlinear optimization model for integrated flight scheduling, fleet assignment, and ticket pricing in competitive market”. In: *Journal of Revenue and Pricing Management*, pp. 1–12.
- Lohatepanont, Manoj and Cynthia Barnhart (2004). “Airline schedule planning: Integrated models and algorithms for schedule design and fleet assignment”. In: *Transportation science* 38.1, pp. 19–32.
- Lübbecke, Marco and Jacques Desrosiers (2005). “Selected Topics in Column Generation”. In: *Operations Research* 53.6, pp. 1007–1023. DOI: 10.1287/opre.1050.0234.
- Méndez-Díaz, Isabel, Juan José Miranda-Bront, Gustavo Vulcano, and Paula Zabala (2014). “A branch-and-cut algorithm for the latent-class logit assortment problem”. In: *Discrete Applied Mathematics* 164, pp. 246–263. DOI: 10.1016/j.dam.2012.03.003.
- Morley, Clive, Jaume Rosselló, and Maria Santana-Gallego (2014). “Gravity models for tourism demand: theory and use”. In: *Annals of tourism research* 48, pp. 1–10. DOI: 10.1016/j.annals.2014.05.008.
- Pessoa, Artur, Ruslan Sadykov, Eduardo Uchoa, and François Vanderbeck (2018). “Automation and Combination of Linear-Programming Based Stabilization Techniques in Column Generation”. In: *INFORMS Journal on Computing* 30 (2), pp. 339–360. DOI: 10.1287/ijoc.2017.0784.
- Petropoulos, Fotios, Gilbert Laporte, Emel Aktas, Sibel A Alumur, Claudia Archetti, Hayriye Ayhan, Maria Battarra, Julia A Bennell, Jean-Marie Bourjolly, John E Boylan, et al. (2024). “Operational Research: methods and applications”. In: *Journal of the Operational Research Society* 75.3, pp. 423–617.
- Pita, Joao P, Cynthia Barnhart, and António P Antunes (2013). “Integrated flight scheduling and fleet assignment under airport congestion”. In: *Transportation Science* 47.4, pp. 477–492. DOI: 10.1287/trsc.1120.0442.

