Unified tail assignment and maintenance task scheduling: A decision support framework for improved efficiency and stability

Pescio, L.V.L.; Ribeiro, M.J.; Santos, Bruno F.

DOI
10.1016/j.orp.2025.100363

Publication date
2025

Document Version
Final published version

Published in
Operations Research Perspectives

Citation (APA)
Pescio, L. V. L., Ribeiro, M. J., & Santos, B. F. (2025). Unified tail assignment and maintenance task scheduling: A decision support framework for improved efficiency and stability. *Operations Research Perspectives*, *15*, Article 100363. https://doi.org/10.1016/j.orp.2025.100363

Operations Research Perspectives 15 (2025) 100363

Contents lists available at ScienceDirect

Operations Research Perspectives

journal homepage: www.elsevier.com/locate/orp

Unified tail assignment and maintenance task scheduling: A decision support framework for improved efficiency and stability

Luigi Pescio, Marta Ribeiro *, Bruno F. Santos

Aerospace Faculty, Delft University of Technology, Delft, The Netherlands

ARTICLE INFO

Keywords:
Tail assignment
Maintenance task scheduling
Fleet availability
Time-space network

ABSTRACT

Flight and maintenance scheduling pose conflicting objectives: while maintenance is vital for ensuring aircraft airworthiness, it comes at the cost of taking aircraft out of operation. In current operations, airlines manually handle tail assignment and maintenance task scheduling separately, missing an opportunity to strike a better balance. This division leads to wasted maintenance resources, restricted fleet availability for schedule flexibility, inconsistent planning, and neglect of schedule resilience. This study presents a novel approach that integrates tail assignment and maintenance scheduling into a unified decision-support framework. An integer program, tailored to meet airline-specific requirements and constraints, is combined with an innovative time-space network (TSN). The TSN incorporates two distinct spaces for maintenance and network activities. The primary objective is to generate feasible plans that increase schedule efficiency (i.e., no cancellations, high fleet availability, high fleet health, and optimal use of maintenance resources) and schedule stability (i.e., limited number of late arrival disruptions during operations) the day before operation. Additionally, this framework addresses overlooked aspects in the literature: it treats maintenance tasks as variable interval activities based on aircraft-specific needs, departing from the traditional fixed interval approach. The performance of the framework is tested with real-data provided by a major European single hub-to-spoke airline, with a heterogeneous fleet of over 50 wide-body aircraft. Historical data from arrival delays is used to create robust buffers that mitigate delay propagation. A 17% reduction in maintenance time was achieved compared to the airline's current plans, resulting in a 10% increase in fleet availability on the day of operations. This improvement is attributed to higher labour and task interval utilization, indicating the framework's superior efficiency in scheduling maintenance tasks. Lastly, the framework produced plans more resilient to arrival delays, reducing the number of disruptions and delay propagation over 40%. This framework can be used as a decision-support tool for airlines, enabling the creation of schedules that are both robust against delays and optimized for fleet utilization.

1. Introduction

Airline schedules suffer from disruptions daily. Poor weather conditions, airport congestion, unavailable personnel, and unplanned mechanical failures have led to 30.1% of flights not running on schedule in Q2 of 2023 according to [1]. Consequently, the tasks of tail and maintenance must constantly be adjusted to new information arriving. However, most airline planners still do this manually, which is complex and time consuming. Airline operators manually mitigate disruptions by altering aircraft schedules, crew schedules, and passengers' itineraries. For large airlines with extensive networks and complicated schedules, an operator cannot find manually an optimal solution.

Additionally, neither tail assignment nor maintenance task scheduling are static decisions. Instead, these are highly intertwined tasks with conflicting interests. Tail assignment planners decide where each available aircraft should fly, decreasing operating and cancellation costs. In turn, maintenance schedulers guarantee the airworthiness of the aircraft by removing it from operations. This inefficiency in collaboration leads to sub-optimal plans, with increased variability in schedule performance. Additionally, due to the complexity of the problem, involving a multitude of maintenance tasks, flights, and aircraft, all subject to stringent operational constraints, planners often end up overlooking robustness.

Leading airlines are looking for fast operating tools that provide robust and integrated decision support for tail and maintenance task assignment, all with the aim of facilitating disruption management.

* Corresponding author.
E-mail address: m.j.ribeiro@tudelft.nl (M. Ribeiro).

https://doi.org/10.1016/j.orp.2025.100363
Received 10 June 2025; Received in revised form 16 October 2025; Accepted 5 November 2025
Available online 19 November 2025

However, despite growing interest, extensive modelling of simultaneous tail assignment and maintenance task scheduling is lacking in the current literature. Predominantly, focus has been reserved for stable and/or flexible network plans, however, often disregarding maintenance constraints. Most advanced models, for example, [2,3], treat maintenance as a fixed time, fixed interval, non-aircraft-specific activity when creating stable plans. While robust maintenance planning models, such as that presented by Sadjadi et al. [4], disregard network operations altogether. In reality, network and maintenance are highly intertwined activities that if integrated can potentially increase schedule stability and efficiency. For instance, more efficient task planning may decrease maintenance time, thereby increasing useful buffer time, helpful for preventing disruptions from flight late arrivals.

Implementing a framework for simultaneous tail assignment and maintenance task scheduling can support airline planners in the decision-making process, improving the trade-off between useful buffer time and maintenance time. A well-designed framework has the potential to create more efficient, stable plans that reduce operating costs and disruptions on the day of operations. However, to be suitable for implementation, a decision support framework must comply with numerous requirements set by the airline. It should have a short execution time, allowing planners to make adjustments as new information arrives, such as new maintenance tasks. The framework should create feasible plans to ensure aircraft airworthiness and that adhere to the airline's tail restrictions, maintenance schedule, and resource availability. This paper addresses these aspects. The innovative framework developed in this work assigns tails and schedules maintenance tasks simultaneously. At the core of the framework is an efficient linear programming (LP) optimization model that follows a hierarchical structure set by airline planners. It ensures continuous aircraft airworthiness and respects airline-specific rules and schedules. The assignment of tasks is restricted by the availability of material, machinery, method, and manpower (4M). Unlike the literature and airline practices, maintenance slots are modelled as tail-specific activities based on individual aircraft needs, with the intention of improving schedule efficiency. Finally, this study incorporates stability using historical flights' arrival delay data. Thereby, this study provides new insights into integrated tail assignment and maintenance task scheduling, with the aim of creating more efficient and stable plans in the process.

This paper follows the subsequent structure. The operational context in this work is defined in Section 2. In Section 3, a literature overview provides insights into network planning, maintenance planning, and robust planning research. This is followed by the framework's methodology, which includes the time-space network (TSN), slot generation process, and LP model, in Section 5. Afterwards, the hypotheses for the results in this framework are presented in Section 6. The case study that the framework is applied is described in Section 7, alongside the costs employed. Subsequently, Section 8 presents the results obtained by implementing the framework in a real airline case study, critically demonstrating its validity and benefits in practice. Validation of the hypotheses are discussed in Section 9, alongside recommendations for future research. Lastly, Section 10 concludes this work.

2. Operational context

This section presents the operational context for the framework developed. Although this study focuses solely on the unification of the tail assignment and maintenance task period before the day of operations, these are part of a larger operational sequence which defined the input and outputs of these phases. Thus, their definition is of relevance to the understanding of this work. This section highlights the fact that airline network and maintenance planning are complex and multi-stage processes beginning months before the actual day of operations, as shown in Fig. 1. Both are highly intertwined from the start, and highly influence each other, but are often treated as separate decision processes in the literature and by airlines.

2.1. Network planning

Network planning starts with schedule design, which is the process in which airline marketing and network divisions use demand forecasts to decide which Origin–Destination (O–D) pairs to fly [5]. Schedule design begins many months before the day of operation. Additionally, due to the complexity and size of airline network planning, it is created in sequential steps, both in practice and in the literature. The next step, fleet assignment, assigns aircraft fleets to individual flights according to passenger demands and operating costs such that total profit is maximized [6,7]. Third, a few days before operations up until the day of operations, tail assignment assigns flights to individual registrations, ensuring that maintenance requirements are respected, with often the objective of minimizing operating costs [8,9]. Finally, on the day of operation disruption may lead to modifications to the strategic schedule [10,11].

2.2. Maintenance planning

Maintenance planning is equally a multi-stage problem. It begins with check scheduling in which long-term maintenance activities are assigned to tails, often for multiple years [12,13]. Next, during the flight schedule design phase, time slots, are reserved for maintenance. A few days before the day of operations, along with tail assignment, maintenance tasks are assigned to maintenance slots, while ensuring the 4M constraints [14–16]. Lastly, on the day of operations, maintenance disruption management is performed to cope with new tasks, missing resources, and/or sudden schedule changes [17].

Disruptions which become known on the day of operations may include: aircraft late arrival (e.g., weather conditions restrict the airport's landing capacity causing the aircraft's inbound flight to be delayed leading to delay propagation into its next planned job), an aircraft is unavailable (e.g., a fault during the aircraft's last flight may cause an Aircraft on Ground (AOG) situation), slot schedule changes (e.g., unavailable towing services may cause late delivery of aircraft, shortening the duration of the maintenance slot), maintenance resources may be unavailable (e.g., due to decreases in workforce availability due to unforeseen absence and/or unavailability of material results in lower maintenance task execution rate). These disruptions require planners to make changes to the original plans to guarantee smooth operations, aircraft airworthiness, and minimal passenger inconvenience. They may employ a range of solutions including aircraft swapping, delaying flights or maintenance slots, postponing maintenance if the aircraft's health permits, deploying a reserve aircraft, and, as a last resort, cancelling flights. When choosing which recovery strategy to implement, planners aim to reduce recovery costs and minimize deviation from the original plan. The result is a revised tail and task assignment.

3. Literature overview

This section presents the current state-of-the-art for the concepts previously defined in Section 2. Given the fact that most of these are treated separately in research, different sub-sections are presented below.

3.1. Network planning

Has been covered extensively in literature, specifically the fleet assignment and tail assignment problems. The former has been typically solved for maximum profits using Integer Programming (IP), as demonstrated by Abara [6], Barnhart et al. [18]. The latter is often based on a Line of Flight (LOF), where aircraft are assigned to predefined strings of flight legs. The challenge is computing efficient and feasible LOFs - a large solution space often requires authors to use methods to simplify it. Barnhart et al. [8] used a Mixed Integer Linear Programming (MILP)

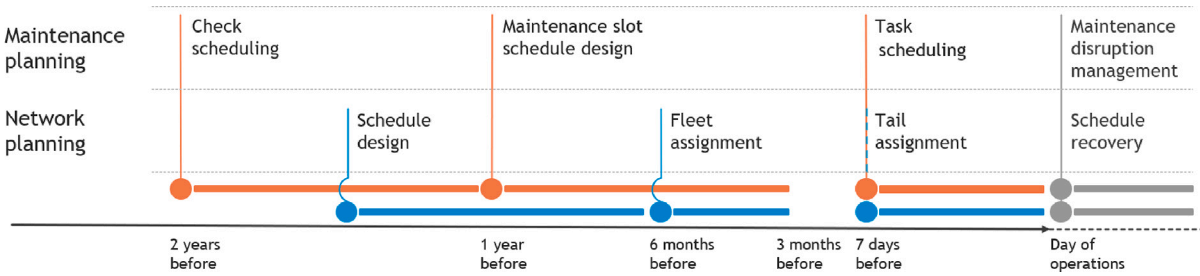

Fig. 1. The general time frame of airlines' operations planning processes. Maintenance and network planning time frames are represented in orange and blue, respectively. On the day of operations, disruptions will directly affect maintenance and network planning (represented in grey).

model and a branch-and-price algorithm, while [9] used a similar MILP model and a heuristic method to simplify the problem.

Network planning literature largely disregards maintenance. It is often only considered at a strategic level in the tail assignment problem to ensure feasible plans, being treated as a fixed interval, fixed duration activity, equal for each aircraft in the fleet. Barnhart et al. [8] assign LOFs that start and end at maintenance stations. Alternatively, [9] assumed overnight maintenance at maintenance stations in their LOF-based MILP model. However, maintenance is aircraft-specific, as each aircraft has different maintenance requirements and tasks. Sarac et al. [19] proposed a route-building MILP model, matted to a branch-and-price algorithm, that respects aircraft's allotted flight hours by routing aircraft with tasks due the next day to a maintenance station. Nonetheless, detailed task assignment is still ignored.

3.2. Maintenance planning

Maintenance planning has been largely neglected from research, with the only instances when its effect is considered in tail assignment studies. A few works focus on maintenance task scheduling. Papakostas et al. [16], Shaukat et al. [20], Witteman et al. [15] solve the task assignment problem in the strategic or operational phase considering the uncertainty in the arrival of tasks, with the objectives of minimization of costs and task interval waste. Van Kessel et al. [17] created a MILP task rescheduling model that copes with disruptions. However, they all disregard network considerations. Lagos et al. [21] integrates tail assignment and maintenance task scheduling with dynamic programming and a Markov decision process to schedule LoFs and maintenance tasks into nighttime slots. While this work prevents maintenance tasks from going due, it disregard key aspects such as the availability of resources and realistic tail assignment complexity.

