Mitigating Uncertainty in an Extended-Arrival Manager Environment

Evaluating Pop-Up Flight Mitigation Strategies Using Stochastic Uncertainty Models

Jorn van Beek

Delft University of Technology

TU Delft

Mitigating Uncertainty in an Extended-Arrival Manager Environment

Evaluating Pop-Up Flight Mitigation Strategies Using Stochastic Uncertainty Models

by

Jorn van Beek

to obtain the degree of Master of Science at the Delft University of Technology,

Supervisor:	Dr. J. Ellerbroek
External Supervisor (LVNL):	F. Dijkstra
Thesis Committee:	Dr. M.M. van Paassen
Thesis Committee:	Dr. P. Proesmans
Faculty:	Control & Operations, Aerospace Engineering, Delft

Cover: RWY 27 Final Stack (Ronald Stella)

An electronic version of this thesis is available at http://repository.tudelft.nl/.

TU Delft

Preface

This paper marks the final milestone of my Aerospace Engineering degree at TU Delft, as well as the completion of my graduation project at the KDC of LVNL.

During this project, I have learned a great deal about myself and about where my interests truly lie, but above all about Air Traffic Management, LVNL, and the complex operational side of air traffic control, with arrival management in particular. I am proud to have been able to contribute, even in a small way, to improving the operational process, and I hope that this research can make a meaningful difference.

I would like to sincerely thank Joost Ellerbroek for his support throughout the process, for the challenging questions, valuable feedback, and continuous guidance. I would also like to thank Ferdinand Dijkstra for all his enthusiasm and for his help in understanding the operational side of the project. Without the two of you, this project would never have gotten off the ground.

In addition, I would like to thank my fellow graduate students at iLabs for the many "bakkies", "uitsmijters" and the new insights and distraction during those. I would also like to thank my girlfriend, Iris, for her patience, support and listening ear throughout this project, as well as my parents for their support during my entire studies.

Finally, I would like to thank you, the reader of this thesis, for taking the time to read my work. I hope that, in one way or another, it proves useful to you. If there is anything I can do to help you make use of this work, please do not hesitate to contact me!

Jorn van Beek
Delft, April 2026

Contents

Mitigating Uncertainty in an Extended-Arrival Manager Environment

Evaluating Pop-Up Flight Mitigation Strategies Using Stochastic Uncertainty Models

Jorn van Beek, Dr. Ir. Joost Ellerbroek
Aerospace Engineering, TU Delft
Delft, The Netherlands

Ferdinand Dijkstra
LVNL (Air Traffic Control the Netherlands)
Schiphol, The Netherlands

***Abstract*—Extended Arrival Management (E-AMAN) enables earlier sequencing and upstream delay absorption by extending the freeze horizon, but increases exposure to prediction uncertainty and pop-up flights, which can degrade planning stability. This study presents the BlueSky AMAN Simulator (BAMS) and evaluates how stochastic uncertainty and pop-up mitigation strategies affect E-AMAN performance at Amsterdam Airport Schiphol.**

Paired Monte-Carlo experiments are performed for freeze horizons of 14, 20, and 25 minutes under stochastic take-off, departure-route, and en-route uncertainty. Multiple mitigation strategies are evaluated, including Back-of-the-Line (BOL) scheduling, delayed-slot scheduling, and enabling planning at take-off.

The results show a clear trade-off between earlier planning and stability. Pop-up flights primarily drive sequence disruptions, while trajectory prediction uncertainty mainly increases temporal instability through repeated Expected Approach Time (EAT) revisions. Among the evaluated strategies, BOL scheduling at a 20-minute freeze horizon provides a balance between stability and delay performance, while longer horizons show diminishing returns due to uncertainty propagation. In addition, planning at take-off is shown to improve performance in both current and extended AMAN operations.

The findings indicate that feasible E-AMAN implementation at Schiphol requires either reduced uncertainty or stability-preserving scheduling strategies.

***Keywords*—Extended Arrival Management (E-AMAN), AMAN, Pop-up flights, BlueSky, Freeze horizon extension, Air traffic flow management (ATFM)**

I. Introduction

Arrival Manager (AMAN) systems are widely used at major hub airports to regulate inbound traffic flows and support runway sequencing [1]. By assigning landing slots and Expected Approach Times (EATs), AMAN enables delay absorption outside the Terminal Maneuvering Area (TMA), reducing controller taskload, fuel burn, and the need for tactical interventions close to the airport [1]. In particular, absorbing delay earlier in the trajectory increases the opportunity to use small speed adjustments instead of more disruptive measures such as vectoring or holding [1]–[3].

An AMAN sequence becomes fixed once aircraft cross a predefined freeze horizon (FH). At the freeze horizon, aircraft are assigned a landing slot and EAT that are treated as fixed. The assigned EAT is used by air traffic control to issue speed, vectoring, or holding instructions, ensuring that aircraft arrive at the Initial Approach Fix (IAF) within the required time window. Beyond this horizon, sequence changes only occur when necessary to preserve planning effectiveness [4]. The robustness of this process is directly influenced by uncertainty in arrival time predictions [3].

This uncertainty originates across multiple flight phases, including take-off planning and en-route trajectory prediction [5], [6]. One operational consequence is the occurrence of *pop-up flights*, defined as aircraft that become visible to the AMAN only after the relevant part of the sequence has already been frozen. Such flights trigger re-planning actions, reducing planning stability and increasing controller workload [7].

Extended Arrival Management (E-AMAN) shifts the planning process further upstream by increasing the freeze horizon. A longer freeze horizon provides more time to absorb delay before aircraft reach the TMA, increasing the opportunity to use speed control instead of less efficient measures such as vectoring or holding. However, this upstream shift also increases exposure to prediction uncertainty [5]. As the horizon is extended, more nearby departures behave as pop-up flights [7], while trajectory prediction uncertainty increases due to the longer remaining flight segment that must be predicted [6]. These effects directly degrade E-AMAN performance, particularly in terms of planning stability [7], highlighting the trade-off between earlier planning and increased uncertainty.

Under SESAR Common Project One (CP1), E-AMAN must be implemented with freeze horizons up to 180 NM from the arrival airport at major European airports [8]. As a result, operations must cope with increased uncertainty at longer planning horizons, which creates a need for effective mitigation strategies for pop-up flights.

The challenge is particularly relevant for Amsterdam Airport Schiphol. Due to the proximity of major departure airports, including Brussels and several airports in western Germany and the United Kingdom, horizon extension is expected to increase the occurrence of pop-up flights. Schiphol currently operates the Advanced Schiphol Arrival Planner (ASAP), which uses a time-based freeze horizon of 14 minutes to the Initial Approach Fix (IAF) [4], [9]. Even under this baseline configuration, some nearby departures already behave as pop-up flights. In particular, flights departing from Brussels typically have a flight time between radar visibility

and the IAF of only 10 to 12 minutes, causing them to enter the planning after the sequence has been frozen.

As a result, nearby departures may enter the planning late, while at the same time extending the freeze horizon increases the uncertainty on which planning decisions are based. E-AMAN performance at Schiphol is therefore highly sensitive to both traffic structure and uncertainty propagation, making it a relevant case for investigating the interaction between horizon extension and stochastic uncertainty.