- Santini, Alberto (2025). “Destination selection and flight scheduling for regional airlines at slot-constrained airports”. In: *International Transactions in Operational Research* 32.3, pp. 1400–1421.
- Sen, Alper, Alper Atamtürk, and Philip Kaminsky (2017). *A conic integer programming approach to constrained assortment optimization under the mixed multinomial logit model*. arXiv preprint 1705.09040. arXiv: 1705.09040 [math]. URL: <https://arxiv.org/abs/1705.09040>.
- Sherali, Hanif D, Ki-Hwan Bae, and Mohamed Haouari (2010). “Integrated airline schedule design and fleet assignment: Polyhedral analysis and Benders’ decomposition approach”. In: *INFORMS Journal on Computing* 22.4, pp. 500–513. DOI: 10.1287/ijoc.1090.0368.
- Sherali, Hanif D, Ki-Hwan Bae, and Mohamed Haouari (2013a). “A benders decomposition approach for an integrated airline schedule design and fleet assignment problem with flight retiming, schedule balance, and demand recapture”. In: *Annals of Operations Research* 210, pp. 213–244. DOI: [Spr10.1007/s10479-011-0906-3](https://doi.org/10.1007/s10479-011-0906-3)inger.
- Sherali, Hanif D, Ki-Hwan Bae, and Mohamed Haouari (2013b). “An integrated approach for airline flight selection and timing, fleet assignment, and aircraft routing”. In: *Transportation Science* 47.4, pp. 455–476. DOI: 10.1287/trsc.2013.0460.
- Sherali, Hanif D, Ebru K Bish, and Xiaomei Zhu (2005). “Polyhedral analysis and algorithms for a demand-driven reflecting model for aircraft assignment”. In: *Transportation science* 39.3, pp. 349–366. DOI: 10.1287/trsc.1040.0090.
- Sherali, Hanif D and Xiaomei Zhu (2008). “Two-stage fleet assignment model considering stochastic passenger demands”. In: *Operations Research* 56.2, pp. 383–399. DOI: 10.1287/opre.1070.0476.
- Sohoni, Milind, Yu-Ching Lee, and Diego Klabjan (2011). “Robust airline scheduling under block-time uncertainty”. In: *Transportation Science* 45.4, pp. 451–464.
- Soumis, François and Anna Nagurney (1993). “A stochastic, multiclass airline network equilibrium model”. In: *Operations research* 41.4, pp. 721–730.
- Swan, William M and Nicole Adler (2006). “Aircraft trip cost parameters: A function of stage length and seat capacity”. In: *Transportation Research Part E: Logistics and Transportation Review* 42.2, pp. 105–115. DOI: 10.1016/j.tre.2005.09.004.
- Vanderbeck, Fran cois (2005). “Implementing mixed integer column generation”. In: *Column generation*. Ed. by Guy Desaulneirs, Jacques Desrosiers, and Marius Solomon. Springer. DOI: 10.1007/0-387-25486-2_12.
- Vanderbeck, Fran cois and Laurence Wolsey (1996). “An exact algorithm for IP column generation”. In: *Operations Research Letters* 19 (4), pp. 151–159. DOI: 10.1016/0167-6377(96)00033-8.
- Vaze, Vikrant and Cynthia Barnhart (2012). “Modeling airline frequency competition for airport congestion mitigation”. In: *Transportation Science* 46.4, pp. 512–535. DOI: 10.1287/trsc.1120.0412.
- Wang, Desmond Di, Diego Klabjan, and Sergey Shebalov (2014). “Attractiveness-Based Airline Network Models with Embedded Spill and Recapture”. In: *Journal of Airline and Airport Management* 4.1, pp. 1–25. DOI: 10.3926/jairm.20.
- Wei, Keji, Vikrant Vaze, and Alexandre Jacquillat (2020). “Airline Timetable Development and Fleet Assignment Incorporating Passenger Choice”. In: *Transportation Science* 54.1, pp. 139–163. DOI: 10.1287/trsc.2019.0924.
- Wolsey, Laurence (2021). *Integer Programming*. Second. Wiley. ISBN: 9781119606536.
- Xu, Yifan, Nicole Adler, Sebastian Wandelt, and Xiaoqian Sun (2023). “Competitive integrated airline schedule design and fleet assignment”. In: *European Journal of Operational Research*. DOI: 10.1016/j.ejor.2023.09.029.
- Xu, Yifan, Sebastian Wandelt, and Xiaoqian Sun (2021). “Airline integrated robust scheduling with a variable neighborhood search based heuristic”. In: *Transportation Research Part B: Methodological* 149, pp. 181–203. DOI: 10.1016/j.trb.2021.05.005.
- Yan, Chiwei, Cynthia Barnhart, and Vikrant Vaze (2022). “Choice-based airline schedule design and fleet assignment: A decomposition approach”. In: *Transportation Science* 56.6, pp. 1410–1431. DOI: 10.1287/trsc.2022.1141.

Yan, Shangyao, Ching-Hui Tang, and Tseng-Chih Fu (2008). “An airline scheduling model and solution algorithms under stochastic demands”. In: *European Journal of Operational Research* 190.1, pp. 22–39. DOI: 10.1016/j.ejor.2007.05.053.

A Details of the Network-Wide Model Extensions

This section formalizes the model extensions introduced in Section 4.2.1 using the sets and parameters described in Table 7.

Route selection. Changing the equality constraint (6b) to a less-than-or-equal-to constraint changes the domain of the corresponding dual values $\lambda_{ij}^{(6b)}$ from \mathbb{R} to \mathbb{R}_0^+ but does not otherwise affect the column generation procedure.