A collection of airline operation planning research has been looking at a relatively new concept i.e., robust planning, with the objective of facilitating and alleviating pressure on schedule recovery on the day of operation. Robust planning begins from the inception of network and maintenance planning up until the day before operations. It has two objectives: to increase schedule stability, by decreasing disruptions on the day of operations; and/or increase schedule flexibility, by making disruptions easier to recover. This concept has been divided into two sub-concepts: proactive absorption and proactive recovery [22]. The former has the goal of designing plans resistant to disruptions, ensuring schedule stability. In the past, researchers have achieved schedule stability in multiple ways, the most common being: time buffers, flight retiming, and robust delay routing. Proactive recovery aims to design plans that are easily adjusted when disruptions occur, and thus ensure schedule flexibility. The most common method to promote schedule flexibility has been increasing opportunities to swap aircraft. However, in general, methods are few and not widely researched. The following paragraphs present more information on the existing research.

3.3. Proactive absorption

AhmadBeygi et al. [23] proposed an IP model that assigns slack with the objective of minimizing flight delay during the schedule design phase. Later, a mathematical model for integrated fleet assignment and flight retiming was developed, which determines flight departure times by maximizing a surrogate robustness measure [24]. Most recently, [3] built a two-stage model for robust multi-fleet aircraft routing. First, a MILP finds optimal aircraft routes. Second, a heuristic approach based on Monte Carlo simulation flights are re-timed with the aim of maximizing on-time performance.

Robust delay routing often involves using historical delay data to create more stable aircraft routes. Lan et al. [25] were among the first to study this concept. Using a column generation algorithm, they assigned LoFs to aircraft intending to minimize the expected propagated delay, determined from historical airline data. Aircraft routing was then followed by a MILP retiming model that minimizes passengers' misconnections. A similar model was developed by Liang et al. [2], who also considered maintenance, planned overnight every three days. Lastly, [26] proposed a novel approach to proactive absorption, scenario-based recoverable robust routing. By using LP in tandem with column generation and bender decomposition, aircraft routing minimized recovery costs from a set of historical disruption scenarios.

Finally, researchers have created stability in maintenance plans in comparable ways. Weide et al. [27] used a genetic algorithm for robust check-scheduling, running ten scenarios with check duration uncertainty, while [4] used Monte Carlo simulations to take uncertainty samples from task duration sets in their robust maintenance task scheduling ϵ-Conservative model. However, in this field, detailed maintenance and network planning are also solved separately.

3.4. Proactive recovery

Ageeva [28] pioneered the concept of using swapping opportunities to increase schedule flexibility, by creating a LOF-based column generation aircraft routing model with the objective of maximizing swapping opportunities. The method was re-proposed by Burke et al. [29], who solved the tail assignment problem with a multi-meme memetic algorithm by integrating multiple robustness strategies, for both stability and flexibility. They re-timed flights to optimize stability, given by the probability that each flight departs on time, and flexibility, given by the probability that each flight involved in a swap departs on time. Lastly, [30] presented an optimization algorithm that suggests the best swapping opportunities between LoFs of aircraft. Their objective was to maximize a novel measure called maintenance reachability, by reducing aircraft requiring maintenance but assigned to an LOF with no maintenance opportunities. Nonetheless, they treated maintenance as a fixed-interval activity.

Few other methods for proactive recovery have been researched. One is partitioning the schedule into sub-networks to safeguard high-revenue itineraries [31]. Another is to promote station purity (i.e., reducing the number of aircraft types visiting an airport) to facilitate

Table 1
Overview of airline operations planning literature. Abbreviations in the table: BD: Bender Decomposition; BP: Branch-and-Price Algorithm; CG: Column Generation; DP: Dynamic Programming; DRR: Delay robust routing; ϵ-C: ϵ-Conservative; FA: Fleet assignment; FR: Flight retiming; GA: Genetic Algorithm; IP: Integer Programming; MA: Memetic Algorithm; MC: Monte Carlo Simulations; MCDM: Multi-Criteria Decision Making; MILP: Mixed Integer Linear Programming; MSU: Maintenance scheduling with uncertainty; PA: Proactive absorption; PR: Proactive recovery; RRR: Recoverable robust routing; SC: Short cycles; SD: Schedule design; SH: Search Heuristics; SN: Sub-networks; SO: Swapping opportunities; SP: Station purity; TA: Tail assignment; WFD: Worst-Fit Decreasing Algorithm.

Paper	SD	FA	TA	Method	Maintenance	PA	PR
[6]		✓		IP			
[8]			✓	IP & BP	Fixed interval activity		
[28]			✓	IP & CG			SO
[18]		✓		IP			
[19]			✓	IP & BP	Aircraft specific need		
[9]			✓	IP & SH	Fixed interval activity		
[31]		✓	✓	IP & SH			SN
[33]		✓		IP			SC
[25]			✓	IP & CG		FR & DRR	
[32]		✓		IP & CG			SP
[23]	✓			IP		FR	
[24]		✓		IP		FR	
[16]				MCDM	Task scheduling		
[29]			✓	MA	Fixed interval activity	FR & DRR	SO
[30]			✓	IP	Fixed interval activity		SO
[35]			✓	IP & CG		FR & DRR	
[26]			✓	IP & CG & BD		RRR	
[2]			✓	IP & CG	Fixed interval activity	FR & DRR	
[3]			✓	IP & MC	Fixed interval activity	FR	
[21]			✓	DP	Task scheduling		
[20]				IP	Task scheduling		
[4]				ϵ-C & MC	Task scheduling	MSU	
[15]				WFD	Task scheduling		
[27]				GA	Check scheduling	MSU	
[17]				IP	Task rescheduling		
[34]			✓	IP	Fixed interval activity		
[36]			✓	MILP			
This paper			✓	IP	Task scheduling	DRR	

aircraft swapping [32]. Lastly, the idea behind short cycle scheduling together with hub-isolation (i.e., LoFs with few flights which include only one hub airport) is to prevent disruptions from spreading across hubs and affecting many later flights [33]. Overall, most of these methods are not applicable to all types of airline networks, are less effective in creating flexibility, and still disregard maintenance.

3.5. Research gap

Based on this overview of the literature, it is concluded that robust planning that integrates both tail assignment and maintenance task scheduling is currently not addressed in the literature. An overview is presented in Table 1. Building stable tail assignment plans, often by re-timing flights and/or using historical delay data, has been the focus. For the most part, maintenance in the tail assignment problem has either been excluded or treated as a fixed time, fixed interval, non-aircraft specific activity. However, maintenance is aircraft-specific. Works such as [3,21], the more similar to the framework in this study, perform tail assignment and maintenance scheduling without considering the availability of resources, essential in a real operational scenario. Recent studies [17,34] cover maintenance to a limited extent, focusing on either task rescheduling when disruptions occur or treating maintenance tasks as fixed interval activities. Such falls short of more dynamic decisions where maintenance tasks are scheduled with dynamic intervals for the best balance with tail assignment.

To the best of the author's knowledge this is first work where (1) tail assignment and maintenance task scheduling are approached simultaneously, (2) maintenance scheduling has the added complexity of being non-fixed and based on the needs of the specific tail, and (3) scheduling takes into account resource availability and historical delays. Simultaneous tail assignment and task scheduling has the potential decrease airlines' operating costs (schedule efficiency) and recovery costs (schedule stability). This research, therefore, aims to solve these two processes simultaneously, by assigning tails considering aircraft-specific, variable interval maintenance tasks. Finally, by considering historical arrival delays, this framework will mitigate the propagation of delays throughout rotations.

4. Tail assignment and maintenance tasks scheduling framework

The Tail Assignment and Maintenance Tasks Scheduling framework developed in this study is structured as shown in Fig. 2. On the day before operations, with the inputs shown on the left, airline planners assign tails and schedule maintenance tasks. They do so often for the upcoming two to three days. At the end of the day, planners deliver their final plan to the operations control team. The latter will manage disruptions and revise the tail and task assignments on the day of operations. As such, this framework is designed to support tail assignment and maintenance task scheduling on the day prior to operations. The remaining parts of this section will cover the framework's formulation in detail. Section 4.1 describes the inputs of the framework. Section 4.2 discusses the constraints for both tail assignment and maintenance task scheduling. Definition of the maintenance slots is presented in Section 4.3. Lastly, Section 4.4 shows the main objectives of this framework.

4.1. Inputs

Airline planners assign tails and schedule maintenance tasks using generally three inputs: a network schedule, a maintenance schedule, and a backlog of maintenance tasks. These same inputs are used by the framework:

- *Network schedule*: Is composed of flights. In the framework, flights have been combined in short sequences of flights (often just two) starting and ending at the same airport called rotations. This suits the network of hub-and-spoke carriers, with mostly only flights from and to their hubs.
- *Maintenance schedule*: Provides insight into the maintenance slot schedule and fixed maintenance activities. Each is explained in further detail below:
  - *Maintenance slot schedule*: It is a plan, often one week long, outlining maintenance scheduling opportunities. From it, one can obtain maintenance slots start and end times, work location (hangar or platform), assigned full time equivalents (FTEs) per hour, duration, and suitable aircraft type.
  - *Fixed maintenance Activities*: Include long-term letter checks, modifications, and other maintenance activities scheduled several weeks in advance. In the framework, all fixed maintenance activity must be scheduled, and tails cannot be assigned to overlapping jobs during these checks.
- *Maintenance tasks backlog*: A backlog of aircraft-specific maintenance tasks ready for scheduling. Following common airline practice, maintenance tasks are subdivided into two categories: preventive and corrective tasks. The differences are as follows:
  - *Preventive tasks*: Recurring tasks executed following interval requirements imposed by regulatory agencies (e.g., the European Aviation Safety Agency (EASA)). Airlines also have the opportunity to add additional preventive tasks. Once a task is executed the due date is reset. An example is visual inspection of aircraft brake wear, which repeats at a fixed number of take-off and landing cycles.

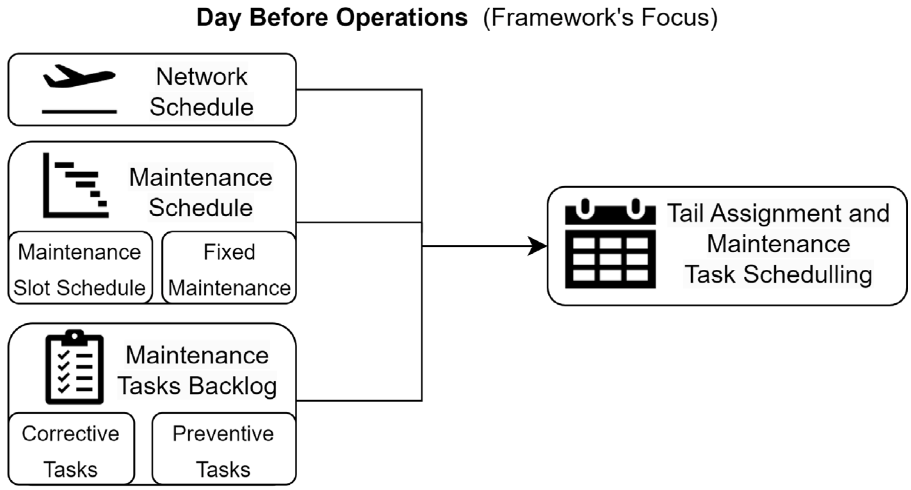

Fig. 2. Overview of the tail assignment and maintenance task scheduling framework.

Table 2
Common subdivision of maintenance tasks used in practice and in this modelling framework.

		Preventive	Corrective
Mandatory	Tasks with fixed due dates	Requirements: Minimum Required Item (MRI), Engineering Order (EO), and Structural Defect Report (SDR)	Ad-hoc, Minimum Equipment List (MEL)
Deferrable	Tasks with soft due dates		Non Safety Related Equipment (NSRE)

– *Corrective tasks*: One-off tasks executed only once before their due date. Corrective tasks also derive from regulatory agencies (e.g., Minimum Equipment List (MEL)) or the airline (e.g., Non-Safety Related Equipment (NSRE)). The replacement of brake pads after a preventive check is an example of a MEL, while a broken tray table is an example of an NSRE.

These tasks are further divided into mandatory and deferrable tasks. This division is based on the type of task and the harshness of its due date. Table 2 provides an overview of the task division.

4.2. Constraints

On the day prior to operations, the framework will perform the tail assignment and maintenance task scheduling, complying with the following constraints consequently defined.

Tail assignment: Involves assigning aircraft to rotations. It is constrained by airline-specific rules, aircraft airworthiness constraints, and aircraft balance. Each airline has rules restricting certain registrations from flying to certain destinations and/or limiting the number of daily quick turnarounds. Airworthiness constraints connect the tail assignment and task scheduling problems. These ensure that an aircraft is assigned to a flight only if it has no outstanding maintenance task expiring during the flight. Flow balance must be guaranteed as tails cannot be assigned to two overlapping flights. Lastly, it is assumed that the framework's tail assignment is not restricted by decisions taken the previous day. Every day, it can create a tail assignment from scratch.