Several mitigation concepts have therefore been proposed to limit the operational impact of horizon extension. These include time-based and rule-based metering strategies [10] and Required Time of Arrival (RTA) procedures [11]. Van Welsenaere [7] also evaluated simpler mitigation approaches, such as pre-planning nearby departures and scheduling concepts that implicitly leave more room between slots depending on the level of uncertainty. While such approaches can improve predictability or reduce the impact of pop-up flights on performance, their applicability in dense European airspace is often limited. Feasibility studies further indicate that Extended AMAN performance is strongly governed by uncertainty propagation across flight phases [12].

A clear research gap remains. Existing work has identified the main uncertainty sources and has described the operational risks of horizon extension, but there is still no integrated evaluation of how stochastic uncertainty, freeze horizon extension, and operationally realistic pop-up mitigation strategies interact in a closed-loop E-AMAN environment. This gap is particularly relevant for Schiphol, where nearby departures, pop-up behavior, and trajectory uncertainty interact strongly. There is also limited understanding of whether earlier flight visibility, for example through EFD-based planning at take-off, can reduce the negative effects of pop-up behavior.

This study addresses that gap by combining phase-specific uncertainty modeling based on operational ETFMS Flight Data (EFD) messages with a BlueSky-based AMAN simulation environment [13].

The paper makes three main contributions. First, it derives stochastic uncertainty models for take-off planning, departure route, and en-route trajectory prediction using one year of operational data. Second, it integrates these uncertainty models into the BlueSky AMAN Simulator (BAMS), enabling reproducible paired Monte-Carlo experiments under identical stochastic conditions. Third, it evaluates the effects of freeze horizon extension and multiple pop-up mitigation strategies, including back-of-the-line scheduling, delayed-slot planning, planning at take-off using EFD messages, and the influence of Brussels pop-up flights.

II. Uncertainty Modeling & Data Analysis

This section describes how operational temporal uncertainty is quantified from ETFMS Flight Data (EFD) messages and subsequently translated into stochastic uncertainty models for use in the AMAN simulation framework. The objective of this section is to represent temporal uncertainty in a data-driven and phase-consistent manner, such that it can be realistically reproduced in Monte-Carlo simulations further in this research.

In earlier AMAN-related research, uncertainty is often treated as a single aggregated source of prediction error. In this study the uncertainty is split into four separate phases. This phase-based decomposition is used because the dominant error sources and their statistical characteristics differ substantially across flight phases [6]. In particular, take-off prediction uncertainty is considerably larger than the later trajectory-prediction uncertainties and is driven by different operational processes. Representing all uncertainty as a single error source would therefore hide relevant differences in both magnitude and mechanism.

The analysis is based on one year of EFD messages, comprising approximately 260,000 flights inbound to Amsterdam Airport Schiphol, reconstructed into per-flight prediction and actual histories. Using this data, the total arrival-time prediction error at the IAF is decomposed into four components: take-off planning error, departure route error, trajectory prediction (TP) error outside the FIR, and TP error within the FIR. The latter two correspond to distinct operational contexts: outside the FIR, controller intent is largely unknown, while inside the FIR the aircraft is subject to the local Air Navigation Service Provider (ANSP) and AMAN environment.

Each of these phases is analyzed in detail in the following subsections. Fig. 1 provides a summary overview of the resulting error distributions based on the phase-specific definitions and methods described below. This figure is intended to give an overview of the relative magnitude and distribution of the different error sources. The phase-based decomposition is also supported by the global error characteristics shown in Fig. 1. The figure presents the general prediction error distributions for the modeled uncertainty phases and shows that these phases differ clearly in both spread and shape. Take-off planning uncertainty has the largest spread, while the departure-route and en-route trajectory prediction errors are smaller and more concentrated. At the same time, the departure-route segment represents only a relatively short part of the total trajectory. Even though its absolute contribution to total arrival-time uncertainty is therefore more limited, it remains relevant for nearby departures and pop-up behavior. These differences in magnitude, operational origin, and time scale motivate modeling the uncertainty in separate phases rather than as a single aggregated error source.

A. ETFMS Flight Data

The uncertainty analysis is based on ETFMS (Enhanced Tactical Flow Management System) Flight Data (EFD) messages, which provide updates of flight planning and execution. Each message contains a timestamp, route information, and estimated times over trajectory points. EFD messages are updated both periodically and event-driven [**?**]. Typically, updates are provided at hourly intervals several hours before departure, increasing to approximately 15-minute intervals within two hours of departure, with additional updates triggered by changes in planning or predicted times [**?**].

The dataset does not directly contain explicit Estimated or Actual Take-Off Times (ETOT/ATOT). Instead, predicted times are obtained from the estimated time over relevant

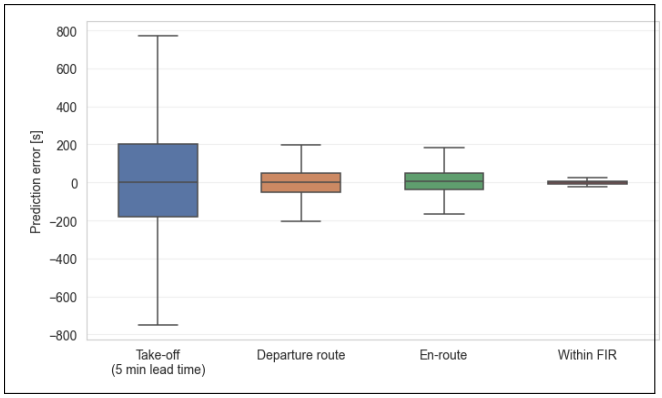

Figure 1. Overview of the prediction error distributions for the modeled uncertainty phases. The figure highlights differences between take-off planning uncertainty, departure-route uncertainty, en-route trajectory prediction uncertainty outside the FIR, and en-route trajectory prediction uncertainty within the FIR.

trajectory points, while actual times are reconstructed from correlated position report (CPR) messages, which provide updated timestamps along the flown trajectory [14].

Prediction errors are defined as the difference between predicted and reconstructed actual times over a given point. The prediction lead time is derived from the message timestamp, representing the time between the prediction and the corresponding event. This enables reconstruction of per-flight time histories and phase-specific prediction errors used in the following analysis. The specific definitions and processing steps for each uncertainty component are described in the corresponding subsections.

B. Take-Off Planning Uncertainty

Take-off prediction error has been defined per prediction as the difference between the actual take-off time (ATOT) and the estimated take-off time (ETOT),

$$E_{\mathrm{TO}} = T_{\mathrm{ATOT}} - T_{\mathrm{ETOT}}, \tag{1}$$

while the prediction lead time (lookahead) is defined as

$$L_{\mathrm{TO}} = T_{\mathrm{ETOT}} - T_{\mathrm{msg}}. \tag{2}$$

For take-off planning uncertainty, prediction lead time is an essential variable, because the uncertainty is strongly affected by variable ground processes and update timing [6]. Errors are therefore grouped into 1-minute lead-time buckets to capture how uncertainty evolves toward take-off. In addition to the discrete per-message representation, a continuous representation is constructed by keeping the error constant between updates, for simulation implementation.