Airport slot restrictions. We can limit the number of flights departing and arriving during each period by adding the following two constraints:

$$\sum_{j:(i,j) \in L} \sum_{S \in \mathcal{S}_{ijt}} y_S \leq \tau_{it}^+ \quad \forall i \in D, \forall t \in T \quad (12)$$

$$\sum_{j:(j,i) \in L} \sum_{S \in \bar{\mathcal{S}}_{jit}} y_S \leq \tau_{it}^- \quad \forall i \in D, \forall t \in T. \quad (13)$$

The reduced cost (7) of a segment schedule $S \in \mathcal{S}_{ij}$ is modified by subtracting the following terms:

$$\sum_{t \in T_S} \lambda_{it}^{(12)} + \sum_{t \in \bar{T}_S} \lambda_{jt}^{(13)},$$

where $\lambda_{it}^{(12)} \geq 0$ and $\lambda_{jt}^{(13)} \geq 0$ are the dual variables associated with (12) and (13). In the pricing subproblem, we subtract quantity $\lambda_{it}^{(12)} + \lambda_{j,t+F_{ij}^a}^{(13)}$ from coefficient \bar{c}_{ijt}^a used in (10).

Mainline capacity. We can impose flight capacity constraints on mainline flights by adding the following inequality:

$$\sum_{j \in L_g^+} \sum_{t \in U_{gj}^+} \sum_{S \in \mathcal{S}_{gjt}^+} \bar{x}_{gjtS}^+ y_S + \sum_{j \in L_g^-} \sum_{t \in U_{gj}^-} \sum_{S \in \mathcal{S}_{gjt}^-} \bar{x}_{gjtS}^- y_S \leq P_g \quad \forall g \in \Gamma. \quad (14)$$

Constraint (14) ensures that capacities are respected on mainline flights preceding and following the connecting regional flights. The first sum counts regional passengers arriving at the departure hub of mainline flight g and then transferring to g . The second term counts regional passengers arriving at the arrival hub of mainline flight g and then transferring to a regional flight. The reduced cost of a segment schedule $S \in \mathcal{S}_{ij}$ is modified by subtracting term

$$\sum_{q \in Q_{ij}} \sum_{t \in T_S \cap T_{ijq}^+} \bar{x}_{ijqtS} \lambda_{g_{ijqt}^+}^{(14)}$$

if j is a hub ($j \in H$) or term

$$\sum_{q \in Q_{ij}} \sum_{t \in T_S \cap T_{ijq}^-} \bar{x}_{ijqtS} \lambda_{g_{ijqt}^-}^{(14)}$$

if i is a hub ($i \in H$), where $\lambda_g^{(14)} \geq 0$ is the dual variable associated with (14). In the pricing subproblem, we subtract quantity $\lambda_{g_{ijqt}^+}^{(14)}$ from coefficient \bar{f}_{ijqt} if j is a hub and quantity $\lambda_{g_{ijqt}^-}^{(14)}$ if i is a hub.

Aircraft utilization. We can impose maximum average utilization constraints using the following inequalities:

$$\sum_{S \in \mathcal{S}} R_S^a y_S \leq u^a \sum_{i \in D^a} w_{i1}^a \quad \forall a \in A. \quad (15)$$

The sum on the right-hand side counts how many aircraft of type a are used in the solution. This number typically equals the fleet size N^a because airlines use all the available aircraft. The reduced cost of a schedule $S \in \mathcal{S}_{ij}$ is changed by subtracting the following term:

$$\sum_{a \in A_{ij}} |T_S^a| \cdot F_{ij}^a \cdot \lambda_a^{(15)},$$

Set	Description
$H \subseteq D$	Hubs served by the regional airline.
Γ	Mainline flights with at least one endpoint at a hub served by the regional airline.
$\Gamma^+ \subseteq \Gamma$	Mainline flights departing from a hub served by the regional airline.
$\Gamma^- \subseteq \Gamma$	Mainline flights arriving at a hub served by the regional airline. Remark: $\Gamma^+ \cap \Gamma^-$ are hub-to-hub mainline flights.
$\Gamma_i^+ \subseteq \Gamma^+$	Mainline flights departing from $i \in H$.
$\Gamma_i^- \subseteq \Gamma^-$	Mainline flights arriving at $i \in H$.
$L_g^+ \subseteq L$	Segments of type (j, i_g^+) for $g \in \Gamma^+$. In words, these are the segments operated by the regional airline with a destination airport that is the origin for a mainline flight $g \in \Gamma^+$.
$L_g^- \subseteq L$	Segments of type (i_g^-, j) for $g \in \Gamma^-$. In words, these are the segments operated by the regional airline with an origin airport that is the destination of a mainline flight $g \in \Gamma^+$.
$U_{gj}^+ \subseteq T$	Departure times for regional flights on segment $(j, i_g^+) \in L_g^+$ such that passengers can then transfer to $g \in \Gamma^+$ at $i_g^+ \in H$.
$U_{gj}^- \subseteq T$	Departure times for regional flights on segment $(i_g^-, j) \in L_g^-$ such that passengers could have transferred from $g \in \Gamma^-$ at $i_g^- \in H$.
$\mathcal{S}_{ijt} \subseteq \mathcal{S}_{ij}$	Schedules for segment $(i, j) \in L$ containing a flight departing in period $t \in T_{ij}$.
$\bar{\mathcal{S}}_{ijt} \subseteq \mathcal{S}_{ij}$	Schedules for segment $(i, j) \in L$ containing a flight arriving during period $t \in T$.
$\mathcal{S}_{gjt}^+ \subseteq \mathcal{S}$	Schedules for segment $(j, i_g^+) \in L_g^+$ containing a regional flight departing during period $t \in U_{gj}^+$.
$\mathcal{S}_{gjt}^- \subseteq \mathcal{S}$	Schedules for segment $(i_g^-, j) \in L_g^-$ containing a regional flight departing during period $t \in U_{gj}^-$.
$T_{ijq}^+ \subseteq T_{ij}$	Departure periods for regional flights in segment $(i, j) \in L$ for which parameter g_{ijqt}^+ (see below) is defined.
$T_{ijq}^- \subseteq T_{ij}$	Departure periods for regional flights in segment $(i, j) \in L$ for which parameter g_{ijqt}^- (see below) is defined.
$T_S \subseteq T$	Departure periods of flights in schedule $S \in \mathcal{S}$.
$\bar{T}_S \subseteq T$	Arrival periods of flights in schedule $S \in \mathcal{S}$.
$T_S^a \subseteq T$	Departure periods of flights operated with an aircraft of type $a \in A$ in schedule $S \in \mathcal{S}$.
Parameter	Description
$\tau_{it}^+ \in \mathbb{Z}^+$	Maximum number of departures from airport $i \in D$ during period $t \in T$.
$\tau_{it}^- \in \mathbb{Z}^+$	Maximum number of arrivals at airport $i \in D$ during period $t \in T$.
$i_g^+ \in H$	Departure hub for $g \in \Gamma^+$.
$i_g^- \in H$	Arrival hub for $g \in \Gamma^-$.
$\bar{x}_{gjtS}^+ \geq 0$	Number of passengers on the regional flight from j to i_g^+ departing during period $t \in U_{gj}^+$ according to schedule $S \in \mathcal{S}_{gjt}^+$, who will later transfer to mainline flight $g \in \Gamma^+$ at hub i_g^+ .
$\bar{x}_{gjtS}^- \geq 0$	Number of passengers on the regional flight from i_g^- to j departing during period $t \in U_{gj}^-$ according to schedule $S \in \mathcal{S}_{gjt}^-$, who transferred from mainline flight $g \in \Gamma^-$ at hub i_g^- .
$g_{ijqt}^+ \in \Gamma_j^+$	Mainline flight out of $j \in H$ taken by passengers of market $q \in Q_{ij}$ travelling on a regional flight in segment $(i, j) \in L$ departing during period $t \in T_{ij}$. This parameter is only defined for periods t such that a regional-to-mainline transfer is possible at j .
$g_{ijqt}^- \in \Gamma_i^-$	Mainline flight to $i \in H$ taken by passengers of market $q \in Q_{ij}$ travelling on a regional flight in segment $(i, j) \in L$ departing during period $t \in T_{ij}$. This parameter is only defined for periods t such that a mainline-to-regional transfer is possible at i .
$P_g > 0$	Capacity of mainline flight $g \in \Gamma$.
$R_S^a \geq 0$	Total flight time of an aircraft of type $a \in A$ in schedule $S \in \mathcal{S}$.
$u^a \geq 0$	Maximum average flight time for $a \in A$.
$F_{ij}^a \in \mathbb{Z}^+$	Flight time (in periods) from i to j , with $(i, j) \in L$, on an aircraft of type $a \in A_{ij}$.