Maintenance task scheduling: Involves assigning aircraft and their maintenance tasks to maintenance slots. The process of scheduling maintenance tasks is constrained by the 4M requirements. These include scheduling a task before its due date, only when material becomes available, and in a sufficiently long maintenance slot at the right location (i.e., hangar or platform). The last 4M requirement, which regards assigning the correct mechanics skills required to perform a task, is not modelled in the framework. Hence it is assumed that the right skills to perform tasks are always available. As only relatively simple maintenance tasks, which require skills possessed by most mechanics, are scheduled by the framework, this assumption is non-limiting. In addition to 4M requirements, task scheduling is constrained by rules that limit the assignment of aircraft to maintenance slots. As airline mechanics are often certified for one aircraft type, maintenance schedules are created per aircraft type (i.e., Boeing 787). Thus it is assumed that maintenance slots are not interchangeable across aircraft types. Moreover, like the tail assignment, the framework does not have to respect the task assignment planned the previous day. Hence, it can create a new assignment each day.

4.3 Maintenance slots

The framework creates its maintenance slots based on aircraft-specific needs. This has advantages over current state-of-the-art approaches which make use of generic fixed times, fixed duration maintenance slots from the airline's maintenance schedule, often suitable for the entire type fleet. For instance, the framework creates maintenance slots specific to one registration, whose duration is based on the aircraft's individual maintenance needs. It is hypothesized that this will result in higher labour utilization, less maintenance time, greater fleet availability, and hence improved schedule efficiency. Nonetheless, this methodology increases complexity of the model as it introduces additional decisions. Whereas in practice, it is based on the assumption that towing services are available at any time, and it requires a revolution in the current maintenance slot planning approach.

Maintenance slots in the framework are created using a two-step process. First, on a given day, the model determines aircraft-specific required maintenance time based on their task backlogs and a days-clean target, i.e., the target aircraft health after a maintenance slot. Subsequently, using the calculated maintenance time duration, the framework creates maintenance slots that can only be assigned to the corresponding registration. It ensures that slots are created within the confines of the existing maintenance slot schedule to respect the airline's resource limitations.

On rare occasions, to prevent infeasibility, off-hour maintenance slots are created. These do not abide by the airline's maintenance slot schedule, and thus in practice would require quick alterations to mechanics' work shifts. These slots are created when the aircraft has

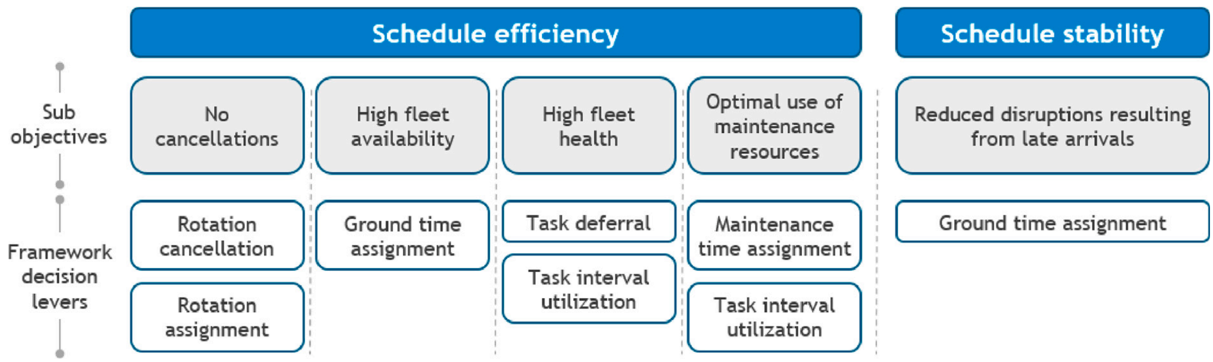

Fig. 3. Breakdown of the framework's objectives into sub-objectives and decision levers.

tasks that are due the next day but there are no suitable maintenance opportunities in the schedule. Otherwise, also when the aircraft's required maintenance time is longer than the longest slot in the schedule. Maintenance planners also create off-hour slots when similar issues arise.

4.4 Main objectives

As previously mentioned, the unified flight and maintenance scheduling framework developed in this study aims to improve schedule efficiency and stability. However, some of the framework's sub-objectives need to be further defined, as shown in Fig. 3. Schedule efficiency induces lower operating costs for an airline. It is characterized by no cancellations, high fleet availability, high fleet health, and optimal use of maintenance resources. On the other hand, improved schedule stability lowers disruptions recovery costs, reduces workload on the day of operations, and leads to indirect revenue gains from more satisfied passengers. In turn, schedule stability is defined by a reduction in aircraft's subsequent jobs starting late due to the late arrival of their previous flight.

In the framework, six unique decision levers enable the achievement of the sub-objectives. These levers will make up the objective function of the framework. Ground time, apart from time on the ground for maintenance, is divided into fleet availability and ground time waste. Fleet availability provides flexibility to the airline in the form of potential aircraft-swapping opportunities. On the contrary, ground-time waste, which is ground time before maintenance slots, offers no flexibility. This is because aircraft often cannot be swapped before planned maintenance. On a given day, the fleet health equals the number of days before the first maintenance task of an aircraft goes due, averaged over the entire fleet. When fleet health is high, most of the fleet does not require maintenance in the next days, thereby offering greater flexibility to the network. Lastly, an optimal use of resources is characterized by improved task interval utilization, reduced maintenance time, and higher labour utilization. This ought to reduce airline maintenance costs.

5 Methodology

This section defines the methods employed to build the unified tail assignment and maintenance tasks scheduling framework previously described in Section 4. Sections 5.1 and 5.4 describe the time-space networks (TSN) used to optimize scheduling and the LP which assigns tails and schedules maintenance tasks, respectively.

5.1 Time-space network

To assign tails and schedule maintenance tasks simultaneously, the framework uses parallel TSN. This is a modelling framework that allows for the direction representation of movement of available resources over time and locations. From this movement, the availability of resources can be derived. As a result, time windows, location transitions, and resource availability do not have to be defined explicitly. Additionally, the TSN construction enables the optimization of schedule efficiency and stability. Advantages of parallel TSNs include individual aircraft routing and ground time assignment. Parallel TSNs have been commonly used in the literature for tail assignment, such as demonstrated by Vink et al. [11]. However, the novelty stands in solving the maintenance task assignment within the TSN.

Fig. 4 depicts a simplified example of an aircraft's parallel TSN employed by the framework, whereby each aircraft has a separate network with independent ground arcs and maintenance slots. Coverage constraints for rotations are used to connect all networks. To decrease the network's size, time is discretized in non-homogeneous time steps. These steps depend on the start and end times of the aircraft's potential jobs (rotations, maintenance slots, and fixed maintenance slots). A Job's duration is rounded to the next nearest minimum time step defined by the airline's network and maintenance schedules. Moreover, rotation arcs are extended to accommodate for a fixed minimum turnaround time (TAT) before departure and after arrival. Maintenance slot arcs are also elongated with the location-dependent fixed tow time before the start and after the end of the slot.

5.2 Schedule efficiency

Part of the TSN's novelty lies in its utilization of two distinct spaces: the air space and the maintenance space. The assignment of rotations takes place at the air space, while scheduling maintenance tasks and maintenance slots is confined to the maintenance space. This dual-space approach serves the purpose of distinguishing between fleet availability and ground time waste. All air space ground arcs account for fleet availability, whereas all those of the maintenance space contribute to ground time waste.

Using connection arcs, aircraft can travel between the two spaces. Aircraft can travel to the maintenance space at their source node or only directly after the rotation they just flew. While aircraft can connect back to the air space at the end of all maintenance slots, including fixed maintenance slots. A side effect of the TSN formulation is that it results in a trade-off between schedule efficiency and schedule stability. Reducing ground time waste before a maintenance slot increases schedule efficiency. However, it also degrades schedule stability as the chances of the previous rotation's arrival delay propagating into the slot increase. Next, it is explained how schedule stability is optimized in the TSN.

5.3 Schedule stability

Stability is incorporated by using additional, non-mandatory ground arcs in the TSN, referred to as robust ground arcs. These are based only on rotations' historical arrival delays, excluding other sources of

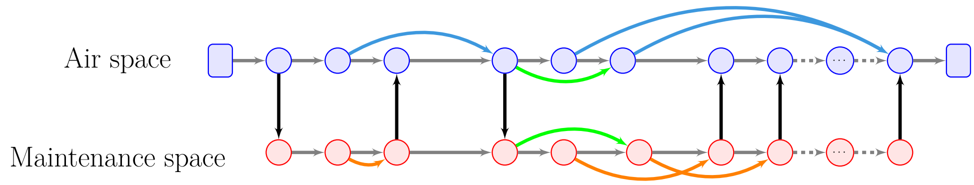

Fig. 4. Parallel time-space network used in the framework, representing the network for an aircraft. Rectangular nodes represent the aircraft's source and sink nodes, thus the start and end of the aircraft's path. The nodes in the *Air Space* represent the possible landing and departure moments. The nodes in the *Maintenance Space* represent maintenance slots and fixed maintenance activities. In light blue, a single rotation is defined, representing the two legs of the rotation. The departure of the second leg can be delayed to defend against potential arrival delays — the corresponding robust ground arc is represented in green. The aircraft may enter and exit maintenance opportunities — this path is represented by the connection arcs in black.

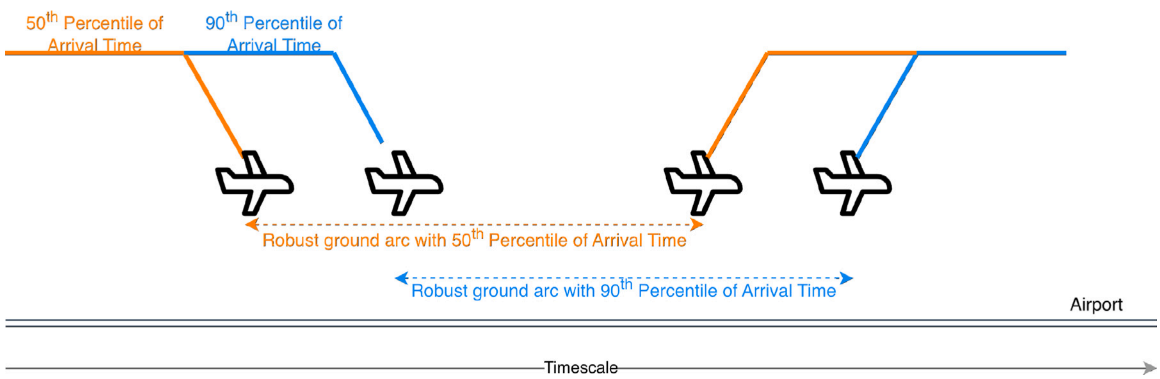

Fig. 5. Example of different percentiles of historical arrival delays to construct robust ground arcs. Selecting a higher percentile shifts the planned departure time of the next leg to a later time, as it accounts for larger historical delays and thus provides greater protection against disruption propagation.

disruptions. The objective of robust ground arcs is to reduce the chances of rotations' arrival delay from propagating into the aircraft's next jobs. However, delay mitigation should not compromise the framework's other objectives. This approach enables to simultaneously optimize schedule stability and efficiency, thereby enabling a trade-off between the two.

Every rotation has an air space and a maintenance space robust ground arc. Each starts at the node corresponding to the rotation's end time at the respective space. The duration of both arcs is set using a delay mitigation parameter that represents the rotation's Nth percentile arrival delay determined from historical airline delay data. For instance, a delay mitigation parameter of 90% implies that 90% of the rotation's historical delays are mitigated if the robust ground arc were to be used. Fig. 5 displays the difference between using the 50th and the 90th percentiles. Utilizing the 90th percentile results in the departure of the second leg of the rotation to planed for later. This mitigates larger values of arrival delay.

A trade-off analysis to determine the delay mitigation parameter will follow in Section 7. Lastly, a rotation's robust ground arc can only be assigned to the same aircraft assigned to fly the rotation. This is enforced by constraints in the LP model.

This approach allows the framework to make strategic decisions regarding the assignment of buffer time, based on historical delays, maintenance needs, and network assignment completion. It was chosen over other techniques due to its implementation simplicity and flexibility: the airline can directly control how conservative the model is by selecting the Nth percentile of arrival times. This parameter can be quickly adjusted, allowing the airline to compare multiple output scenarios from the framework. Responsibility is then shifted to the network planning team to determine how conservative the schedule should be. However, it is worth noting that alternative approaches could be explored to automatically select the most appropriate percentile based on the network schedule and the costs of cancellations [37].

5.4 *Linear programming optimization model formulation*

Core to the modelling framework, is the LP optimization model that simultaneously assigns tails and schedules maintenance tasks to optimize schedule efficiency and stability. LP was selected due to its explainability and optimality. However, scalability of the model is an issue for binary linear problems, which is the case of our framework (see Table 5). The computation time of the framework and its compatibility as a decision support tool will be discussed in Section 9.