Fig. 2 shows the estimated take-off time error spread as a function of prediction lead time. It shows the tendency of reduced uncertainty at shorter prediction lead times. Differences can be observed between different airports. Previous work has concluded that the differences are due to the Collaborative Decision Making airport type [6]. CDM enabled airports provide more specific information and more frequent updates concerning departure status, which in theory could reduce the amount of uncertainty concerning take-off planning [15]. However, the CDM airport type is not the only driver. A more accurate feature to select by is the individual airports themselves. The figure shows quite significant differences between airports. The take-off uncertainty is represented in the simulation individually per airport to accurately capture these differences in uncertainty statistics between airports. Furthermore, prediction lead times are also represented in individual 1-minute distributions in the simulation.

C. Departure Route Trajectory Prediction Uncertainty

Departure route prediction uncertainty captures the error between predicted and realized flight time in the initial climb phase. The analysis compares the predicted and actual time to reach a fixed downstream point (above FL180, or FL100 for departures from Brussels) and normalizes the resulting error by predicted flight time. This enables comparison across routes and representation in the AMAN simulator.

$$E_{\mathrm{DEP}} = (T_{\mathrm{point,act}} - T_{\mathrm{ATOT}}) - (T_{\mathrm{point,pred}} - T_{\mathrm{ETOT}}), \tag{3}$$

$$E_{\mathrm{DEP,norm}} = 100 \cdot \frac{E_{\mathrm{DEP}}}{(T_{\mathrm{point,pred}} - T_{\mathrm{ETOT}})}. \tag{4}$$

Compared to take-off planning, departure route error as shown in Fig. 3 is smaller in magnitude, but still non-negligible for (E-)AMAN performance, especially for near-horizon departures. In this phase, prediction lead time is not used as a feature, since deviations from the filed departure route are mainly driven by procedures and ATC instructions rather than by the time of prediction. The analysis indicates that differences between CDM types are less pronounced in this phase, while differences between individual airports can still be substantial. Departure route uncertainty is therefore modeled using airport-specific distributions.

D. En-Route Trajectory Prediction Uncertainty

En-route trajectory prediction uncertainty is treated separately for predictions outside the FIR and within the FIR, since these correspond to different operational conditions. Outside the FIR, controller and pilot intent are largely unknown, which yields wider and potentially skewed prediction errors. For this phase, prediction errors are evaluated using a fixed lookahead window of 12–16 minutes prior to the coordination point (COP). Fig. 4 shows the resulting normalized en-route flight time error up to the COP.

For the within-FIR phase, trajectory prediction uncertainty is represented by a normal distribution with a standard deviation of 1.5% of the predicted segment flight time. This value is based on SESAR/EUROCONTROL predictability targets for the overall flight [16] and it is in line with the value used within LVNL for internal trajectory prediction estimates. The value also reflects the reduced uncertainty associated with locally controlled intent.

E. Stochastic Representation of Operational Uncertainty

For each flight phase, empirical error distributions are derived from the EFD message analysis. In the case of Take-off planning error and departure route prediction error, distributions are modeled using a Johnson-SU probability

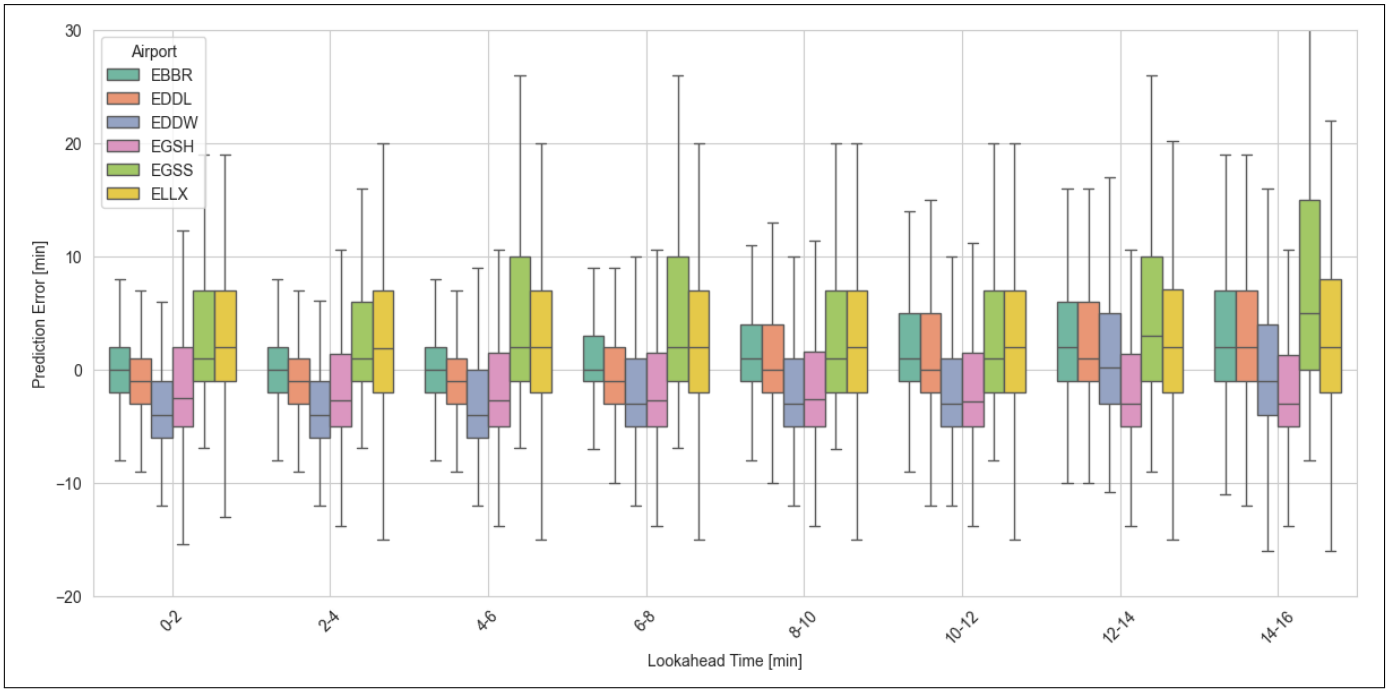

Figure 2. Estimated Take-Off Time (ETOT) prediction error as a function of prediction lead time for nearby departure airports. Each airport is shown separately to illustrate both the reduction in uncertainty toward take-off and the differences in error spread between airports.

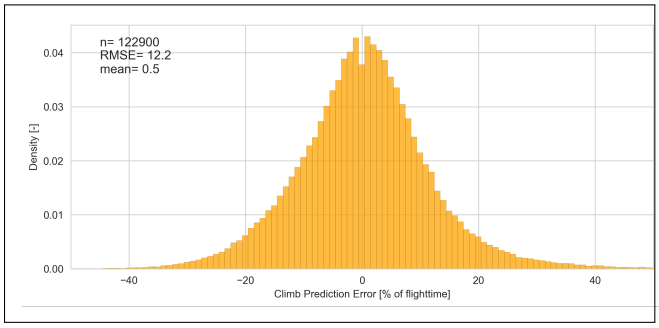

Figure 3. Distribution of normalized departure-route flight time prediction error for nearby departure airports. The error is expressed relative to predicted segment flight time, enabling comparison between routes and airports with different nominal flight times.