Table 7: Sets and parameters used in Appendix A.

Set		Description
Q	★	Set of all markets, including the new ones created when the split market assumption is relaxed.
T'	★	Itinerary start periods for passengers. It is a superset of T because an itinerary could start before our planning horizon.
\hat{T}_{ij}		Departure times of flights on segment $(i, j) \in L$ in the optimized schedule.
Ω_{qt}		Set of tuples of the form (i, j, s, a) describing the regional flights with flight origin i , flight destination j , flight departure time s and aircraft type a , which are used by a passenger in market $q \in Q$ starting their itinerary at time $t \in T'$. Since we include both non-stop and one-stop passengers, the size of set Ω_{qt} can be 0, 1 or 2.
Parameter		Description
\hat{z}_{ijt}^a		Binary parameter taking value one if there is a flight on segment $(i, j) \in L$ operated by an aircraft of type $a \in A_{ij}$ at time $t \in T_{ij}^a$ in the optimized schedule.
\hat{a}_{ijt}		Aircraft type used to operate a flight on segment $(i, j) \in L$ at time $t \in \hat{T}_{ij}$ in the optimized schedule.
f_{qt}	★	Fare for a passenger in market $q \in Q$ starting their itinerary at time $t \in T'$ flying on the segment(s) operated by the regional airline. Unlike in Section 4, an itinerary in this passenger allocation model can include two regional airline segments.
α_{qt}	★	Attractiveness of the itinerary that starts at time $t \in T'$ in market $q \in Q$.
β_q	★	Outside option attractiveness for market $q \in Q$.
M_q	★	Size of market $q \in Q$.

Table 8: Sets and parameters used in Appendix C. A ★ next to the parameter name indicates that this parameter has been redefined with a related but slightly different meaning compared to Section 4.

where $\lambda_a^{(15)} \geq 0$ are the dual variables associated with (15). In the pricing subproblem, we subtract $F_{ij}^a \lambda_a^{(15)}$ from coefficients \hat{c}_{ijt}^a .

B Early Termination

To use our algorithm for faster decision-making or to perform multiple scenario analyses on large instances, we might want to reduce its computational runtime by terminating the column generation procedure early. We would then obtain a primal solution via the reduced master heuristic. Unfortunately, we would not have a valid dual bound to assess the solution's quality because we did not solve [MP] to optimality. In the following, we derive a Lagrangian bound (see, e.g., Lübbecke and Desrosiers 2005) that is valid even in case of early termination. Let $Z^* \in \mathbb{R}$ be the unknown optimal objective value of [MP], and $Z \in \mathbb{R}$ be the optimal objective value of an instance of the reduced [MP]. Let $\Delta_{ij} > 0$ be the highest reduced cost of a column of \mathcal{S}_{ij} , according to the duals of the reduced [MP]. We can obtain Δ_{ij} by solving all pricing subproblems to optimality. We then choose, for each subproblem, the generated column with the highest reduced cost, if positive. Otherwise, if a subproblem for (i, j) proves that no positive reduced cost column exists, $\Delta_{ij} = 0$. We can then use a Lagrangian bound to obtain that $Z^* \leq Z + \sum_{(i,j) \in L} \Delta_{ij}$, where constraints (6b) act as convexity constraints on each set \mathcal{S}_{ij} (see, e.g., Wolsey 2021, §11.5.1).