The following sections present the characteristics of the LP model. Section 5.4.1 introduces the LP model's sets, parameters, and decision variables. The objective function is defined in Section 5.4.2. Finally, the constraints are discussed in Section 5.4.3.

5.4.1 *Notations*

The LP model's notation is introduced first. First, definitions for all the model's indices and sets can be found in Table 3. Then, the model's parameters are presented in Table 4, and can have either values that are fixed or dependent on the properties of model sets. Lastly, the model's decision variables are found in Table 5.

Tail assignment is characterized by two decision variable types, one to assign tails to arcs and the other to cancel rotations. Similarly, maintenance task scheduling is possible through two variables, one to assign tasks to maintenance slots and the other to defer tasks.

5.4.2 *Objective function*

The objective function is formulated as a cost-minimization weighted sum objective as shown in Eq. (1). It takes inspiration from common airline objectives preferences. An explanation of the formulation of weight functions will follow in Section 7. The LP mathematical formulation of the objective function is given by Fig. 3. It is composed

Table 3
Definitions of the LP model's sets and indices.

Set	Definition
$s \in S$	Set of all TSN arcs s (rotations, maintenance slots, ground arcs, etcetera)
$f \in F \subset S$	Set of rotations f
$m \in M \subset S$	Set of maintenance slots m
$u \in U \subset S$	Set of fixed maintenance slots u
$g \in G \subset S$	Set of ground arcs g
$b \in B \subset G$	Set of air space ground arcs (fleet availability) b
$w \in W \subset G$	Set of maintenance space ground arcs (ground time waste) w
$c \in C \subset S$	Set of connection arcs c
$d \in D$	Set of days d in the planning window
$pd \in PD_d \subset D$	Set of days preceding day d including day d
$s \in S_d \subset S$	Set of arcs s on day d
$k \in K$	Set of aircraft types k (e.g. Boeing 777)
$a \in A$	Set of aircraft a
$a \in A_s \subset A$	Set of aircraft a that can be assigned to arc s
$a \in A_k \subset A$	Set of aircraft a of type k
$s \in S_a \subset S$	Set of arcs s that can be assigned to aircraft a
$t \in T$	Set of tasks t open at the start of the planning time
$t \in T_{Due} \subset T$	Set of open mandatory tasks t due on the first day of the plan (i.e., day of operation)
$t \in T_a \subset T$	Set of open tasks t of aircraft a
$t \in T_m \subset T$	Set of open tasks t that can be scheduled in maintenance slot m
$t \in T_d \subset T$	Set of open mandatory tasks t going due on or after day d
$n \in N$	Set of nodes n
$n \in MN \subset N$	Set of maintenance space nodes n
$n \in AN_a \subset N$	Set of air space nodes n with a terminating air space ground arc g and an originating connection arc c for aircraft a
$l \in L$	Set of maintenance locations l (hangar and platform)
$m \in M_t \subset M$	Set of maintenance slots m in which task t can be scheduled
$m \in M_{d,l} \subset M$	Set of non-off-hours maintenance slots m at maintenance location l on day d
$m \in M_{n,l,k} \subset M$	Set of non-off-hours maintenance slots m suitable for aircraft subtype k at maintenance location l passing through maintenance node n
$m \in M_a \subset M$	Set of maintenance slots m that can be assigned to aircraft a
$m \in G_a \subset G$	Set of ground arcs g that can be assigned to aircraft a
$g \in RG_{a,r} \subset G$	Set of robust ground arcs g of rotation f that can be assigned to aircraft a
$s \in O_{n,a} \subset S$	Set of arcs s originating from node n that can be assigned to aircraft a
$s \in T_{n,a} \subset S$	Set of arcs s terminating at node n that can be assigned to aircraft a
$f \in QTA_d$	Set of quick turnaround rotations on day d

Table 4
Definitions of the LP model's parameters.

Parameter	Unit	Definition
f_{QTA}	[–]	Quick turn around copy of rotation f
$FB_{n,a}$	[–]	Flow balance at node n for aircraft a: 1 at source nodes, −1 at sink nodes, and 0 at intermediate nodes
FTE_m	[/h]	FTE required for maintenance slot m
$FTE_{n,l,k}$	[/h]	FTE available at node n at location l for aircraft type k
$MaxM_{d,l,k}$	[–]	Maximum number of slots at location l and suitable for aircraft type k that can be assigned on day d
$MaxQTA$	[–]	Maximum number of daily quick turnaround rotations
L_t	[hr]	Labour hours required to complete task t
$L_{max,m}$	[hr]	Maximum labour hours that can be assigned in maintenance slot m
$BigM$	[–]	A relatively large number
$c_{a,n}$	[–]	Connection arc c originating from air space node n for aircraft a
$b_{a,n}$	[–]	Air space ground arc b terminating at air space node n for aircraft a
$DL_{t,d}$	[days]	Number of days remaining from day d until task t is due
H_f	[days]	Minimum required aircraft health to operate rotation f (e.g., for a rotation leaving today and returning tomorrow $H_f = 1$)
$C_{Type,t}$	[–]	Task type criticality coefficient of task t
$W_{a,f}$	[–]	Cost of assigning aircraft a to rotation f
$W_{a,m}$	[–]	Cost of assigning aircraft a to maintenance slot m
$W_{a,g}$	[–]	Cost of assigning aircraft a to ground arc g
$W_{Canx,f}$	[–]	Cost of cancelling rotation f
$W_{t,m}$	[–]	Cost of assigning task t to maintenance slot m
$W_{Defer,t}$	[–]	Cost of deferring task t

Table 5
Definitions of the LP model's decision variables.

Decision variable	Definition
$\delta_{a,s}$	Binary variable equal to 1 if aircraft a is assigned to arc s, 0 otherwise
$\delta_{Canx,f}$	Binary variable equal to 1 if rotation f is cancelled, 0 otherwise
$\delta_{t,m}$	Binary variable equal to 1 if task t is assigned to maintenance slot m, 0 otherwise
$\delta_{Defer,t}$	Binary variable equal to 1 if task t is deferred, 0 otherwise

of six terms, one for each decision lever presented in Section 4.

$$
\begin{aligned}
Min:\ & \overbrace{\sum_{f\in F} \delta_{Canx,f} W_{Canx,f}}^{\text{Rotation Cancellation}} + \overbrace{\sum_{f\in F}\sum_{a\in A_f} \delta_{a,f} W_{a,f}}^{\text{Rotation Assignment}} \\
& + \overbrace{\sum_{t\in T\neg T_{Due}} \delta_{Defer,t} W_{Defer,t} C_{Type,t}}^{\text{Maintenance Task Deferral}} + \\
& \underbrace{\sum_{a\in A}\sum_{m\in M_a} \delta_{a,m} W_{a,m}}_{\text{Maintenance Time Assignment}} + \underbrace{\sum_{m\in M}\sum_{t\in T_m} \delta_{t,m} W_{t,m} C_{Type,t}}_{\text{Maintenance Task Interval}} + \underbrace{\sum_{a\in A}\sum_{g\in G_a} \delta_{a,g} W_{a,g}}_{\text{Ground Time Assignment}}
\end{aligned} \tag{1}
$$

In more detail, the objective function terms are:

- *Rotation Cancellation*: Network planners can decide to cancel a rotation if it cannot be operated or to safeguard the future schedule. However, this induces high costs to the airline and is often a remedy of last resort. This severity is represented in the model with the highest weight, $W_{Canx,f}$.
- *Rotation Assignment*: In practice, and in the LP model, this decision depends on the aircraft's fuel efficiency and the rotation's block time.
- *Maintenance Task Deferral*: Maintenance resources are limited and/or performing maintenance may mean cancelling a flight. Therefore sometimes the best choice is to defer a task. This decision depends on the task hierarchy, implemented with $C_{Type,t}$ as done by Van Kessel et al. [17], and on the interval utilization (i.e., the moment in time a task is executed relative to its due date). Regarding the latter, the following two distinctions are made:
  – For mandatory tasks, the deferring cost increases close to the due date, as diminishing scheduling opportunities lead to a higher risk of grounding the aircraft.
  – For deferrable tasks, the deferring cost is constant, as there is no risk of grounding the aircraft as the task can be postponed.
- *Maintenance Time Assignment*: The assignment of maintenance slots is penalized with $W_{a,m}$. The cost of maintenance depends on the slot's duration. Additionally, extra costs are incurred for off-hours maintenance slots, as these require changes to mechanics rosters.
- *Maintenance Task Interval Utilization*: Maintenance tasks scheduling should optimize maintenance resources, specifically interval utilization and fleet health. Maintenance planners can schedule a task in multiple maintenance slots. Choosing the most suitable slot is a decision that depends on the task hierarchy and interval utilization. Hierarchy between tasks is achieved again with $C_{Type,t}$, while, when it comes to interval utilization, the following two distinctions are made:

– Preventive tasks should be scheduled as close as possible to their due date to minimize interval waste.
– Corrective tasks should be scheduled as soon as possible to not hinder future maintenance plans.

• *Ground Time Assignment*: Fleet availability is maximized, while ground time waste is minimized through $W_{a,g}$. The importance given to fleet availability as opposed to ground time waste is airline-specific. Moreover, using $W_{a,g}$, robust ground arcs are assigned with the intent of improving schedule stability.

5.4.3 Constraints

The constraints of the LP model are formulated as follows. Constraints (2) is a coverage constraint that ensures that either rotation f or its quick turnaround copy f_{QTA} are assigned to at most one aircraft or cancelled. No other coverage constraints are required because all other arc types are aircraft-specific, i.e., can be assigned only to one tail. Constraints (3) ensure that all fixed maintenance slots are assigned. Flow balance is guaranteed with constraints (4) at every node and for every aircraft.

$$\sum_{a\in A_f}\delta_{a,f}+\sum_{a\in A_{f_{QTA}}}\delta_{a,f_{QTA}}+\delta_{Canx,f}=1,\quad \forall f\in F \tag{2}$$

$$\sum_{a\in A_u}\delta_{a,u}=1,\quad \forall u\in U \tag{3}$$

$$\sum_{s\in O_{n,a}}\delta_{a,s}-\sum_{s\in T_{n,a}}\delta_{a,s}=FB_{n,a},\quad \forall n\in N,a\in A \tag{4}$$

Using constraints (5) the model guarantees that a task is scheduled to a maintenance slot only if the corresponding aircraft is assigned to that same slot. Constraints (6) prohibit empty slots from being assigned, while constraints (7) prohibit the number of scheduled labour hours in a slot from exceeding that slot's maximum allowed. Tasks with due dates not on the first day of the plan should be scheduled in a maintenance slot or deferred, while tasks going due have to be scheduled. The model enforces the former with constraints (8) and the latter with constraints (9).

$$\sum_{t\in T_m\cap T_a}\delta_{t,m}\le Big_M\delta_{a,m},\quad \forall a\in A,m\in M_a \tag{5}$$

$$\sum_{a\in A_m}\delta_{a,m}\le\sum_{t\in T_m}\delta_{t,m},\quad \forall m\in M \tag{6}$$

$$\sum_{t\in T_m}L_t\delta_{t,m}\le L_{max,m},\quad \forall m\in M \tag{7}$$

$$\sum_{m\in M_t}\delta_{t,m}+\delta_{Defer,t}=1,\quad \forall t\in T:t\notin T_{Due} \tag{8}$$

$$\sum_{m\in M_t}\delta_{t,m}=1,\quad \forall t\in T_{Due} \tag{9}$$

The following two constraints ensure that the airline's maintenance schedule is respected. Constraints (10) ensure that the available FTE for an aircraft type at a given maintenance location and node is never exceeded. Constraints (11)guarantee that on any given day the number of assigned maintenance slots at a given maintenance location for a particular aircraft type do not exceed that maximum in the schedule.

$$\sum_{m\in M_{n,l,k}}\sum_{a\in A_m}FTE_m\delta_{a,m}\le FTE_{n,l,k},\quad \forall k\in K,n\in MN,l\in L \tag{10}$$

$$\sum_{a\in A_k}\sum_{m\in M_a\cap M_{d,l}}\delta_{a,m}\le MaxM_{d,l,k},\quad \forall k\in K,d\in D,l\in L \tag{11}$$

The next constraints, constraints (12), prohibit aircraft from being assigned to more than one maintenance slot per day. Constraints (13) ensures that the number of rotations assigned to each aircraft does not exceed the maximum number of daily quick turnarounds. Constraints (14) and constraints (15) prevent certain movements in the TSN. The former makes sure that aircraft can travel to the maintenance space only directly after the end of rotations they were assigned to, meaning that an aircraft can only go into maintenance after completing all the flights to which it is assigned. The latter ensures aircraft continuity, i.e. the robust ground arc of a rotation can only be assigned to the aircraft that is assigned to the rotation. Note that, as previously mentioned in Section 5.3, robust ground arcs are based rotations' historical arrival delays. The objective is to allocate enough ground time to compensate for commonly seen delays.

