Integrated Stochastic Optimization of Maintenance Scheduling and Tail Assignment with Health-Aware Models in Aviation

Benno Käslin[1,2,3]**, Marta Ribeiro**[1]**, Dimitrios Zarouchas**[1]
and Manuel Arias Chao[1,2]

[1] *Delft University of Technology, 292F+VC Delft, Netherlands*
b.kaeslin@tudelft.nl
m.j.ribeiro@tudelft.nl
d.zarouchas@tudelft.nl
m.a.c.ariaschao@tudelft.nl

[2] *Zurich University of Applied Sciences, 8400 Winterthur, Switzerland*
xka2@zhaw.ch
aria@zhaw.ch

[3] *Swiss International Air Lines Ltd. 8302 Kloten, Switzerland*
benno.kaeslin@swiss.com

Abstract

Efficient maintenance management is essential—not only to reduce costs but also to maximize aircraft availability and uphold safety standards. This requires balancing maintenance scheduling (MS), which drives downtime, with tail assignment (TA), which governs aircraft utilization. While recent research has explored the integration of MS and TA, these efforts have largely neglected the role of Condition-Based Maintenance (CBM) and the uncertainty inherent in prognostic models. This research proposes a novel, unified framework that jointly optimizes MS, TA, and CBM using stochastic programming and health-aware models. The approach leverages sensor-derived prognostic information to forecast component degradation and incorporates its probabilistic nature directly into the planning process. By accounting for uncertainty in remaining useful life (RUL) predictions, the model produces robust flight and maintenance schedules that reduce the risk of unplanned disruptions. Preliminary experiments using real-world airline data demonstrate that explicitly modeling health uncertainty leads to more reliable scheduling outcomes, while improving operational efficiency and reducing maintenance costs. Compared to current industry practice, the integrated framework enables data-driven, future-oriented decision-making at the interface between fleet operations and maintenance planning. This work advances the state-of-the-art by holistically addressing TA, MS, and CBM within a scalable and interpretable optimization model—closing a critical gap in the practical deployment of CBM strategies in civil aviation.

1. Introduction

Efficient maintenance management is crucial in aviation, impacting both operational reliability and cost, with maintenance comprising approximately 11% of total airline expenses (International Air Transport Association, 2024). Traditionally, Maintenance Scheduling (MS) and Tail Assignment (TA) have been addressed in sequential, decoupled processes: MS seeks to minimize downtime through optimal task allocation, while TA maximizes fleet utilization by assigning aircraft to flights. This separation often results in suboptimal solutions, particularly under operational constraints.

The industry is increasingly adopting data-driven and predictive strategies such as Condition-Based Maintenance (CBM), which leverages real-time sensor data and prognostic models to estimate Remaining Useful Life (RUL) and dynamically schedule interventions. While these methods enhance reliability and reduce unnecessary maintenance, their integration with operational planning remains limited due to data quality challenges and the intrinsic uncertainty of prognostic outputs (Verhagen & Curran, 2023).

Recent studies have sought to jointly optimize MS and TA (Sriram & Haghani, 2003; Lagos, Delgado, & Klapp, 2020), but existing frameworks rarely incorporate CBM or model prognostic uncertainty. Additionally, advances in large language models (LLMs) have shown promise in improving the interactivity and adaptability of scheduling systems (Pallagani, Katz, Marques-Silva, Kumar, & Yeoh, 2024).

This work introduces a unified framework that concurrently optimizes MS, TA, and CBM using stochastic programming and health-aware optimization. By integrating sensor-driven prognostic data and modeling RUL uncertainty, the approach delivers robust, cost-efficient schedules suitable for real-world airline operations.

2. RESEARCH GAPS AND OBJECTIVES

The contributions are:

- A joint optimization model is formulated that integrates tail assignment, maintenance scheduling, and condition-based maintenance.
- Uncertainty from prognostic health models is explicitly incorporated using stochastic programming, enabling realistic and risk-aware planning.
- The proposed approach is validated on real-world airline data, demonstrating that health-aware scheduling results in fewer disruptions and reduced maintenance costs compared to current industry practice.
- The integration of LLMs is assessed to facilitate natural language interaction within the scheduling framework, thereby enhancing usability for planners.

This work aims to close a critical gap between CBM research and airline operations by offering a scalable and interpretable solution that supports the practical adoption of CBM in real-world planning environments.