C Passenger Allocation Model for the Regional-Regional Split Market Assumption

In this appendix, we present the approach used to allocate passengers to the optimized schedule under the split market assumption (see Section 7.4) when new markets are introduced. Recall that Q denotes the set of all markets (sets and parameters used in this appendix are defined in Table 8). We retain the

symbols used in the main body of the paper (α and β for the attractiveness of all itineraries and the outside option, f for the fare, M for the market size), but we give them different indices and slightly different meanings as explained in Table 8. We now solve a network-wide passenger allocation model for the entire network rather than solving a segment-by-segment pricing subproblem as in Section 4.1; therefore, we update the meaning of the above parameters to reflect this change. Similarly, we continue to use the letter x for the following variables: The variable $x_{q0} \geq 0$ for $q \in Q$ denotes the number of passengers in market q who choose the outside option. The variable $x_{qt} \geq 0$ for $q \in Q$ and $t \in T'$ denotes the number of passengers in market q who start their itinerary at time t and use one or more regional flights. We note that the passenger's itinerary can start with a regional flight (eventually followed by a mainline or another regional flight) or with a mainline flight (followed by a regional one). Indeed, we include in our model the following itinerary types: (a) Regional; (b) Regional + Mainline; (c) Mainline + Regional; (d) Regional + Regional. Whereas our main model included itineraries of types (a)–(c), we also include itineraries of type (d) in the following passenger allocation model.

A Linear Programming formulation for the network-wide passenger allocation model reads as follows.

$$\begin{aligned} \max \quad & \sum_{q \in Q} \sum_{t \in T'} f_{qt} x_{qt} & (16a) \\ \text{subject to} \quad & \beta_q x_{qt} \leq \alpha_{qt} x_{q0} & \forall q \in Q, \forall t \in T' & (16b) \\ & x_{q0} + \sum_{t \in T'} x_{qt} = M_q & \forall q \in Q & (16c) \\ & \sum_{\substack{(q,t) \in Q \times T': \\ (i,j,s,\hat{a}_{ijs}) \in \Omega_{qt}}} x_{qt} \leq P^{\hat{a}_{ijs}} & \forall (i,j) \in L, \forall s \in \hat{T}_{ij} & (16d) \\ & x_{qt} \geq 0 & \forall q \in Q, \forall t \in T' & (16e) \\ & x_{q0} \geq 0 & \forall q \in Q. & (16f) \end{aligned}$$

The objective function (16a) maximizes the fare revenue. Constraints (16b), (16c) and (16d) are analogous, respectively, to (2c), (2d) and (2e), but now adapted to the network-wide passenger allocation model. We remark that market sizes M_q for new markets were estimated according to the formula

$$M_q = k \cdot \frac{\text{out}(o) \cdot \text{in}(d)}{\text{dist}(o, d)},$$

where o and d are the market origin and destination, $\text{out}(o)$ is the number of passengers in markets with origin o , $\text{in}(d)$ is the number of passengers in markets with destination d , $\text{dist}(o, d)$ is the geodesic distance between airports o and d , and k is a parameter estimated using the DB1B dataset.

D Case Study Setup

We test our modeling and computational approach on instances based on real-world data on the U.S. domestic market gathered from multiple data sources, namely the Airline On-Time Performance (AOTP) dataset, the Airline Origin and Destination Survey (DB1B) and the Form 41 dataset from the Bureau of Transportation Statistics 2023a; Bureau of Transportation Statistics 2023b; Bureau of Transportation Statistics 2023c, and the aircraft registration database from the Federal Aviation Administration 2023. We used data from June 2022 and considered United Airlines (UA) as the mainline carrier and Mesa Airlines (YV), Republic Airways (YX) and SkyWest Airlines (OO) as regional airlines. This choice is motivated by the greater data availability, partly due to data reporting obligations for airlines above a certain size, compared to alternatives. We consider both the optimization of each regional airline by itself and jointly. In the latter case, our instance uses seven hubs (ORD, EWR, IAD, IAH, LAX, SFO and DEN) and 149 spokes for a total of 534 non-stop segments and 3312 markets (see Figure 2). We now describe how to estimate the main model parameters from the above data sources.