$$\sum_{m\in M_a\cap M_d}\delta_{a,m}\le 1,\quad \forall d\in D,a\in A \tag{12}$$

$$\sum_{f\in QTA_d}\sum_{a\in A_f}\delta_{a,f}\le MaxQTA,\quad \forall d\in D \tag{13}$$

$$\delta_{a,c_{n,a}}+\delta_{a,b_{n,a}}\le 1,\quad \forall a\in A,n\in AN_a \tag{14}$$

$$\delta_{a,g}\le\delta_{a,f},\quad \forall f\in F,a\in A_f,g\in RG_{a,f} \tag{15}$$

Lastly, aircraft airworthiness is guaranteed with constraints (16). For every day and every tail, it checks that the days left in the intervals of all the tail's open mandatory tasks going due on or after the day in question are higher than the aircraft health required to operate a given rotation. Each task that does not satisfy this requirement should be scheduled in a maintenance slot before the rotation, otherwise, that rotation cannot be operated by the tail in question. Constraints (17)–(20) define the decision variables as binary. Finally, the LP model is solved using a commercial tool.

$$\frac{1}{|M_t\cap M_{pd}\forall pd\in PD_d|}\sum_{pd\in PD_d}\sum_{m\in M_t\cap M_{pd}}(1+BigM\delta_{t,m})DL_{t,d}>H_f\delta_{a,f},$$
$$\forall d\in D,a\in A,t\in T_a\cap T_d,f\in F_a\cap F_d \tag{16}$$

$$\delta_{a,s}\in\{0,1\},\quad \forall s\in S,a\in A_s \tag{17}$$

$$\delta_{Canx,f}\in\{0,1\},\quad \forall f\in F \tag{18}$$

$$\delta_{t,m}\in\{0,1\},\quad \forall m\in M,t\in T_m \tag{19}$$

$$\delta_{Defer,t}\in\{0,1\},\quad \forall t\in T:t\notin T_{Due} \tag{20}$$

6 Hypotheses

The following hypotheses are defined regarding the performance of the unified tail assignment and maintenance tasks scheduling framework developed in this work:

Hypothesis 1: Instead of using aircraft generic fixed duration maintenance slots, the unified framework creates tail-specific slots, tailored to aircraft maintenance needs. This is expected to reduce total maintenance time and, in turn, increase fleet availability.

Hypothesis 2: Rather than allocating tails and scheduling maintenance jobs separately, the framework uses its unique TSN formulation to do both simultaneously. It is expected that this will maximize fleet availability over ground time waste, hence optimizing the ground time assignment.

Hypothesis 3: The framework's computation capabilities are expected to improve the utilization of maintenance resources, specifically task interval and labour utilization, compared to the airline's manual approach.

Hypothesis 4: The LP model provides consistency in the decision-making process for assigning maintenance tasks and tails, which would otherwise be unattainable in two large planning teams with conflicting goals. This is expected to reduce variance in planning performance, especially regarding planned fleet health.

Hypothesis 5: Delay robust routing using rotation's historical arrival delay data, as integrated into the framework's TSN, will decrease unexpected disruptions, resulting in easier-to-recover plans.

These hypotheses will be validated with the results shown in the following Section 7, and will be further discussed in Section 9.

Table 6
Overview of the dataset use to evaluate the framework developed in this study.

Planning date	Rotations	Open tasks	Fixed maintenance slots
31/10/2023	77	379	9
01/11/2023	90	379	8
02/11/2023	92	374	9
03/11/2023	93	344	12
04/11/2023	126	345	13
06/11/2023	94	363	12
07/11/2023	82	387	9

7 Case study

A case study is performed to understand whether the framework can provide realistic and real-time decision support to airline planners the day before operations. Moreover, by comparing the framework's plans with those of the airline, it is assessed if the model can contribute to improved schedule efficiency and stability. Although planning is performed manually by the airline, the framework models the same method currently adopted in practice, thereby it can be a solid benchmark for comparison. This section follows the subsequent structure. The training data is defined in Section 7.1. Section 7.2 determines the model parameters with the help of trade-offs.

7.1 Test procedure

The objective of the case study is to assess if the framework can create realistic plans in real time that improve schedule efficiency and stability. The evaluation focuses on plans crafted for seven consecutive planning dates, spanning from October 31st, 2023 to November 6th, 2023. For each day within this timeframe, the framework generates stand-alone tail and task assignments for a predetermined number of subsequent days. All models employ identical input information as the airline's planners have available at 5 PM of each day, as shown in Table 6. To ensure a fair comparison, only rotations planned by the airline's planners are scheduled, using the same tail restriction rules. Note that each rotation corresponds to two flights, with the aircraft departing and returning to the hub. Additionally, maintenance is planned seven days ahead of the planned date considered, mirroring the practices of the airline's maintenance planners.

The case study airline is a major European single hub-to-spoke airline with a fleet of 54 wide-body aircraft of four subtypes and two types. The fleet assignment has been solved in advance and provided as an input. Therefore, type and subtype swaps are not allowed, meaning that rotations can only be assigned to the subtype fleet decided during the fleet assignment. Not all maintenance tasks are part of the case study. Short-interval turnaround maintenance is excluded and assumed to be scheduled during the aircraft's TAT, even if the aircraft's following job is a maintenance slot. Long-interval letter-check maintenance tasks are also excluded. These are planned in fixed maintenance slots, which are forcefully assigned by the LP model. Moreover, fixed maintenance slots are assumed to be full, hence prohibiting the assignment of any further maintenance task in these slots. These decisions align with practices at the case study airline. Lastly, the maintenance schedule is also provided by the airline.

7.2 Parameter definitions

As the case study involves comparing the airline's manual plan against that of the framework, parameters and weight functions are defined according to the airline's preferences. Section 7.2.1 explains the rules used to create a hierarchy between weights. Sections 7.2.2 to 7.2.4 outline the definition and fine-tuning process of the objective weights, which were formulated as cost functions. An overview of the weights is provided in Table 7.

Table 7
Overview of the framework parameters, defined according to the airline's preferences.

Parameter	Value
$MaxQTA$	2
$W_{Canx,f}$	10^7
W_{AOG}	100
C_{fuel}	1 /kg
C_{FTE}	100 /hour
P_{QTA}	2
$P_{off-hours\ FTE}$	2
H_{min}	3 days
Days-clean target, δ	10 days
Delay mitigation parameter	95%

7.2.1 Objective function hierarchy

To reflect the airline planners' priorities, the framework's objectives are characterized by four hierarchies or rules found in practice. These hierarchies help set the weights' orders of magnitude relative to each other. First, airlines aim to cancel as few rotations as possible. Thus cancelling a rotation is much more expensive than flying one. Second, deferring a task is more expensive than scheduling it in a maintenance slot, as continuous deferral of maintenance tasks can result in unplanned maintenance and/or too large task backlogs in the long run. The third rule reflects the trade-off between flying and performing maintenance. It states that cancelling a flight is worse than deferring tasks. Lastly, ground time is never prioritized over rotations, maintenance activity, and task interval utilization, and is the by-product of the planners' tail and task assignments. These rules as graphically shown in Fig. 6 and detailed in the following sections.

7.2.2 Network planning costs

Costs associated with network planning are the cost of operating a rotation, and the cost of cancelling a rotation. The cost of flying $W_{a,f}$ is based on the rotation's block time BT_f in hours, the aircraft fuel consumption FE_a in kg/hour, and the cost of fuel C_{fuel} in /kg, as follows:

$$W_{a,f} = BT_f FE_a C_{fuel}, \tag{21}$$

Where both the fuel cost and the aircraft-specific fuel consumptions are provided by the case study airline. Given the airline's data and rotations, flying a long-haul rotation typically costs around 10^5. All aircraft within a subtype fleet are assumed to have the same fuel efficiency, as planners ignore individual registration fuel efficiencies.

For quick TA rotations, $W_{a,f}$ is multiplied by a penalty, P_{QTA}. The case study airline allows for a maximum of two quick TAs per day by removing 60 min of turnaround time at the start of the rotation. The selected penalty value ensures that the additional cost of the rotation is significantly higher than the benefit of scheduling an extra 60 min of maintenance or fleet availability. Thus, quick TA is only used to prevent cancellations, as it is in practice.

Cancelling a rotation is the last resort decision taken by planers. To represent this in the framework, the cost of cancelling a rotation, $W_{Canx,f}$, equals 10^7 or around 100 times higher than the rotation's operating cost. $W_{Canx,f}$ is the same for all rotations, as the goal is to prevent cancellations, rather than optimally selecting which rotations to cancel.

$$W_{Canx,f} = 10^7 \tag{22}$$

7.2.3 Maintenance planning costs

Maintenance planning consists of three parts: task deferral, task scheduling, and maintenance time assignment. The cost of deferring a task, $W_{Defer,t}$, is directly linked to the cost of cancelling a flight, $W_{Canx,t}$. This relation dictates the magnitude of $W_{Defer,t}$. It is ensured

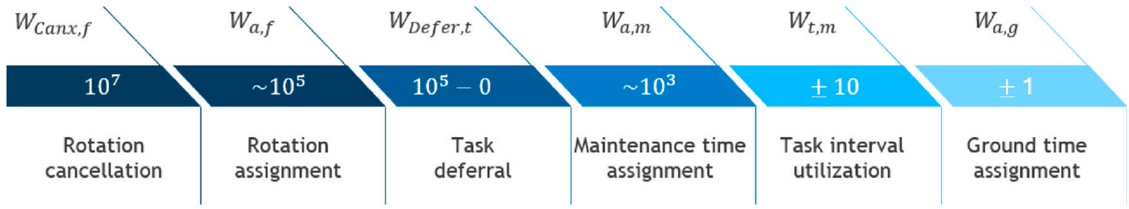

Fig. 6. Objective function hierarchy based on airline planning priorities (image inspired by the work of Van Kessel et al. [17]). The order of greatness of the different weights is shown in decreasing order of magnitude. The highest and lowest importances are given to *Rotation cancellation* and *Ground time assignment*, respectively.

that the framework will always prefer task deferral over flight cancellation. Only if operating the rotation would lead to more than 100 maintenance tasks being deferred does flight cancellation become the more viable option. However, this is a highly unlikely situation.

Moreover, deferring a task becomes more expensive as its due date approaches. This is because fewer opportunities to schedule it will arise and the risk of AOG increases. Therefore, $W_{Defer,t}$ is assumed to follow a linear function that increases with decreasing days left to the task's due date (DL) given by Eq. (23). When the days until the task's due date are more than the days-clean target (δ) the deferral cost is zero, to prevent early scheduling of preventive tasks. Afterwards, $W_{Defer,t}$ increases linearly, with the deferral cost drastically increasing when the remaining days are less than the minimum health target (H_{min}) to ensure that the task is scheduled.

$$W_{Defer,t} = \begin{cases} 10^5 - \frac{(10^5-10^4)DL}{H_{min}} & \text{if } DL \leq H_{min} \text{ and t is mandatory} \\ 10^4 - \frac{(10^4-10^3)(DL-H_{min})}{\delta - H_{min}} & \text{if } H_{min} < DL \leq \delta \text{ and t is mandatory} \\ 0 & \text{otherwise} \end{cases} \tag{23}$$

The cost of scheduling a task has two components: the cost of task interval utilization ($W_{t,m}$) and the cost of maintenance time ($W_{a,m}$). The formulation of the former is inspired by the work of Van Kessel et al. [17]. It depends on the task's type and due date (Due_t) and the maintenance slot's start date ($Start_m$). $W_{t,m}$ is defined by a linearly decreasing function for preventive tasks, while a linearly increasing function for corrective tasks. This is to optimize the task interval (I_t) utilization. Following hierarchical rules, weights range from 10 to 0 for preventive tasks and from −10 to 0 for corrective tasks. The weight functions are given, respectively, in Table 8 and Eq. (25). Moreover, maintenance planners prefer not to schedule tasks too close to their due date (Due_t). Thus, a penalty (W_{AOG}) is implemented when the days left in the task's interval are less than the minimum health target. Finally, Table 8 presents the values of $C_{Type,t}$, which define the hierarchy between task types. These are taken from the work of Van Kessel et al. [17] and are determined based on the impact on airworthiness of the type of task.

$$W_{Preventive\ t,m} = \begin{cases} 10\frac{Due_t - Start_m - H_{min}}{I_t - H_{min}} & \text{if } Start_m \geq H_{min} \\ W_{AOG}(1 - \frac{Due_t - Start_m}{H_{min}}) & \text{otherwise} \end{cases} \tag{24}$$

$$W_{Corrective\ t,m} = \begin{cases} -10\frac{Due_t - Start_m - H_{min}}{I_t - H_{min}} & \text{if } Start_m \geq H_{min} \\ W_{AOG}(1 - \frac{Due_t - Start_m}{H_{min}}) & \text{otherwise} \end{cases} \tag{25}$$

The cost of assigning a maintenance slot depends on the slot's scheduled duration, D_m, and FTE, FTE_m, as well as the cost of one FTE hour. This is the same in the framework, where $W_{a,m}$ is given by Eq. (26). The cost of one FTE hour C_{FTE} is a sum of labour, opportunity, and material costs. Off-hours maintenance slots are penalized by multiplying $W_{a,m}$ by $P_{off-hours\ FTE}$. The value of the penalty is chosen to ensure that task deferral remains more expensive.