3. METHODOLOGY

The proposed framework (Fig. 1) integrates three primary input streams:

- **Maintenance Data:** Task lists, available slots, and workforce capacity, enabling efficient resource allocation under operational constraints.
- **Flight Plan:** Scheduled flights and routes, ensuring optimal tail assignment and seamless integration of maintenance without disrupting operations.
- **Prognostic Models:** Sensor-driven, uncertainty-aware predictions of component end-of-life, supporting risk-informed, condition-based maintenance decisions.

The framework leverages LLMs to translate optimization outcomes into interpretable language for planners, and enables interactive updates for real-time solution refinement.

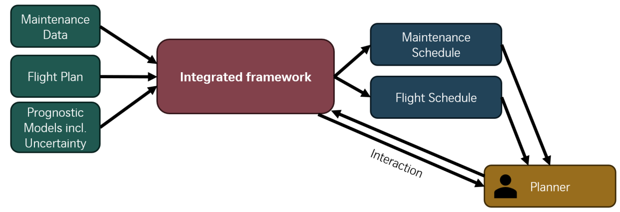

Figure 1. Integrated framework overview: maintenance, flight plan, and prognostic models jointly inform the optimization of MS and TA, supporting both automated and planner-involved scheduling.

3.1. Integrated Stochastic Optimization

At its core, the framework employs a unified optimization model that simultaneously addresses:

1. **Maintenance Scheduling (MS):** Assignment of tasks to slots considering constraints and deadlines.
2. **Tail Assignment (TA):** Optimal aircraft-to-flight allocation respecting operational requirements.
3. **Condition-Based Maintenance (CBM):** Scheduling of prognostic maintenance tasks under uncertainty.

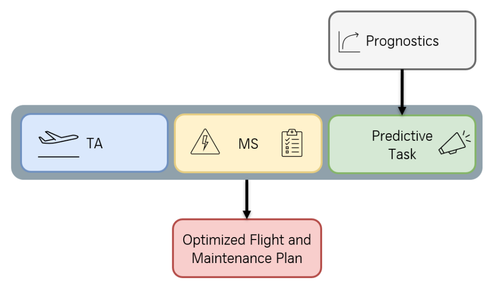

Figure 2. Unified MILP formulation integrating MS, TA, and CBM with explicit uncertainty modeling.

Unlike sequential or decoupled methods, the problem is formulated as a single extensive framework, capturing interdependencies between maintenance and flight operations while explicitly modeling prognostic uncertainty.

3.2. Uncertainty Modeling in Prognostic Tasks

A distinguishing feature is the explicit incorporation of prognostic uncertainty. For each monitored component, the model includes:

- **Estimated Remaining Useful Life (RUL):** Data-driven failure forecasts,
- **Uncertainty Distributions:** Confidence intervals for RUL estimates,
- **Scenario Sampling:** Multiple realizations of health trajectories drawn from the uncertainty distribution.

This scenario-based formulation supports robust optimization, allowing maintenance schedules to proactively manage operational risk and adapt to the probabilistic nature of component degradation.

4. INITIAL RESULTS

Building upon the foundational work of Oremans (Oremans, n.d.), this study advances the joint optimization of flight and maintenance scheduling under prognostic uncertainty. From her work and the further development of the framework, flight cancellations were identified as the predominant cost driver within the objective function. Incorporating uncertainty into prognostic maintenance tasks reveals that certain realizations of component health trajectories may cause additional cancellations—a phenomenon that deterministic models fail to capture.

To address this challenge, the risk of cancellations arising from uncertainty in component health predictions is modeled using two methodological approaches:

- A neural network-based model to estimate cancellation probabilities as a function of maintenance scheduling decisions and stochastic health scenarios;
- A compact mini-MILP formulation to assess schedule robustness across multiple uncertainty scenarios.

These methods are evaluated on a real-world case study from SWISS International Air Lines Ltd., involving 5 short-haul aircraft, 92 scheduled maintenance tasks, 5 prognostic tasks, and 132 flights over a 5-day horizon. The following solution strategies are compared:

- **Deterministic:** MILP integrating maintenance scheduling (MS), task assignment (TA), and prognostic tasks, without uncertainty;
- **Stochastic:** An approach that explicitly models uncertainty in component health;
- **Reality:** The actual schedule implemented by the airline.

The resulting deterministic schedule is illustrated in Figure 3, with flights depicted in blue and maintenance slots in yellow.

Figure 3. Flight (blue) and Maintenance (yellow) schedule resulting from deterministic optimization. Uncertainty is not considered in this framework.

A central finding is that cancellations predominantly arise when the RUL of components is overestimated. To quantify the operational impact of RUL uncertainty, probabilistic cancellation costs are modeled under various prognostic scenarios. For illustration, a simplified example with two aircraft and three operational days is analyzed, highlighting the sensitivity of cancellations to deviations from mean RUL predictions (see Figure 4).