Fleet composition, costs, fares and market sizes. We use our data sources to determine how many aircraft each regional airline employs to operate flights ticketed under United Airlines’ *United Express* brand. In addition to collaborating with United Airlines, Mesa also operated under the American Airlines *American Eagle* brand until 2023, Republic operates under the *American Eagle* and Delta’s *Delta Connection* brands, and SkyWest operates under the *Alaska Airlines*, *American Eagle*, and *Delta Connection* brands. The aircraft registration database and the airline company websites report the aircraft capacities. We estimate operating costs using the short-haul formulas of Swan and Adler (2006), adjusted for inflation. We use fares data from the DB1B database. For multi-leg itineraries, only the total fare is available; in this case, we assume that the regional leg fare is proportional to its length. We derive flight times and minimum turnaround times from AOTP data; in particular, we have minimum turnaround times that are functions of both the aircraft type and the airport. Note that, as is the typical case for regional airlines, there are only a small number of distinct aircraft types across the entire fleet of our case study networks of regional airlines. They consist primarily of Canadair Regional Jets (CRJs) and Embraer aircraft. We generate our data on market sizes and reasonable paths through the network using the DB1B database. Since DB1B includes a 10% sample of passenger tickets, we scaled it up by multiplying by 10, and then by a further factor of 1.1 to account for passengers who used the outside option, in order to obtain the total market sizes.

Attractiveness data. We compute attractiveness values α_{ijqt} using the method proposed by Barnhart, Fearing, et al. (2014) and assuming a minimum connection time of thirty minutes for one-stop itineraries. To do this, we used the mainline schedules of a reference day, June 29th 2022. Given the market sizes, the attractiveness values α , and the number of passengers carried from our data sources, if the flights were not subject to seating capacities, we could apply eq. (1) to recover the attractiveness values of the outside option β . However, because of the seating capacity limits, the ratio between the number of passengers taking each option (including the outside option) does not necessarily match the ratio between the options’ attractiveness values. Therefore, we use the following approach to derive realistic β . First, we produce estimates $\hat{\beta}$ disregarding seating capacities. Note that if these estimates were correct, solving problem (2a)–(2h) with the fixed real-world schedule would give passenger numbers that closely match those observed in the data. Otherwise, we should adjust the values of $\hat{\beta}$ to make the numbers match. However, each β_{ijq} can potentially require a different adjustment factor, making the search for such factors computationally expensive and prone to overfitting to the data. To overcome this issue, we adjust all estimates $\hat{\beta}$ simultaneously using a single multiplicative factor from a small search grid (fifteen logarithmically distributed values between 10^{-3} and 1, in our case). To select the best factor, we quantify the similarity between the passenger numbers obtained from (2a)–(2h) and the real ones using their mean absolute deviation.

E Impacts of Demand Uncertainty and Revenue Management practices

The results presented in Section 7 have assumed that the demand for each passenger itinerary is deterministic and that the average fares for each itinerary are known and fixed. In practice, passenger demand can fluctuate considerably from one day to another, and this uncertainty can impact the actual number of passengers carried by the airline. Too few passengers on a given day can lead to reduced revenue, while too many passengers compared to those estimated can lead to some passengers not being accommodated (or getting “spilled”) due to aircraft seating capacity limits. Spill estimation models usually make some assumptions about the demand distribution and compute the expected number of passengers as the expected value of the demand distribution truncated by the number of seats. As mentioned in Section 7, our results have indirectly captured the effects of demand uncertainty by imposing a hard upper limit of 93% on the maximum number of seats sold in each segment.

The load factor is a non-decreasing function of the *demand factor*. The *demand factor* is the ratio of passenger demand divided by seat capacity. All else being equal, adding more flights to a segment, thus increasing the seat capacity, would decrease the demand factor. Our “Freq ± 1 ” and “Freq ± 2 ” schedules have more segments with an increase rather than a decrease in flight frequency and the

	Without correction			With correction		
	Revenue	Cost	Profit	Revenue	Cost	Profit
Status Quo	3,387,407	1,535,941	1,851,465	3,387,407	1,535,941	1,851,465
Freq ± 0	3,471,732 (+2.49%)	1,538,942 (+0.20%)	1,932,790 (+4.39%)	3,458,547 (+2.10%)	1,538,942 (+0.20%)	1,919,605 (+3.68%)
Freq ± 1	3,713,894 (+9.64%)	1,670,419 (+8.75%)	2,043,474 (+10.37%)	3,589,626 (+5.97%)	1,589,970 (+3.52%)	1,999,656 (+8.00%)
Freq ± 2	3,790,470 (+11.90%)	1,692,362 (+10.18%)	2,098,108 (+13.32%)	3,637,440 (+7.38%)	1,622,231 (+5.62%)	2,015,209 (+8.84%)

Table 9: Comparison of key financial metrics on the joint network instance (YV + YX) with and without applying the ICO process described in Appendix E.