Table 8
Task type weighting factor [17].

Task type	$C_{Type,t}$
Requirements	4
MEL	4
Adhoc	2
NSRE	1
Other	1

$$W_{a,m} = FTE_m D_m C_{FTE} \tag{26}$$

The choice of the days-clean target significantly influences task deferral costs, maintenance time, and, consequently, fleet availability. To justify the selection of a days-clean target that optimizes schedule efficiency, a trade-off analysis was conducted. The results are presented in Fig. 7. This trade-off analysis used a rolling horizon technique to simulate scenarios spanning five weeks (from July 31st, 2023 to September 4th, 2023) while varying the days-clean target. This approach allowed for an assessment of the long-term impact of scheduling with a specific target. As anticipated, the choice of the days-clean target affects schedule efficiency, as it influences fleet health, maintenance time, and fleet availability. A decrease in the days-clean target results in lower fleet health due to reduced aircraft health after a maintenance slot. In the long term, lower fleet health corresponds to an increase in weekly maintenance time, as aircraft require maintenance more frequently. This, in turn, translates into a decrease in overall fleet availability.

Conversely, higher fleet health in the long term results in more maintenance time. This is driven by fewer task deferrals, a higher number of scheduled tasks, and poorer task interval utilization, as illustrated in Fig. 8. This leads to a higher maintenance frequency necessary to achieve the demanding days-clean target. It is important to note that the utilization of corrective tasks' intervals remains largely unaffected. This is because they are typically scheduled well before the days to their the due date become fewer than the days-clean target. Thereby, the days-clean target has little effect on corrective task's interval utilization. Schedule efficiency is optimized when planning between 10 and 14 days clean based on fleet health, maintenance time, and fleet availability. For the case study, the days-clean target is set to 10 days, as this falls within the optimal range and aligns with the airline's days-clean target.

7.2.4 Ground time costs

Airline planners are indifferent to where and when ground time is assigned as they prioritize flying and maintenance. Consequently, to reflect its low priority, ground time is associated with the lowest weight ($W_{a,g}$). However, unlike airline planners, the model optimizes ground time allocation, by maximizing fleet availability and schedule stability. Thus, a distinction is made between air space ground arcs (ground

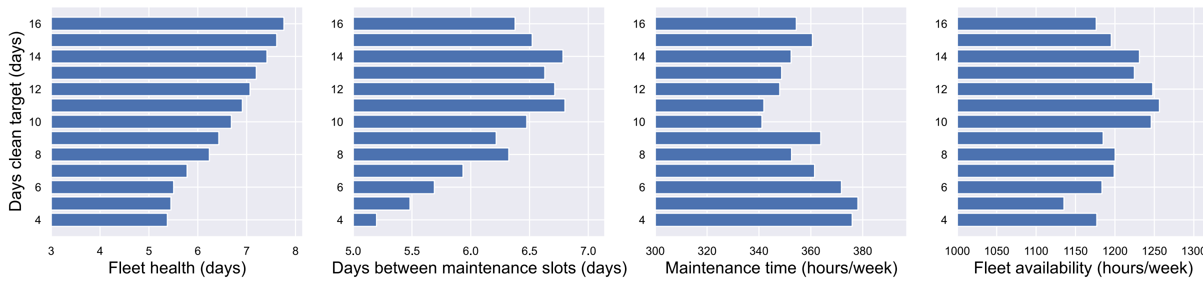

Fig. 7. Average fleet health, days between maintenance slots, weekly maintenance time, and weekly fleet availability for several days-clean targets.

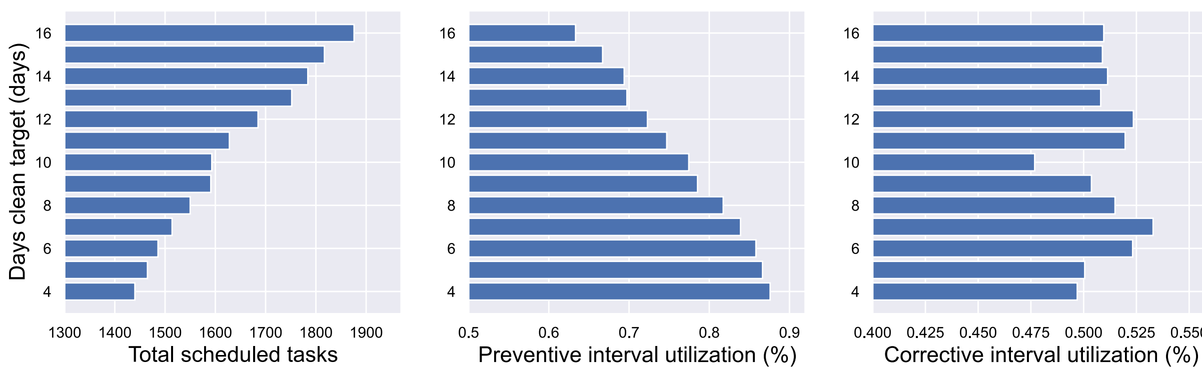

Fig. 8. Scheduled tasks, average preventive interval utilization, and average corrective interval utilization for several days-clean targets.

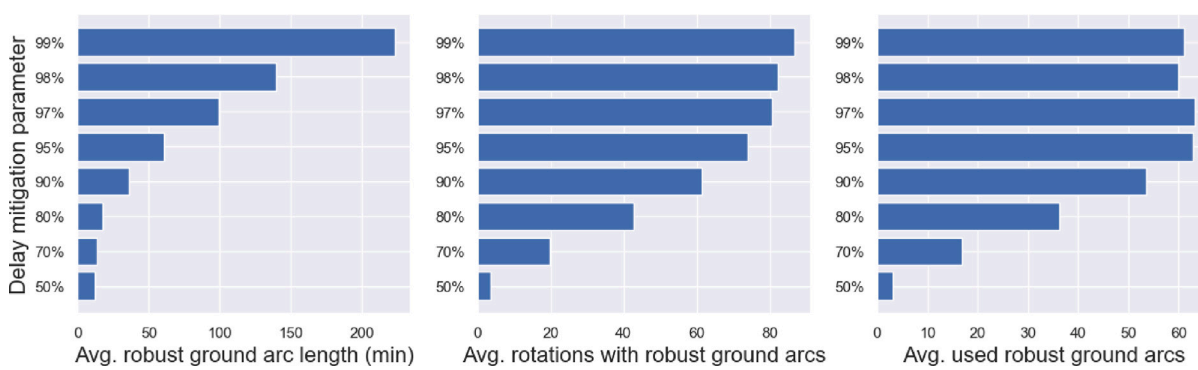

Fig. 9. Average robust ground arc length, number of rotations with robust ground arcs, and number of assigned robust ground arcs for several delay mitigation parameters.

time waste) and maintenance space ground arcs (fleet availability), respectively costing 0/h and 1/h of ground time.

The schedule stability is implemented by assigning robust ground arcs. These arcs, when assigned, can help reduce disruptions. Hence, they have a cost of −1/h of ground time, such that the framework actively seeks to assign them whenever possible. This cost formulation prioritizes scheduling stability over fleet availability. The delay mitigation parameter has a direct implication on schedule stability. It affects the length and number of robust ground arcs. As the delay mitigation parameter increases, robust ground arcs are expected to become longer and more. The trade-off results presented in Fig. 9 illustrate this dynamic. Nevertheless, there exists a threshold beyond which the arcs become excessively long, resulting in a reduction in their assignment. Ideally, to increase schedule stability, one should maximize the number of used robust ground arcs. Hence, the delay mitigation parameter is set to 95%. Note that this analysis is specific to the case study airline and its delay distributions. Hence, maximum stability is not always achieved with a 95% delay mitigation parameter.

8 Results

This section shows the final results with the application of the framework to this case study. Section 8.1 assesses schedule efficiency. Lastly, Section 8.2 evaluates schedule stability on aircraft late arrivals from late rotations through stochastically generated disruption scenarios. The unified framework created realistic plans, validated by the airline's planers. To benchmark the framework's performance, four different reference models are used:

- *Airline*: The schedule is produced manually by planners at a major European airline. The planners use the same inputs as the framework. However, they do not consider historical delays when defining the ground time between rotations. Commonly used turn around times are employed. Used as baseline comparison to test the efficiency and stability of the framework.
- *Fixed Tail Assignment (TA)*: The schedule produced by the LP model when scheduling maintenance tasks only and fixing the tail assignment to the one created by the airline's planners.

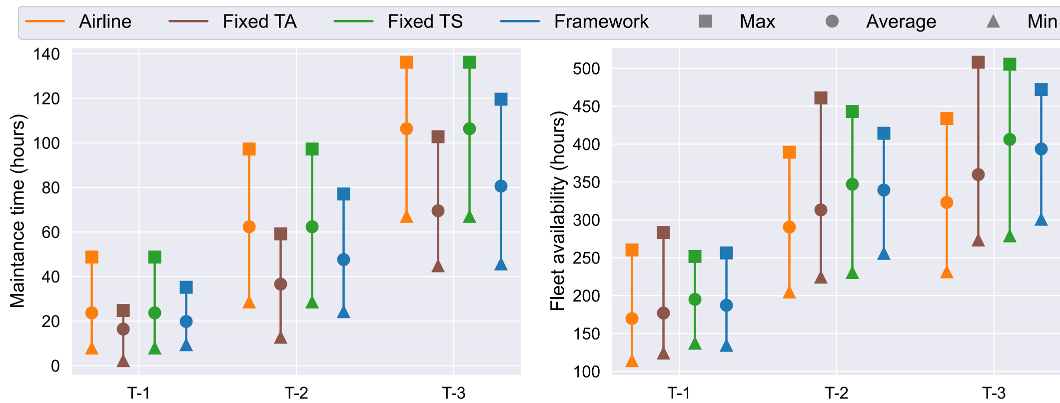

Fig. 10. Planned fleet availability and maintenance time for T-1, T-2, and T-3 for all four models.

- *Fixed Task Scheduling (TS)*: The schedule produced by the LP model when assigning tails only and fixing the task assignment to the one created by the airline's planners.
- *Framework*: the schedule provided by the LP model considering that both tail assignment and maintenance task scheduling are optimized together.

The purpose of the *TA* and *TS* models is to assess the benefits, if any, of integrating tail assignment and maintenance task scheduling. Although airline planners assign tails and schedule tasks in separate teams, they still occasionally communicate and collaborate. Thus, in this respect, the airline model is not the perfect benchmark. The Fixed TS and Fixed TA models are constructed by forcing decision variables based on the airline's decisions. Therefore, the LP model only solves the other part of the problem. For example, for the fixed TA model, the airline tail assignment is enforced by restricting decision variables, thereby only allowing the LP model to change the maintenance task scheduling. The following subsections cover the performance of the sub-objectives of the framework.

Table 9
Distribution of the ground time waste divided into short ((0, 2]), medium ((2, 12]), and long (>12) duration in hours for all four models.

Ground time waste	Short	Medium	Long
Airline	40.9%	36.4%	22.7%
Fixed TA	42.4%	31.8%	25.8%
Fixed TS	45.6%	35.3%	19.1%
Framework	39.5%	39.5%	21.0%

Table 10
Distribution of the ground time planned after rotations in hours for all four models.

Ground time planned (h)	(2, 8]	> 2
Airline	23.7%	45.2%
Fixed TA	21.8%	47.7%
Fixed TS	32.4%	54.9%
Framework	32.4%	55.9%

8.1 Schedule efficiency analysis

The framework increases fleet availability by reducing maintenance time. Fig. 10 summarizes the fleet availability and maintenance time planned in the subsequent one, two, and three days of a plan, respectively, denoted by T-1, T-2 and T-3. Thanks to the framework's aircraft-specific maintenance slots, the average maintenance labour utilization increases from the airline's 67% to 81%, reducing wasted maintenance time. Consequently, compared to the airline, the framework plans on average 15 additional hours of fleet availability on the day of operations, or an increase of 10%. Notably, the minimum fleet availability planned by the framework is significantly higher than that of the airline. Thus, the framework proves especially beneficial in dense plans, when fleet availability is limited, but highly desired. However, there are operational constraints, such as mechanics skills assignment and sudden labour and/or towing services shortages, that are outside the scope of the framework. These could increase maintenance time and consequently decrease fleet availability.

The framework's ability to integrate tail assignment and maintenance task scheduling helps reduce maintenance time and increase fleet availability. Fixing the tail assignment restricts maintenance task scheduling in the Fixed TA model, resulting in additional ground time waste. As a result, average fleet availability falls even when maintenance time is decreased. The Fixed TS model implements the same task and slot assignment as the airline, and hence plans the same amount of maintenance time. But compared to the airline, it significantly increases fleet availability. This highlights the advantages of employing the optimization model for the tail assignment. Additionally, compared to the framework, it manages to slightly increase average fleet availability. Given that the Fixed TS approach plans more maintenance time, this may appear counterintuitive. But because in the framework task interval utilization is given priority over ground time assignment, the framework wastes more ground time in order to optimize task interval utilization. Consequently, this decreases the availability of the fleet. Nonetheless, the framework's emphasis on task interval utilization will in the long term save maintenance time, increase fleet availability, and ultimately boost schedule efficiency.