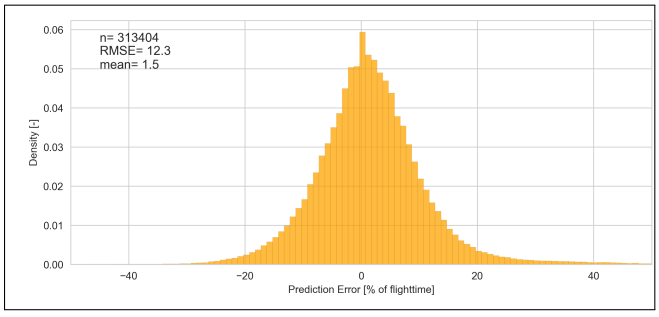

Figure 4. Normalized en-route flight time prediction error to the coordination point (COP), evaluated using a fixed lookahead window of 12–16 minutes. The error is expressed relative to predicted segment flight time.

density function (PDF) [6]. Johnson-SU distributions are used to capture skewness and heavy-tailed behavior occurring in operational prediction errors. Within-FIR trajectory prediction uncertainty is represented using a normal distribution, reflecting the reduced variability associated with locally controlled intent. These distributions form the stochastic input models used in the Monte-Carlo simulation framework. The mechanisms by which the error stemming from this uncertainty are introduced into the simulation environment are described in Section III. Although prediction errors are not strictly independent across phases, independence is assumed for the stochastic modeling.

The phase-based decomposition is essential for the subsequent simulation experiments, as different uncertainty sources affect AMAN performance through different mechanisms. Take-off planning uncertainty primarily determines the predictability of nearby departures before radar visibility, and is therefore particularly relevant when evaluating pre-planning of pop-ups. Departure route and en-route trajectory prediction uncertainty mainly affect the predicted time at IAF with continuous updates. By separating these effects in the stochastic model, the simulation framework can distinguish between instability caused by late flight visibility and instability caused by continuous prediction updates. This enables a more targeted evaluation of freeze horizon extension and pop-up mitigation strategies in the following sections.

III. Simulation Framework

To evaluate the impact of stochastic operational uncertainty on Extended Arrival Manager (E-AMAN) performance, a dedicated simulation framework was developed in BlueSky [13]. This framework integrates trajectory prediction, arrival

manager logic, and air traffic control behavior within a modular architecture. The main components of this framework are described in the following subsections.

The resulting simulation environment, referred to as the **BlueSky AMAN Simulator (BAMS)**, is designed to reproduce the operational principles of the Arrival Manager currently used at Amsterdam Airport Schiphol (Advanced Schiphol Arrival Planner (ASAP)), while incorporating Schiphol-specific procedures and controller strategies. At the same time, the framework is structured generically, allowing application to other airport environments with limited adaptation. BAMS is intended primarily for relative comparison between configurations and not intended for absolute performance prediction.

A. Simulation Environment: BlueSky

The simulation framework is implemented using the BlueSky open-source air traffic simulator. BlueSky provides a physics-based aircraft performance model, flight management logic, and air traffic control interfaces, enabling the representation of realistic aircraft trajectories and controller interventions [13]. Its modular plugin architecture allows direct integration of custom components, including the trajectory predictor (TP), Arrival Manager (AMAN/E-AMAN), and ATC logic used in this study.

B. Trajectory Predictor (TP)

The TP serves as the primary source of arrival-time information for the AMAN. It provides Estimated Times Over (ETO) along the flight trajectory, including the Estimated Time Over the Initial Approach Fix (ETO IAF), and the corresponding Estimated Time of Arrival (ETA). Nominal trajectory predictions are generated using BlueSky by propagating each flight along its planned route under deterministic conditions. This produces internally consistent baseline ETO and ETA values prior to uncertainty introduction. During simulation, TP estimates are updated whenever tactical ATC instructions modify the aircraft trajectory. The predictions are deterministic, as the current traffic situation, including the active aircraft route, speed profile, and known control instructions, is propagated forward in BlueSky without introducing additional stochasticity within the predictor itself. Uncertainty therefore does not originate inside the TP, but enters the simulation through the separate stochastic error realization mechanisms.

C. Uncertainty Injection and Error Realization

Operational uncertainty is introduced into BAMS using phase-consistent error realization mechanisms derived from the stochastic models described in Section II. Rather than directly perturbing predicted arrival times, uncertainty is injected by altering take-off time and groundspeed.

1) Sampling: Phase-specific uncertainty realizations are sampled from the Johnson SU probability density functions derived in Section II. To ensure statistical consistency and full reproducibility across Monte-Carlo runs, a deterministic sampling mechanism is applied.

Uncertainty realizations are sampled using deterministic pseudo-random sequences derived from a scenario seed and aircraft identifier, ensuring reproducibility and paired stochastic conditions across configurations.

2) Take-off Uncertainty Realization: Take-off planning uncertainty is modeled as a deviation between the Estimated Take-off Time (ETOT) and the Actual Take-off Time (ATOT):

$$ATOT = ETOT + E_{\mathrm{TO}}, \tag{5}$$

where E_{TO} is sampled from the airport- and lead-time-dependent take-off error distribution. This error is relevant only when planning pop-up flights before take-off. This capability is not part of current Schiphol operations and is therefore excluded from the baseline representation.

3) Airborne Uncertainty Realization: For airborne phases, uncertainty is introduced by modifying the effective ground speed within the relevant trajectory segment:

$$V_{\mathrm{ground,eff}} = V_{\mathrm{ground,nom}} \cdot (1 - E_{\mathrm{segment}}), \tag{6}$$

where E_{segment} represents the sampled normalized prediction error associated with the departure-route, en-route outside FIR, or en-route within FIR phase. The modification of ground speed emulates uncertainty in lateral path execution, which translates into uncertainty in the arrival time at the IAF. The lateral ground track itself, and its exact uncertainty, are not directly relevant for AMAN, since the system is affected only by the resulting temporal uncertainty.

The values of ETO IAF and ETA are updated according to this error propagation for accurate representation of the prediction error in the AMAN. Since uncertainty is realized through an effective ground-speed deviation over the remaining trajectory segment, the corresponding timing error is not constant. Instead, the absolute prediction error decreases as the aircraft approaches the end of the affected segment. As a result, the AMAN observes an ETA error that gradually reduces approximately linearly toward zero over time.

Once aircraft intent becomes known, either through FIR entry or the application of a control instruction (including those initiated by adjacent centers), the prediction error model transitions to the within-FIR uncertainty representation. Subsequent trajectory prediction uncertainty is therefore governed by the corresponding normal distribution, consistent with the reduced stochastic variability occurring under local control conditions and known intent.