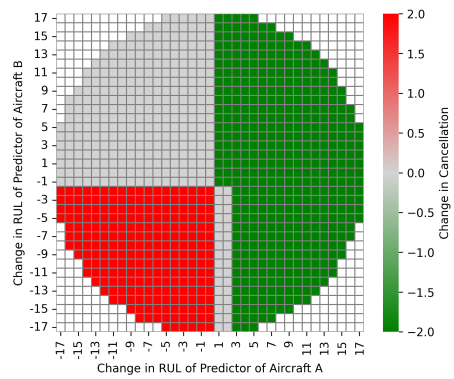

Figure 4. Impact of RUL variation on flight cancellations for two aircraft over three days. Red indicates an increase in cancellations relative to the mean RUL scenario, while green indicates a reduction, hence less cancellations compared to the mean RUL. Mean RULs of aircraft A and B are 2 and 3 flight cycles, respectively; 15 flights are scheduled in total.

Preliminary results indicate that the stochastic framework—whether based on neural network or mini-MILP—yields schedules that are more robust to prognostic uncertainty, substantially reducing the risk of cancellation due to unforeseen component failures. The additional computational effort is justified by the resulting gains in operational resilience.

5. RESEARCH PLAN AND FUTURE WORK

Ongoing and future research is organized into the following work packages:

5.1. WP1: Advanced Stochastic Programming Approaches

Extending the current implementation, advanced stochastic programming methods are investigated to more accurately represent uncertainty in prognostic maintenance, as proposed by Dumouchelle et al. (Dumouchelle et al., 2022). In their approach, a neural network is used to estimate second-stage costs, which are subsequently linearized for MILP integration. Key focus areas include:

- Applying scenario reduction to manage computational complexity,

- Adopting risk-averse optimization to address worst-case outcomes.

5.2. WP2: Reinforcement Learning and Hybrid Approaches

Recent work (Tseremoglou & Santos, 2024; Silva, Alves, Ribeiro, Rizzotto, & Weigang, 2023; Song, Liu, Qin, Wang, & Chen, 2024) demonstrates the efficacy of reinforcement learning (RL) for complex maintenance scheduling. Notable approaches combine dynamic scheduling via POMDPs and deep RL, adaptive rescheduling, and DRL-augmented genetic algorithms. Building on these insights, the following directions are proposed:

- Apply Advantage Actor-Critic (A2C) algorithms to dynamic scheduling,
- Investigate hybrid MILP–RL frameworks integrating mathematical optimization and RL.

Further decomposition of the scheduling problem into bin-packing and slot/flight assignment components (Witteman, Santos, & Leal de Matos, 2021) is proposed, with reinforcement learning applied to both subproblems using delayed rewards to capture interdependencies.

5.3. WP3: Scaling to Different Fleets

Practical deployment requires scaling to realistic airline fleets characterized by:

- Multiple aircraft types with diverse maintenance regimes,
- Varied operational profiles (short-haul and long-haul),
- Numerous components, each with unique maintenance requirements.

Decomposition, parallelization, and hierarchical optimization will be investigated to ensure tractability for large, heterogeneous fleets.

5.4. WP4: Enhanced Explainability with Large Language Models

To enhance framework usability for planners, the integration of large language models (LLMs) is explored to:

- Translate optimization outputs into interpretable natural language explanations,
- Enable user interaction via natural language, supporting real-time updates and solution refinement.

6. CONCLUSION

This research introduces a novel, integrated framework for the stochastic optimization of maintenance scheduling and tail assignment with condition-based maintenance in aviation. The current gap in the literature is addressed by jointly optimizing maintenance scheduling and tail assignment while incorporating condition-based maintenance and its inherent uncertainty within a single, comprehensive optimization model.

Initial results indicate that the proposed approach improves operational efficiency and reduces maintenance costs compared to current industry practice. The framework's ability to schedule additional maintenance tasks while maintaining robust flight operations underscores its practical value for airlines.

Ongoing research is focused on enhancing stochastic modeling capabilities, exploring hybrid optimization–learning approaches, scaling the framework to diverse fleets, and improving explainability through integration with large language models. By addressing these challenges, a comprehensive and practical solution is being developed to enable airlines to optimize maintenance operations under uncertainty.

ACKNOWLEDGMENT

The authors thank SWISS International Air Lines Ltd. for providing data and support for this research.