The location of fleet availability within the schedule is just as crucial as its quantity. Table 9 displays the distribution of ground time waste divided into short, medium and long duration based on the airline's preferences. The framework increases medium duration ground time waste because it prioritizes robust ground arcs and task interval utilization over minimizing ground-time waste. This explains why the Fixed TS model, which reduced medium and long duration ground time waste, manages to increase fleet availability over the framework. Still, the framework reduced long duration ground time waste in comparison to the airline. This is particularly advantageous as long duration ground time waste represents the largest loss of flexibility. Moreover, Table 10 shows that the framework improves ground time allocation after rotations. The framework allocates ground time more evenly, represented by increased rotations followed by ground time between 2 and 8 h. Additionally, it increases rotations followed by more than 2 h of ground time, which is useful to mitigate arrival delays. Both are preferences of the airline's planners.

Although the model plans on average 17% less maintenance time than the airline, it schedules a similar amount of maintenance tasks

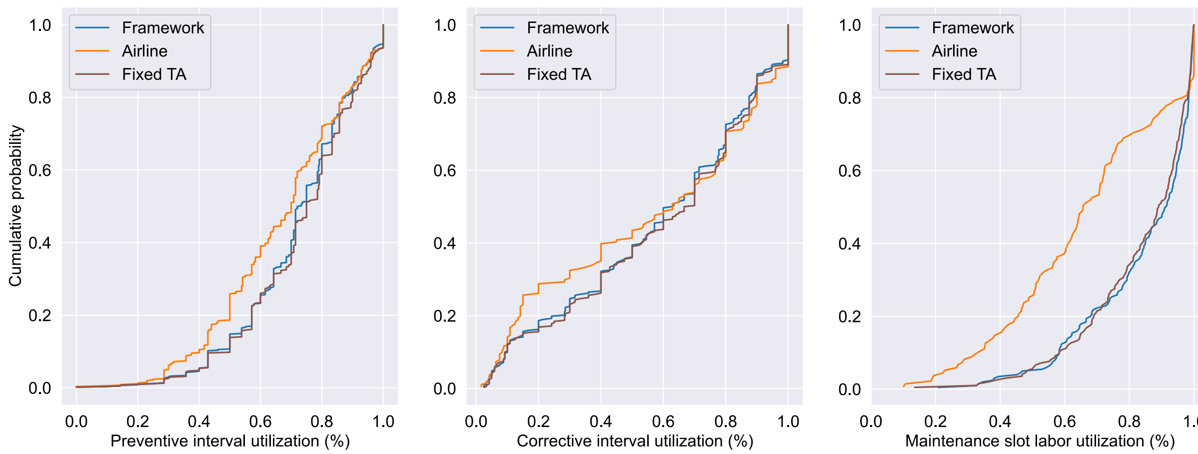

Fig. 11. Cumulative distributions of preventive maintenance tasks' interval utilization, corrective maintenance tasks' interval utilization, and maintenance slots' labour utilization. Note that results for the Fixed TS model are not shown, as they are exactly the same as the Airline model.

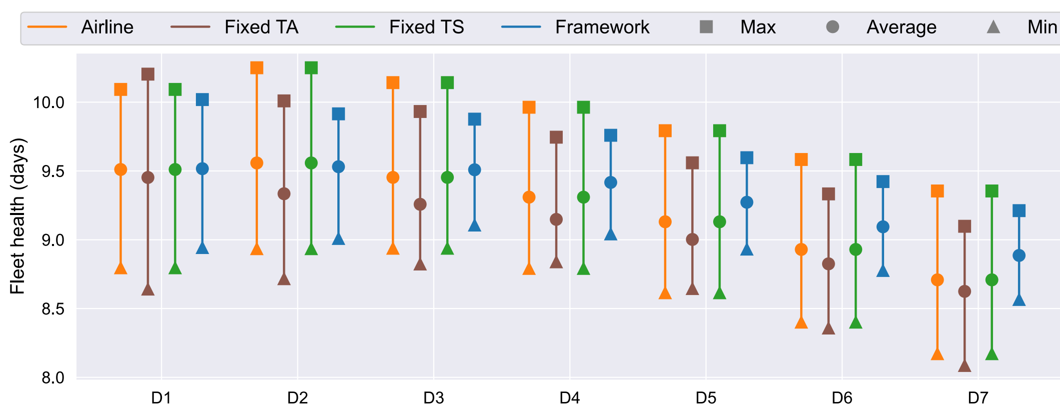

Fig. 12. Planned fleet health for the seven days from the planning date.

each planning day. This indicates that the framework is more efficient at scheduling maintenance tasks. Indeed, it improves slot labour utilization and schedules preventive tasks closer to their due date, improving their interval utilization, as shown in Fig. 11. However, no interval improvements are obtained for corrective tasks. This is likely due to their lower priority set with $C_{Type,t}$. The model's resulting improvements in the utilization of maintenance resources lower airline operating maintenance costs. However, the framework's higher labour utilization could harm work packages' completion rate, resulting in more spilt tasks and additional unplanned maintenance time and costs.

Fig. 12 displays the fleet health planned by each model for seven days from the planning date. The framework and the airline planers plan an almost identical average fleet health from D1 to D3. However, after D3 the framework plans more health. This becomes even more evident when comparing the framework to the fixed TA model. Fixing tail assignment greatly restricts maintenance task scheduling, deteriorating fleet health. A healthier fleet offers more flexibility to network operations as it means less maintenance time is needed in the subsequent days. Moreover, the decrease in the range between max and min planned fleet health, suggests that the framework is more consistent with the task assignment.

8.2 Schedule stability analysis

This subsection evaluates the impact of the framework on schedule stability. To dive into schedule stability, the plans of all four models are subject to aircraft late arrival disruptions. These disruptions are derived from historical rotation arrival delay data, including early arrivals, from August 2014 to November 2023 provided by the case study airline. The data set contains only instances of rotations inbound to the airline's hub. Rotations instances are grouped based on their departure airport and season (winter or summer), excluding flight numbers and/or flown aircraft subtypes as differentiators.

Each plan is tested with 100 different disruption scenarios. For each scenario, all planned rotations have a unique arrival delay randomly generated from their corresponding arrival delay distributions. On a given planning date, all four models are assessed with the same 100 disruption scenarios. To ensure a fair comparison, no maintenance disruptions are introduced, as the framework does not account for maintenance uncertainty while planning. The aim is to reduce disruptions arising from late arrivals, subsequently reducing the necessity for schedule adjustments. Additionally, minimizing propagated delay is a critical factor in schedule stability.

Fig. 13 summarizes the number of disrupted LoFs and propagated delay in T-1 and T-2 for all four models. A LoF is disrupted when the arrival delay extends into the aircraft's subsequent job, be it a rotation or a maintenance slot. On average, the framework creates plans that result in fewer disruptions and less propagated delay. Compared to the airline, average disrupted LoFs are reduced by 42% on the day of operations, while 43% in the subsequent two days. Importantly, the maximum number of disruptions is lower, thereby reducing planners' maximum recovery workload. Moreover, the average propagated delay is decreased by 48 and 88 min in T-1 and T-2 respectively. The Fixed TS model comes closer to the framework's performance as it still has significant freedom in the ground time allocation. On the other hand, fixing the tail assignment, done in the Fixed TA model, greatly

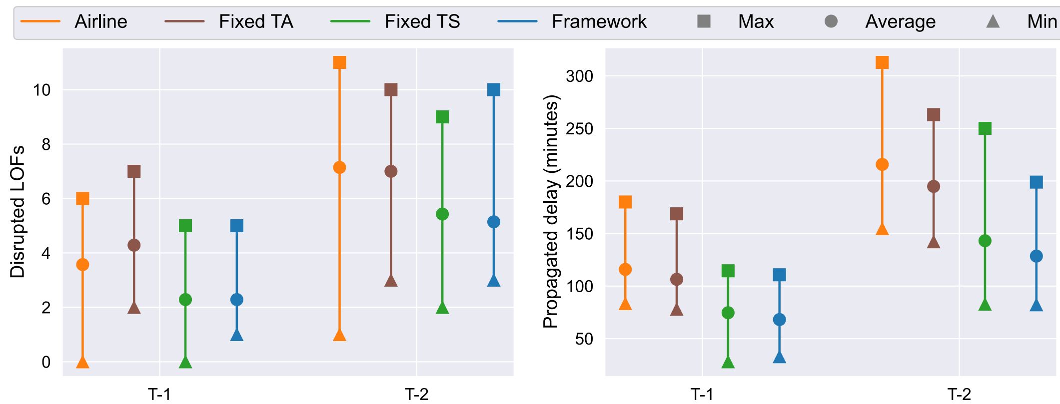

Fig. 13. Disrupted LOFs and propagated delay in the subsequent day (T-1) and subsequent two days (T-2) for all four models.

Table 11
Percentage of jobs that suffer delay in regard to the amount of delay in minutes.

	> 0	(0, 30]	(30, 60]	(60, 90]	> 90
Airline	8.96%	5.46%	1.40%	1.82%	0.28%
Fixed TA	7.56%	4.18%	1.75%	1.08%	0.54%
Fixed TS	6.42%	3.35%	1.40%	1.40%	0.28%
Framework	5.43%	2.38%	1.59%	1.19%	0.26%

Table 12
Distribution of the expected propagated delay (EPD) when disruptions occur in minutes.

	(0, 15]	(15, 30]	(30, 60]	> 60	avg. EPD
Airline	75.3%	16.9%	3.4%	4.9%	13.6 min
Fixed TA	75.3%	16.4%	2.7%	5.5%	13.9 min
Fixed TS	81.3%	14.6%	2.1%	2.1%	10.2 min
Framework	86.5%	13.5%	0%	0%	8.4 min

degrades stability. This is because most buffer time is pre-assigned. Lastly, although the framework's minimum number of disruptions is higher, the minimum propagated delay is lower. This hints at the model having shorter delays.

Table 11 presents the percentage of jobs that suffer delays in regard to the delay duration. The number of delayed subsequent jobs (i.e., job with a delay duration larger than 0 min) is reduced from the airline's 8.96% to less than 5.43%. Moreover, delays are reduced in each possible delay range, with the slight exception of delays between 30 and 60 min which increase from 1.40% to 1.59%. This is a result of some of the larger delay cases moving down to this bracket range. Especially beneficial is the reduction of delays longer than 60 min from 1.85% to 1.19%, which are the most detrimental and the hardest to recover. Thus, framework plans require fewer adjustments, due to fewer disruptions, and are easier to recover, due to shorter delays. This reduces recovery costs and has indirect revenue benefits from more satisfied passengers. Importantly, these advantages are not only obtained over the airline, but also over the Fixed TS and Fixed TA models. Hence, this justifies the benefits of integrating tail and task assignment from a literature perspective.

Table 12 presents a distribution of the expected propagated delay (EPD). EPD is not influenced by the disruptions scenarios — it is solely based on the arrival delay distribution of rotations and the scheduled ground time following them. The framework never suffers a EPD greater than 30 min, contrasting with the other three models. However, we consider that the reduction in the average EPD by 5.2 min obtained by the framework, 38% lower than the airline's, is not significant. Further testing of the model in a denser network with a higher potential for delay propagation is warranted.

Table 13
Average deferred new faults for the subsequent day (T-1) and subsequent two days (T-2).

	T-1 deferred faults	T-2 deferred faults
Airline	15	25
Fixed TA	15	25
Fixed TS	15	25
Framework	15	26

The previous results show that the framework improves network stability. However, these improvements must not compromise maintenance stability. Table 13 displays the number of new faults arriving in T-1 and T-2 that are deferred, i.e. not performed and postponed in time. These deferrals occur when the faults cannot be scheduled within the current plan. This may happen either because the aircraft was not assigned a maintenance slot before the fault's due date or because the assigned slot/slots is/are already full. Unlike network stability, the framework does not incorporate any forms of robustness in the task assignment, such as additional maintenance time buffers for unexpected non-routine work during the slot. Thus, the primary objective is not to outperform the other models. Rather, it is to achieve a similar number of deferrals. This is indeed the case, as average new fault deferrals are the same in T-1 and only one additional fault is deferred in T-2.

9 Discussion

The case study has shown that the framework improves the efficiency of the airline schedule. Firstly, like the airline, it does not cancel any rotation. Higher fleet availability increases useful buffer time for aircraft swapping, making the framework's plans easier to recover. Moreover, improved labour utilization reduces wasted maintenance time and wasted labour resources, reducing costs in the long term. Maintenance time, and thereby costs, are further reduced by the framework's higher fleet health and preventive task interval utilization. The efficiency gains of the framework are attributable to the integration of tail assignment and maintenance task scheduling.

Additionally, the schedule created by the framework increases stability by actively assigning ground time when possible. In case of disruptions, the framework uses this ground time to better mitigate delay compared to the airline's plans. As a result, a schedule produced by the framework would require fewer modifications in the event of disruptions.