	Status Quo	Without correction			With correction		
		Freq ± 0	Freq ± 1	Freq ± 2	Freq ± 0	Freq ± 1	Freq ± 2
Number of flights	218	218	240	246	218	228	236
Avg flight distance	571.0	571.0	584.7	587.4	571.0	579.2	581.1
RASM	36.5	37.3	35.5	35.2	37.3	36.0	35.8
CASM	16.6	16.5	16.0	15.7	16.6	16.2	15.9
Passengers carried	14908	14920	16211	16666	14922	15676	16151
Non-stop	66%	66%	66%	66%	66%	66%	66%
One-stop	34%	34%	34%	34%	34%	34%	34%
Avg fare	\$227	\$233	\$229	\$227	\$231	\$229	\$227
Non-stop	\$223	\$229	\$227	\$225	\$228	\$227	\$225
One-stop	\$242	\$246	\$238	\$236	\$244	\$240	\$238
Light retiming	—	63%	58%	60%	65%	61%	60%

Table 10: Key operational metrics of the solutions of the joint network instance (YV + YX) with and without applying the ICO process described in Appendix E.

number of seats per segment. This implies that we should expect a slightly lower demand factor. Of course, by offering flights at desirable departure times, we hope to attract more passenger demand, thus increasing the numerator of the demand factor. Our results in Section 7.1 simplify these complexities using a hard upper limit on load factors.

In addition, nearly all airlines practice revenue management (RM) methods, which greatly affect the fares of spilled passengers because they try to ensure that spilled passengers are predominantly low-fare passengers in a given segment. Our results in Section 7.1 simplify this effect by assuming identical fares for all passengers in a given market. Belobaba and Farkas 1999 showed that revenue management practices affect spilled passenger fares and the number of spilled passengers. In this section, we use the approach proposed by Belobaba and Farkas 1999 to approximately capture the effects of demand uncertainty and revenue management practices on the number of passengers carried, the average fares of the passengers carried, and consequently the profit of the airline. Similar to the process described in Section 5.4 of Vaze and Barnhart 2012, we change the fare of each market in each segment by applying the correction suggested by Belobaba and Farkas 1999. This results in a new instance with different values of the fare parameters f_{ijqt} . We then re-optimize this instance by running our column generation algorithm again without imposing any hard limit on the load factors, thus obtaining a new schedule informed by the new corrected fares and the corresponding changes in passenger demand. However, the new optimized schedule adds and removes flights (and thus changes seat capacity) in some or all segments. This, in turn, introduces a feedback effect because it alters the extent of the fare correction. Therefore, we propose an iterative procedure in which a correction phase and a reoptimization phase alternate until convergence. In our case, two iterations are enough to converge to stable fare values.

Table 9 presents the key financial performance metrics of the solutions obtained for the joint network instance (YV + YX) by our approach without (second to fourth column) and with (fifth to seventh column) the aforementioned Iterative Correction and Optimization (ICO) process. In addition, Table 10 presents the key operational metrics for the same instance without (third to fifth column) and with (sixth to eighth column) the ICO process. First, as expected, the operational and financial performances of the “Freq±0” instances with and without the ICO process are quite similar to each other. However, in the “Freq±1” and “Freq±2” instances, the ICO process leads to fewer and shorter flights with slightly higher RASM and CASM values, slightly fewer passengers and similar average fares, leading to slightly lower profits. Nevertheless, across all three scenarios (“Freq±0”, “Freq±1”, and “Freq±2”), we continue to see a significant profit increase of between 3.68% and 8.84% relative to the “Status Quo” solutions. This shows the relative stability of the overall improvements generated by our approach under the effects of demand uncertainty and RM practices.