D. Arrival Manager

The Arrival Manager allocates landing slots based on predicted arrival times provided by the trajectory predictor. Upon entering the planning horizon, aircraft are sequenced according to their ETA. Slot allocation follows a separation-based mechanism in which each assigned slot is determined relative to the previously scheduled aircraft:

$$\mathrm{slot}_i = \max\left(\mathrm{slot}_{i-1} + \mathrm{separation},\ \mathrm{ETA}_i - t_{\mathrm{early}}\right), \tag{7}$$

where *separation* is derived from the runway capacity, and in this simulation explicitly does not include separation

constraints for different wake turbulence category aircraft. t_{early} represents the configured early-planning bias which is set at 60 seconds.

The Expected Approach Time (EAT) is directly derived from the assigned slot, accounting for the predicted Terminal Maneuvering Area (TMA) transit time:

$$\text{EAT}_i = \text{slot}_i - TMA_i. \tag{8}$$

Aircraft initially receive a *preplanned* slot. In this state, slots remain flexible and may be updated following trajectory predictor revisions. Once the predicted time-to-go to the IAF falls below the freeze horizon, the aircraft becomes eligible for planning within the frozen sequence and the slot transitions to the *frozen* state. In the baseline configuration representing current Schiphol operations, this occurs approximately 14 minutes prior to the predicted ETO IAF [4].

Aircraft that become visible to the AMAN only after entering the planning horizon, and are immediately within the freeze horizon, are classified as pop-up flights. Since no flexible planning phase is available, these flights are inserted directly into the frozen sequence, often requiring re-planning of already assigned slots [7].

Arrival planning is performed by Approach Control (APP), which is responsible for managing the AMAN. Using the AMAN, APP ensures that inbound traffic is sequenced in an orderly manner and that delay is absorbed as much as possible outside of the Terminal Maneuvering Area (TMA), maintaining manageable traffic conditions within the TMA [4], [9].

Once a slot is frozen, the corresponding Expected Approach Time (EAT) is communicated to Area Control (ACC). ACC then issues instructions to the aircraft to ensure that the predicted time over the IAF remains within the defined approach window, thereby maintaining feasibility of the assigned slot. Slot changes beyond this point occur only when strictly necessary.

The approach window is the time window around the EAT in which each flight is required to arrive at the IAF and enter the TMA. The current approach window used at Schiphol and this simulation is $[-120s, 120s]$.

E. Aircraft Visibility and Radar Representation

For computational efficiency, simulations are limited to a 500 NM radius surrounding the destination airport. Aircraft outside this boundary are not simulated.

In current operations at Schiphol, the AMAN relies on correlated radar data to estimate arrival times of inbound flights and construct the landing sequence. Since radar coverage is limited and strongly dependent on aircraft altitude, aircraft are not visible immediately after departure. Radar range increases with altitude, meaning that flights from nearby airports typically become visible during the climb phase.

To represent this mechanism in BAMS, an airport-dependent radar visibility model is implemented. In general, aircraft are assumed to become visible at FL130, reflecting the increase in radar coverage with altitude, with lower thresholds applied for specific nearby departure airports, as shown in Table I. Aircraft below the corresponding visibility threshold are not yet included in the AMAN planning, even if they are already within the planning horizon. This introduces a realistic source of pop-up behavior, where aircraft become eligible for planning only after satisfying the radar visibility constraint.

Recent operational developments at Schiphol have introduced the use of EFD messages to provide earlier awareness of flights that are airborne but not yet visible on radar. These flights can be pre-planned based on their predicted trajectory, but are not yet actively included in the sequence. Once the aircraft becomes visible to the radar system, it is still considered a pop-up. The information primarily serves as situational awareness for the controller.

The potential use of this information as an active planning input is of particular interest. Instead of using EFD data solely for situational awareness, the airborne status can be used to estimate the ETA and include the flight directly in the AMAN sequence prior to radar visibility. This effectively shifts the planning trigger from radar detection to take-off, allowing nearby departures to be incorporated earlier and potentially reducing the amount and impact of pop-ups. However, this also introduces additional uncertainty, particularly from the departure route phase. In the simulation, this concept is implemented by allowing flights to become eligible for planning at take-off, rather than at the radar-based altitude thresholds shown in Table I.

Table I. Departure-specific radar visibility altitude thresholds used in the simulation. Flights become eligible for AMAN planning only after crossing the listed altitude, introducing airport-dependent pop-up behavior.

Departure Airport	Visibility Altitude (ft)
Norwich (EGSH)	10,000
London Stansted (EGSS)	10,000
London City (EGLC)	10,000
Düsseldorf (EDDL)	2,000
Bremen (EDDW)	2,000
Brussels (EBBR)	1,000
Hannover (EDDV)	5,000
Other	13,000

F. ATC Instruction Logic

To ensure that aircraft adhere to their assigned slot, ATC instructions are issued based on the EAT and the corresponding approach window. In BAMS, this logic is implemented to closely reflect operational practice. Each aircraft is required to cross the Initial Approach Fix (IAF) within this window, which at Schiphol is defined as [-120, 120] seconds around the EAT. As long as the predicted Estimated Time Over the IAF (ETO IAF) remains within this window, no action is taken. If the predicted time over the IAF falls outside the window, an instruction is issued to adjust the trajectory such that the predicted time falls within the approach window, ensuring that the assigned slot remains feasible.

Instruction selection depends on the assumed control authority, modeled through a handover altitude proxy. Below Flight Level (FL) 260, aircraft are considered under local Area Control Center (ACC) responsibility. Above FL260,

aircraft are assumed to be under adjacent center control. Interventions are executed preferentially when aircraft are under local control, as this requires less coordination between control centers.

Instruction selection follows a prioritization consistent with operational controller strategies. For *delay absorption*, the following hierarchy is applied:

1) **Speed reduction** – Preferred method. The maximum possible reduction is assumed to be 50 knots indicated airspeed (KIAS).
2) **Adjacent center speed reduction** – Applied when in-sector speed reduction is insufficient. A standard Mach reduction of 0.04 is used, or a maximum speed reduction of 50 KIAS.
3) **Dogleg vectoring** – Introduced when speed-based mechanisms are insufficient. Increases track miles and controller taskload. The maneuver is modeled by inserting an additional waypoint into the route. The maximum allowable trajectory extension is limited to 30% of the remaining distance to the IAF.
4) **Holding pattern** – Least preferred option. Minimum holding time is assumed to be 120 seconds.

For *time recovery*, the following hierarchy is applied:

1) **Direct-to clearances** – Most favorable, as they reduce path length.
2) **Speed increase** – Applied when route shortening is unavailable or insufficient. The maximum possible increase is assumed to be 25 knots IAS.
3) **Adjacent center speed increase** – Similar to adjacent speed reduction. No Mach increase is assumed, and the maximum possible increase is limited to 25 knots IAS.
4) **Re-planning** – Required only when tactical measures cannot recover sufficient time to meet slot constraints. The aircraft exchanges slots with the subsequent aircraft in the sequence when feasible.