Finally, when presented with the framework's plans alongside their own, without knowledge of which plan belonged to whom, the airline planners found it difficult to distinguish between the two. Furthermore, the average solution time of the framework is around 250 s for each day analysed. Arguably less than what an airline planner takes to assign tails and schedule maintenance tasks. Therefore, the model could be used in an operational environment.

The following subsections further drill into the validation of the proposed hypotheses in Section 6 and potential directions for improvement.

9.1 Hypotheses verification

Given the results for the presented use-case, the following can be stated regarding the hypotheses previously defined in Section 6:

Hypothesis 1: Thanks to maintenance slots tailored to each aircraft's specific needs, maintenance time waste is reduced, contributing to higher average labour utilization, specifically 81% compared to 67%. In turn, the framework reduced the total planned maintenance time on average by 17% and increased fleet availability, offering the airline more flexibility.

Hypothesis 2: Increased fleet availability (10% more on the day of operation), a decrease in long-duration ground time waste (22.7% down to 21.0%), and a more evenly distributed ground time after rotations (23.7% up to 32.4% of rotations followed by ground time between 2 and 8 h) are all indicators of the improved ground time allocation achieved by simultaneous tail assignment and maintenance task scheduling. Improved allocation of ground time provides airlines with increased flexibility and prospects for schedule recovery.

Hypothesis 3: The computational capabilities of the framework provide a more efficient use of maintenance resources than the manual approach used by the airline. Its better labour utilization allowed it to schedule an equivalent number of jobs in a shorter amount of maintenance time. It also increased the interval use of preventive tasks by scheduling them closer to their due date. This results in decreased maintenance operating costs for an airline. The interval utilization for corrective jobs was somewhat worse, though.

Hypothesis 4: Across the seven planning dates, the framework's plans showed less variation among each other considering planned maintenance time, fleet availability, and fleet health. The LP model has this benefit over a team of several planers with differing opinions, aims, and skill levels. As a result, the framework creates plans that enable more consistent operations under less dire circumstances, making operations simpler and more affordable.

Hypothesis 5: Schedule stability was enhanced by the framework. On the day of operations, it decreased the number of disruptions and the propagation delay by 42% and 48 min, respectively. It also significantly decreased the likelihood of long delays. As a result, the framework's plans not only require fewer modifications but are also simpler to recover, which lowers recovery costs and raises passenger satisfaction. This was attributable to simultaneous task and tail assignment (as opposed to the fixed TS model) and delay-robust routing (as opposed to the Airline and fixed TA models).

9.2 Future work

Although the framework's results are promising, the case study had limitations. The case study involved only the airline's long-haul single hub-to-spoke network. However, small adaptations can be made to the framework to accommodate short-haul operations and different types of networks, such as point-to-point. The latter would require a slight modification of the TSN by adding an additional air space per airport and an additional maintenance space per airport at which maintenance can be performed. Additionally, only short-term maintenance checks and tasks were considered. A valuable improvement would be to treat fixed maintenance slots as additional maintenance opportunities for task assignment, given that space permits it, and add minor flexibility to their start times.

Further improvements can be made to the framework itself. Currently, the duration of robust arcs that add buffer time between rotations is fixed, representing the Nth percentile arrival delay determined from historical airline delay data. Note that the airline may choose to run the framework for scenarios representing different values of this percentile. The results of all scenarios may be aggregated for decision-support. Another possibility is a stochastic programming approach, where a second state decision considers the outcome of several uncertainty realization scenarios. The final cost should then be minimized across these scenarios [38]. Finally, Monte Carlo simulations may be used to sample from historical delay data. The method used in this study reflects airlines' emphasis on simplicity of implementation. However, it is valuable to compare these approaches with more advanced techniques capable of accounting for multiple uncertainty scenarios.

Additionally, the research focused exclusively on arrival delays, neglecting other potential sources of uncertainty, such as unexpected labour within maintenance slots and AOG situations. The former could be especially problematic, as the model leads to densely packed slots, increasing the risk of delayed slots. Consequently, there is a need to incorporate robustness against late maintenance slots, potentially through the use of similar robust ground arcs based on historical task duration data. This would require the addition of an input element defining the likelihood of labour shortages. Once again, the impact of the stochasticity on this element can be evaluated by simulating multiple scenarios with varying degrees of uncertainty. This approach allows the framework to assess how different realizations of labour availability affect operational outcomes.

Finally, fleet health is modelled indirectly. Considering its importance as an indicator of airline maintenance planning efficiency, it would be worthwhile to model it with unique decision variables and a specific term in the objective function of the LP model.

10 Conclusion

This paper presents a first example of a robust framework that unifies tail assignment and maintenance task scheduling on the day before operations. Simultaneous optimization of schedule efficiency and stability was made possible by modelling the schedule with an innovative TSN, using two distinct spaces, one dedicated to network and the other to maintenance activities. Additionally, instead of using generic fixed duration maintenance slots, suitable for the entire fleet, the framework creates tail-specific slots, tailored to aircraft's maintenance needs.

Testing the framework within a major European airline case study proved that it is able of providing real-time decision support in an operational environment. Evaluating the framework's plans against those of the airline, showed the framework can improve schedule efficiency and stability. Maintenance time was reduced by 17% as a result of labour utilization increasing from on average 67% to 81%. Consequently, this benefited fleet availability, which on average increased by 10% on the day of operation. Additionally, the framework reduced late arrival disruptions by 42% and propagated delay by 48 min on the day of operations.

This research provides value not only to research but also industry. The computational capabilities of the framework offer an alternative to airline's manual and separate tail and task assignment approach. The framework takes around 250 s for each day analysed. Therefore, the model may be used in an operational environment and serve as a decision-support tool to create schedules that are robust to varying levels of uncertainty and delay. This information can be used to create informed scheduling plans that are robust against delays and maximize fleet utilization. Considering that delays cost airlines on average EUR 0.30 per minute per passenger. Assuming an average of 300 passengers on long-haul flights, the framework can save the airline more than EUR 1.5 millions per year.

CRediT authorship contribution statement

Luigi Pescio: Writing – original draft, Validation, Methodology, Data curation. **Marta Ribeiro:** Writing – review & editing. **Bruno F. Santos:** Writing – review & editing, Supervision.

Declaration of competing interest

The authors certify that they have NO affiliations with or involvement in any organization or entity with any financial interest, or non-financial interest in the subject matter or materials discussed in this manuscript.

Data availability

The authors do not have permission to share data.

References

[1] EUROCONTROL. All-Causes delays to air transport in Europe - Quarter 2 2023. Technical report, 2023.
[2] Liang Z, Feng Y, Zhang X, Wu T, Chaovalitwongse WA. Robust weekly aircraft maintenance routing problem and the extension to the tail assignment problem. Transp Res Part B: Methodol 2015;78:238–59.
[3] Ahmed MB, Ghroubi W, Haouari M, D. SH. A hybrid optimization-simulation approach for robust weekly aircraft routing and retiming. Transp Res Part C: Emerg Technol 2017;84:1–20.
[4] Sadjadi SJ, Safaei N, Shahmoradi-Moghadam H. Robust maintenance scheduling of aircraft fleet: A hybrid SimulationOptimization approach. IEEE Access 2021;9:17854–65.
[5] Cadarso L, Marín Á. Integrated robust airline schedule development. Procedia Soc Behav Sci 2011;20:1041—-1050.
[6] Abara J. Applying integer linear programming to the fleet assignment problem. Interfaces 1989;19(4):20–8.
[7] Hane CA, Barnhart C, Johnson EL, Marsten RE, Nemhauser GL, Sigismondi G. The fleet assignment problem: 1 solving a large-scale integer program. Technical report, 1995.
[8] Barnhart C, Boland NL, Clarke LW, Johnson EL, Nemhauser GL, Shenoi RG. Flight string models for aircraft fleeting and routing. Transp Sci 1998;32(3):208–20.
[9] Sriram C, Haghani A. An optimization model for aircraft maintenance scheduling and re-assignment. Transp Res Part A: Policy Pr 2003;37(1):29–48.
[10] Liang Z, Xiao F, Qianb X, Zhou L, Jin X, Lu X, Karichery S. A column generation-based heuristic for aircraft recovery problem with airport capacity constraints and maintenance flexibility. Transp Res Part B 2018;113:70–90.
[11] Vink J, Santos B, Verhagen W, Medeiros I, Filho R. Dynamic aircraft recovery problem - an operational decision support framework. Comput Oper Res 2020;117.
[12] Deng Q, Santos BF, Curran R. A practical dynamic programming based methodology for aircraft maintenance check scheduling optimization. European J Oper Res 2020;281(2):256–73.
[13] Deng Q, Santos BF. Lookahead approximate dynamic programming for stochastic aircraft maintenance check scheduling optimization. European J Oper Res 2022;299(3):814–33.
[14] Marseguerra M, Zio E. Optimizing maintenance and repair policies via a combination of genetic algorithms and Monte Carlo simulation. Reliab Eng Syst Saf 2000;68(1):69–83.
[15] Witteman M, Deng Q, Santos BF. A bin packing approach to solve the aircraft maintenance task allocation problem. European J Oper Res 2021;294:365–76.
[16] Papakostas N, Papachatzakis P, Xanthakis V, Mourtzis D, Chryssolouris G. An approach to operational aircraft maintenance planning. Decis Support Syst 2010;48(4):604–12.
[17] Van Kessel PJ, Freeman FC, Santos BF. Airline maintenance task rescheduling in a disruptive environment. European J Oper Res 2023;308(2):605–21.
[18] Barnhart C, Kniker TS, Lohatepanont M. Itinerary-based airline fleet assignment. Transp Sci 2002;36(2):199–217.
[19] Sarac A, Batta R, Rump CM. A branch-and-price approach for operational aircraft maintenance routing. European J Oper Res 2003;175(3):1850–69.
[20] Shaukat S, Katscher M, Wu CL, Delgado F, Larrain H. Aircraft line maintenance scheduling and optimisation. J Air Transp Manag 2020;89.
[21] Lagos C, Delgado F, Klappa MA. Dynamic optimization for airline maintenance operations. Transp Sci 2020;54(4):998–1015.
[22] Clausen J, Larsen A, Larsen J, Rezanova NJ. Disruption management in the airline industry—Concepts, models and methods. Comput Oper Res 2010;37(5):809–21.
[23] AhmadBeygi S, Cohn A, Lapp M. Decreasing airline delay propagation by reallocating scheduled slack. IIE Trans 2010;42(7):478–89.
[24] Aloulou MA, Haouari M, Mansour FZ. Robust aircraft routing and flight retiming. Electron Notes Discrete Math 2010;36:367–74.
[25] Lan S, Clarke JP, Barnhart C. Planning for robust airline operations: Optimizing aircraft routings and flight departure times to minimize passenger disruptions. Transp Sci 2006;40(1):15–28.
[26] Froyland G, Maher SJ, Wu CL. The recoverable robust tail assignment problem. Transp Sci 2014;48(3):351–72.
[27] Weide TVd, Deng Q, Santos BF. Robust long-term aircraft heavy maintenance check scheduling optimization under uncertainty. Comput Oper Res 2022;141.
[28] Ageeva YV. Apprroaches to incorporating robustness into airline schedules. Technical report, 2000.
[29] Burke EK, Causmaecker PD, Maere GD, Mulder J, Paelinck M, Berghe GVd. A multi-objective approach for robust airline scheduling. Comput Oper Res 2010;37:822–32.
[30] Lapp M, Cohn A. Modifying lines-of-flight in the planning process for improved maintenance robustness. Comput Oper Res 2012;39(9):2051–62.
[31] Kang LS. Degradable airline scheduling: an approach to improve operational robustness and differentiate service quality. Technical report, 2004.
[32] Smith BC, Johnson EL. Robust airline fleet assignment: Imposing station purity using station decomposition. Transp Sci 2006;40(4):497–516.
[33] Rosenberger JM, Johnson EL, Nemhauser GL. A robust fleet-assignment model with hub isolation and short cycles. Transp Sci 2004;38(3):357–68.
[34] Ding Y, Zhang K, Wang Y, Xu Y, Zhu Y, Xia T, Xi L, Seidelmann J. A joint optimization policy of tail assignment and maintenance scheme for civil aircraft fleet. In: 2023 global reliability and prognostics and health management conference. 2023, p. 1–7.
[35] Dunbar M, Froyland G, Wu CL. An integrated scenario-based approach for robust aircraft routing, crew pairing and re-timing. Comput Oper Res 2014;45:68–86.
[36] Glomb L, Liers F, Rosel F. A stochastic optimization approach for optimal tail assignment with knowledge-based predictive maintenance. CEAS Aeronaut J 2023;14:1001–19.
[37] van der Weide T, Deng Q, Santos BF. Robust long-term aircraft heavy maintenance check scheduling optimization under uncertainty. Comput Oper Res 2022;141:105667.
[38] Hong Y, Choi B, Kim Y. Two-stage stochastic programming based on particle swarm optimization for aircraft sequencing and scheduling. IEEE Trans Intell Transp Syst 2019;20(4):1365–77.