REFERENCES

Dumouchelle, J., Li, F., Khalil, E. B., Nemhauser, G. L., Sokol, J., Wang, Y., & Yuan, S. (2022). Neur2sp: Neural two-stage stochastic programming. *Advances in Neural Information Processing Systems*, *35*, 28775–28789.

International Air Transport Association. (2024). *Annual Review 2024.* IATA Annual Review. Retrieved from `https://www.iata.org/en/publications/annual-review/`

Lagos, C., Delgado, F., & Klapp, M. A. (2020, July). Dynamic optimization for airline maintenance operations. *Transportation Science*, *54*(4), 998–1015. doi: 10.1287/trsc.2020.0984

Oremans, M. (n.d.). *Stochastic optimisation of tail assignment and maintenance task scheduling* (mathesis).

Pallagani, V., Katz, M., Marques-Silva, J., Kumar, A., & Yeoh, W. (2024). Large language models for automated planning and scheduling: Promises and pitfalls. *arXiv preprint arXiv:2402.05233*.

Silva, V. P., Alves, V., Ribeiro, M., Rizzotto, F., & Weigang, L. (2023). Multi-level task scheduling optimization for aircraft component maintenance. *Aerospace*, *10*(1), 77.

Song, Z., Liu, Y., Qin, L., Wang, L., & Chen, X. (2024). A deep reinforcement learning-based genetic algorithm for multi-type satellite observation scheduling problem. *IEEE Transactions on Cybernetics*, *54*(3), 1518–

1532.
Sriram, C., & Haghani, A. (2003). Integrated airline fleet assignment and maintenance scheduling. *Transportation Research Record*, *1844*(1), 60–72.
Tseremoglou, C., & Santos, B. F. (2024). Two-stage dynamic scheduling with stochastic components health condition. *Reliability Engineering & System Safety*, *241*, 109856.
Verhagen, W. J., & Curran, R. (2023). Condition-based and predictive maintenance in aviation: Current status and pathway toward implementation. *Progress in Aerospace Sciences*, *135*, 100869.
Witteman, M. S., Santos, B. F., & Leal de Matos, P. (2021). Optimizing maintenance task allocation within aircraft line maintenance. *Transportation Research Part C: Emerging Technologies*, *131*, 103304.

BIOGRAPHIES

Benno Käslin is a PhD candidate at the Operations and Environment Chair of the Delft University of Technology in the Netherlands since April 2024. He received his MSc in Mechanical Engineering with focus on Aerospace from ETH Zürich in 2020. His research focuses on integrating maintenance scheduling, tail assignment, and condition-based maintenance in aviation, with emphasis on stochastic optimization and AI applications.

Marta Ribeiro is an Assistant Professor in the Operations and Environment Chair at Delft University of Technology in the Netherlands. She holds a PhD in Aerospace Engineering from TU Delft, with a thesis on Conflict Resolution at High Traffic Densities with Reinforcement Learning, and a Master's degree in Avionics from Instituto Superior Técnico in Lisbon. Marta has gained industrial and research experience as an Operations Engineer at Portugália Airlines, a simulator developer at Simteq, and a systems engineer at CGI working on the Galileo satellite navigation project for the European Space Agency. Her research focuses on the application of artificial intelligence and machine learning to optimize aviation operations, including condition-based maintenance, disruption management, airline scheduling, and delay estimation.

Dimitrios Zarouchas is an Associate Professor in the Aerospace Structures & Materials Department at Delft University of Technology in the Netherlands. He leads the PYTHIA research team on Artificial Intelligence for Structures, Prognostics, and Health Management, and is the founder of the Center of Excellence in AI for Structures. His work focuses on developing intelligent systems for real-time diagnostics and prognostics of lightweight structures used in aviation, space, wind energy, and naval industries.

Manuel Arias Chao is a Senior Lecturer at Zurich University of Applied Sciences and an Assistant Professor at the Operations and Environment Chair of the Delft University of Technology in the Netherlands. He has a PhD in Physics-informed Machine Learning for Prognostics and Health Management from ETH Zurich, a Master's degree in Thermal Power from Cranfield University, and a Bachelor's degree in Aeronautical Engineering from the Technical University of Madrid. Manuel has gained valuable industrial and research experience as a visiting researcher at the Diagnostics & Prognostics Group at NASA Ames, Thermodynamics & Performance Lead Engineer at General Electric and ALSTOM Power, and Aero Engine Maintenance Engineer at ITP. In his current role, Manuel also co-leads the Expert Group Smart Maintenance from the Swiss Alliance for Data-Intensive Services.