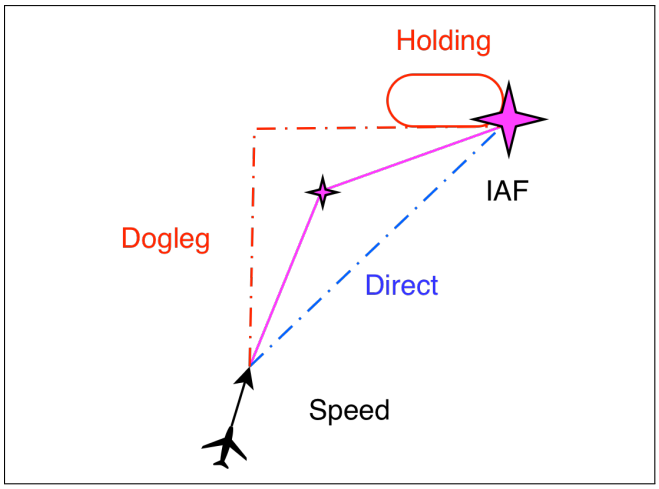

Figure 5. Illustrative solution space for delay absorption and time recovery in the simulated ATC logic. The figure summarizes the tactical control options available to absorb delay or recover time before the IAF.

Speed control is favored because it maintains lateral predictability and is assumed most fuel efficient. Vectoring and holding, while effective, introduce measurable inefficiency and increased workload. The objective of each issued instruction is to guide the aircraft toward the center of the defined approach window. If the maximum amount of an instruction type is used, and the flight adheres to the approach window, no additional instruction is given. A schematic overview is given in Fig. 5.

The ATC logic does not consider separation constraints between aircraft, due to the immense complexity of managing airspace to remain deconflicted.

G. Applicability to different airport environments

Although BAMS is configured to represent Schiphol operations, the simulation framework is intentionally designed as a generic AMAN evaluation environment. Airport-specific behavior is not hard-coded but defined through configurable inputs describing the operational context. These inputs include the TMA route structure, IAF definitions, runway configuration, separation minima, freeze horizon definition, and traffic demand composition. By modifying these parameters, the same AMAN and ATC logic can be applied to different airport environments while preserving methodological consistency. This enables comparative studies of AMAN concepts without altering the underlying simulation mechanics.

IV. Experiment Design

This section describes how the simulation experiments were structured to evaluate the effect of freeze horizon extension, uncertainty representation, and pop-up mitigation strategies on (Extended) Arrival Manager performance.

A. Independent Variables

The experiment evaluates (Extended) Arrival Manager performance as a function of three independent variables: uncertainty representation, freeze horizon (FH), and pop-up mitigation strategy.

Uncertainty representation. To isolate the effects of uncertainty sources, the stochastic uncertainty model is varied systematically. Uncertainty components are either enabled or disabled while preserving identical traffic demand and planning logic. Four uncertainty configurations are considered:

- *No uncertainty* – All stochastic uncertainty sources are disabled. Aircraft trajectories are deterministic. Pop-up flights are still present in the traffic set, but are preplanned and therefore do not appear as late insertions.
- *No pop-up insertions (preplanned)* – The same traffic mix is used, but pop-up flights are preplanned using pre-departure ETAs, thereby removing the late-insertion mechanism while keeping trajectory prediction uncertainty active.
- *No trajectory prediction (TP) uncertainty* – En-route and departure-route trajectory prediction uncertainty are disabled. Take-off uncertainty and pop-up behavior remain active.
- *All uncertainty* – All modeled stochastic uncertainty sources are enabled. This configuration represents the full operational uncertainty environment.

This decomposition enables attribution of performance changes to specific uncertainty mechanisms. Importantly, the traffic demand remains unchanged across these configurations, such that the comparison isolates uncertainty effects rather than scenario differences.

Freeze horizon (FH). Three freeze horizon values are evaluated: 14, 20, and 25 minutes. The 14-minute freeze horizon represents the baseline configuration and is equal to the current operational freeze horizon. The 20-minute freeze horizon is selected as the primary Extended AMAN configuration, as it represents a realistic horizon extension target for Schiphol operations. The 25-minute freeze horizon is included as a further-extended case to assess whether the evaluated mitigation strategies remain effective under more demanding uncertainty conditions.

Pop-up mitigation strategy. Three planning strategies are evaluated:

- *First-Come-First-Serve (FCFS, baseline)* – Pop-up flights are inserted into the existing sequence based on their Estimated Time of Arrival (ETA). In case of a pop-up flight, the frozen sequence requires partial re-planning. This reflects the default behavior where pop-ups are integrated according to their predicted arrival time.
- *Back-of-the-Line (BOL)* – Pop-up flights are appended to the end of the frozen part of the sequence. This keeps the frozen sequence intact.
- *Delayed-slot strategy (Delay)* – Pop-up flights are pre-planned while still on-ground using an expected delay offset derived from the uncertainty statistics. The offset used to set the slot is the 75th percentile delay time in the respective take-off planning distribution, specific to airport and prediction lead time. If departure occurs before the delayed slot is frozen, the back-of-the-line strategy is applied instead. If the flight is not able to reach the IAF within the assigned approach window, the sequence is re-planned accordingly. Each pop-up flight is preplanned only once; if the initially assigned delayed slot is not adhered to, the flight is not preplanned again.

The delayed-slot strategy is included because it explicitly accounts for uncertainty in pop-up flight planning. The 75th percentile was selected as a conservative but not overly restrictive buffer, balancing missed-slot probability against added delay. Together, these three strategies represent a reactive baseline logic (FCFS), a stability-preserving reactive alternative (BOL), and a proactive uncertainty-aware strategy (Delay).

In addition to these three main strategies, two supplementary mitigation measures are evaluated. First, the effect of *planning at take-off* is evaluated. In this condition, flights become eligible for AMAN planning and freezing immediately after take-off, rather than only once they satisfy the radar visibility altitude described earlier. This corresponds to using EFD information not only for situational awareness, but as an active input to the AMAN planning process. The airborne status is used to estimate the ETA and include the flight directly in the sequence prior to radar visibility. This effectively shifts the planning trigger from radar correlation to take-off, allowing nearby departures to be incorporated into the frozen sequence earlier. The measure is evaluated to assess whether earlier inclusion reduces the amount and detrimental effects of pop-ups, while accounting for the additional uncertainty introduced in the departure phase.

Second, the specific influence of *Brussels pop-up traffic* is investigated by evaluating configurations in which pop-up flights originating from Brussels are excluded. These flights are expected to have a relatively strong effect on E-AMAN performance due to their proximity to Schiphol and their tendency to appear late in the planning process. This comparison is included to assess how strongly the instability is driven by this particular traffic source. Although such a measure is conceptually simple, it is not under direct LVNL control and would involve broader operational and commercial considerations.

B. Dependent Variables

Performance is evaluated using a fixed set of Key Performance Indicators (KPIs) extracted from the simulation output. For each scenario–configuration combination, KPIs are aggregated to a single mean value, yielding 12 paired observations per KPI used in the statistical comparisons. The KPIs are grouped according to the Key Performance Areas (KPAs) relevant for (E-)AMAN evaluation.

Stability

- *EAT revisions (–)*: the total number of EAT updates within a scenario. Each revision corresponds to a change in the EAT of an aircraft of at least 10 seconds.
- *Sequence position changes (–)*: the number of sequence disruptions caused by slot insertion or replanning, counted as position changes affecting aircraft already in the frozen sequence and replanned pop-up flights.

Planning effectiveness

- *Delay energy cost (%)*: the additional energy consumption relative to the nominal planned trajectory, capturing inefficiencies introduced by speed control, vectoring, and holding, with energy consumption is obtained from the BADA-based performance model implemented in BlueSky.

$$E_{\text{delaycost}} = 100 \cdot \left(\frac{\sum_i E_i^{\text{actual}}}{\sum_i E_i^{\text{reference}}} - 1 \right) \tag{9}$$

- *Mean total delay (s/ac)*: the cumulative delay absorption per aircraft required to meet the assigned EAT window.
- *Mean vectoring delay (s/ac)*: the delay absorbed via lateral path extension, such as holding or doglegs.
- *Vectored flights delay (ac)*: the number of aircraft affected by lateral path extension.

Taskload

- *Instruction count (–)*: the total number of ATC instructions issued in the scenario. Speed instructions require one instruction to the aircraft, a dogleg requires two, and a holding pattern, independent of holding time, is assumed to require four instructions.

- *Sequence position changes (–)*: also considered a taskload metric, as each sequence disruption may require additional controller intervention.

C. Controlled Variables

All configurations use identical traffic demand, runway use, wind conditions, trajectory prediction logic, separation parameters, aircraft performance assumptions, stochastic sampling framework, and ATC instruction behavior. In addition, the same set of scenario seeds is used across configurations, such that each configuration is evaluated under paired stochastic conditions. Consequently, performance differences can be attributed to the independent variables under investigation rather than to differences in traffic realization, environmental conditions, or controller logic.

D. Traffic Scenarios

To capture the stochastic nature of operational uncertainty under representative peak conditions, a Monte-Carlo experiment design is used. The study includes 12 traffic scenarios, each representing a 3-hour peak period with between 198 and 204 arriving aircraft. The scenarios vary in operational conditions, including wind settings, runway configuration, and traffic composition. Traffic demand corresponds to 34 arrivals per hour per runway, with two arrival runways in use (approximately 68 arrivals per hour in total), while the nominal runway capacity is set at 38 arrivals per hour per runway [17]. Aircraft are distributed over 15-minute intervals to create moderate bunching while maintaining a semi-steady-state arrival demand pattern.

Traffic composition includes a fixed share of departures from nearby airports, matched to representative operational percentages, as shown in Table II. This is particularly relevant for extended freeze horizons, as in- and near-horizon departures are a primary driver of pop-up behavior. Flights within the Dutch FIR are not considered, since no scheduled domestic airline flights occur. Business, repositioning and test flights almost always utilize the Oostbaan runway, and are usually kept outside of the sequence by ASAP.

Table II. Nearby Airport Share in Scenarios

Origin Airport	Flights	Share (%)
Brussels (EBBR)	2	0.98
Frankfurt (EDDF)	5	2.45
Düsseldorf (EDDL)	2	0.98
Bremen (EDDW)	1	0.49
London City (EGLC)	4	1.96
London Heathrow (EGLL)	5	2.45
Norwich (EGSH)	1	0.49
London Stansted (EGSS)	1	0.49
Luxembourg (ELLX)	1	0.49

Flights including their routes and altitudes are obtained from the Eurocontrol Demand Data Repository (DDR2) database [18]. Departure times are adjusted to ensure stable arrival demand. Standard speed scheduling uses ASAP performance data for climb, cruise Mach, and descent, differentiated by aircraft type and airline. Each aircraft passes the IAF at FL100 and flies 250 KIAS below FL100. As flight-path separation is not modeled, no departing traffic is simulated.

E. Monte-Carlo

Each scenario is simulated using 48 seeds per configuration for stochastic error sampling. This number was selected as a practical balance between variance reduction and total experiment size. For every configuration, the seeds are varied across the corresponding scenarios, while seeds are kept identical between configurations. This produces paired stochastic realizations, enabling direct comparison between configurations under identical random conditions. For each scenario–configuration combination, Key Performance Indicators (KPIs) are aggregated into mean values. This results in 12 paired observations per KPI for each configuration.

Prior to hypothesis testing, KPI distributions are standardized using Z-scores to facilitate cross-metric comparability and visual interpretation of relative effect magnitude. Differences between configurations are evaluated using Friedman's Analysis of Variance (ANOVA) by ranks, which is suitable for the repeated-measures design and does not require normality. The test statistics include the chi-square value (χ^2), which reflects the magnitude of differences between configurations, the p-value (p), which indicates whether these differences are statistically significant, and Kendall's coefficient of concordance (W), which represents the effect size and consistency of the observed differences [19]. Paired Wilcoxon signed-rank post-hoc tests are used to identify pairwise differences between configurations [7].

To control the family-wise error rate arising from multiple post-hoc comparisons, Bonferroni correction is applied within each comparison family. The corrected significance threshold is defined as

$$\alpha_{\mathrm{corr}} = \frac{0.05}{m}, \tag{10}$$

where m is the number of post-hoc comparisons considered in the corresponding evaluation. This yields, for example, $\alpha_{\mathrm{corr}} = 0.0167$ for comparisons involving three configurations and $\alpha_{\mathrm{corr}} = 0.0083$ for comparisons involving four configurations.

V. Results

This section presents the results for the evaluated freeze horizons, uncertainty configurations, scheduler strategies, and supplementary mitigation measures related to planning at take-off and Brussels pop-up traffic. Z-score plots are included to visualize the relative separation between configurations across KPIs. Friedman's ANOVA is used to assess whether statistically significant differences exist across configurations, followed by post-hoc tests to identify pairwise differences. Interpretation of the results is provided in the Discussion section.

A. Effect of Horizon Extension

The effect of freeze horizon extension was evaluated as an independent variable with three levels: 14, 20, and 25 minutes.

Fig. 6 shows that the largest differences occur in the stability-related KPIs. Both EAT revisions and sequence position changes increase clearly from FH14 to FH20 and FH25, with a clear separation between configurations.

Table III. Comparison of FH14, FH20, and FH25 under full uncertainty and FCFS scheduling.

KPI	Configuration		
Scheduler	FCFS	FCFS	FCFS
Freeze horizon [min]	14	20	25
Uncertainty	all	all	all
EAT Revisions [-]	1.422	13.318*	55.092*
Sequence Position Changes [-]	1.427	12.694*	45.057*
Delay energy cost [%]	0.212	0.191	0.104
Mean total delay [s/ac]	116.458	111.638	109.088
Mean vectoring delay [s/ac]	15.209	10.361*	9.771*
Vectored flights [ac]	20.660	14.538*	13.781*
Instruction count [-]	150.858	147.167	161.720

a* indicates significant difference versus FH14 (Wilcoxon).
Significance threshold: $\alpha = 0.0167$ (Bonferroni).

Table III confirms these trends and indicates that the increases in EAT revisions and sequence position changes are statistically significant between configurations. In contrast, vectoring-related KPIs decrease with horizon extension, and the number of vectored aircraft follows the same trend. No statistically significant differences are found for mean total delay. Planning effectiveness KPIs show smaller differences overall, consistent with the more limited separation observed in the Z-score figure.

Table IV indicates that freeze horizon length has a statistically significant effect on all KPIs except mean total delay, with the strongest effects observed for the stability-related metrics.

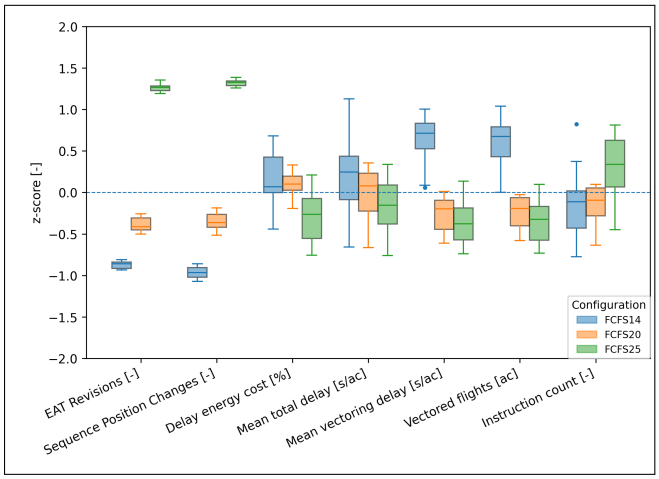

Figure 6. Z-score overview of the freeze horizon extension comparison. KPI values are standardized per metric to show the relative separation between FH14, FH20, and FH25 across stability, planning effectiveness , and taskload.

B. Effect of Stochastic Uncertainties

The influence of stochastic uncertainty was evaluated as an independent variable with four configurations: no uncertainty, no pop-up insertions (preplanned pop-ups), no trajectory prediction (TP) uncertainty, and all uncertainty.

Fig. 7 shows that the largest differences occur in the stability-related KPIs. EAT revisions increase strongly when trajectory prediction uncertainty is introduced, while sequence position changes increase primarily when pop-ups

Table IV. Friedman ANOVA results for the freeze horizon extension comparison. Per KPI, the table reports the test statistic, p-value, and effect size for the comparison between FH14, FH20, and FH25.

KPI	χ^2	p	W	Sig.
EAT Revisions [-]	24.000	0.000	1.000	Yes
Sequence Position Changes [-]	24.000	0.000	1.000	Yes
Delay Energy Cost [%]	10.167	0.006	0.424	Yes
Mean Total Delay [s/ac]	2.167	0.338	0.090	No
Mean Vectoring Delay [s/ac]	15.167	0.001	0.632	Yes
Vectored Flights [ac]	15.167	0.001	0.632	Yes
Instruction Count [-]	6.167	0.046	0.257	Yes

$n = 12$; $k = 3$; $df = 2$; $\alpha = 0.050$

are present. Configurations without pop-ups show clearly lower sequence disruption. Planning effectiveness KPIs show smaller differences overall. Delay energy cost increases when TP uncertainty is introduced, while vectoring-related KPIs remain relatively clustered. Mean total delay shows some variation between configurations, but the magnitude of these differences is limited compared to the stability KPIs. Taskload metrics follow the stability pattern, with higher instruction counts observed when uncertainty mechanisms are active.

Table V confirms these trends and indicates that EAT revisions differ significantly between configurations with and without trajectory prediction uncertainty, while sequence position changes are significantly lower without pop-ups. Delay energy cost also differs significantly across configurations, while mean vectoring delay does not show statistically significant differences.

Table VI indicates statistically significant effects for all KPIs except mean vectoring delay, with the largest effect sizes observed for EAT revisions, sequence position changes, and delay energy cost.

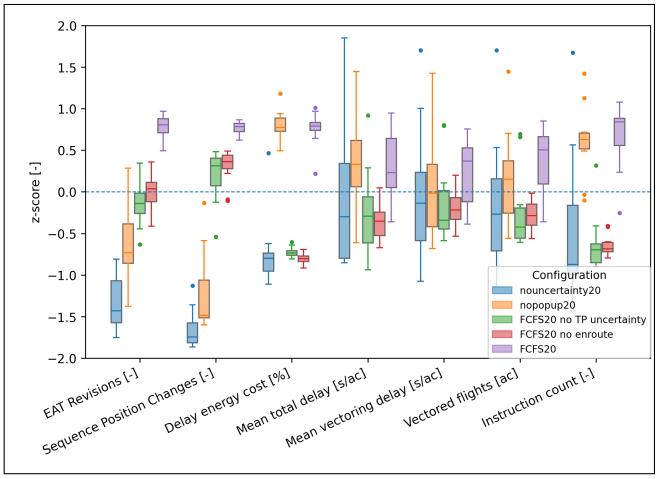

Figure 7. Z-score overview of the uncertainty comparison. KPI values are standardized per metric to show the relative separation between the evaluated uncertainty configurations across stability, planning effectiveness , and taskload.

C. Evaluation of Pop-Up Mitigation Strategies

The effect of pop-up mitigation strategy was evaluated as an independent variable with three levels: First-Come-First-

Table V. Comparison of uncertainty configurations at FH20 with FCFS scheduling.

KPI	Configuration				
Scheduler	FCFS				
Freeze horizon [min]	20				
Uncertainty	no uncertainty	no pop-ups	no TP uncertainty	no en-route uncertainty	all uncertainty
EAT Revisions [-]	0.000	4.620*	7.318*	8.080*	13.215*
Sequence Position Changes [-]	0.000	2.010*	9.946*	10.431*	12.694*
Delay energy cost [%]	-0.177	0.219*	-0.197	-0.216	0.191*
Mean total delay [s/ac]	116.001	118.500	104.958	103.960	111.682
Mean vectoring delay [s/ac]	11.199	11.343	9.192	9.015	10.370
Vectored flights [ac]	14.250	15.451	12.181	11.905	14.552
Instruction count [-]	132.500	157.149*	112.387	111.556	147.220

a* indicates significant difference versus no uncertainty (paired Wilcoxon).
Significance threshold: $\alpha = 0.005$ (Bonferroni corrected).

Table VI. Friedman ANOVA results for the uncertainty comparison. Per KPI, the table reports the test statistic, p-value, and effect size for the evaluated uncertainty configurations.

KPI	χ^2	p	W	Sig.
EAT Revisions [-]	41.867	0.000	0.872	Yes
Sequence Position Changes [-]	46.310	0.000	0.965	Yes
Delay energy cost [%]	36.200	0.000	0.754	Yes
Mean total delay [s/ac]	12.000	0.017	0.250	Yes
Mean vectoring delay [s/ac]	6.200	0.185	0.129	No
Vectored flights [ac]	13.600	0.009	0.283	Yes
Instruction count [-]	29.067	0.000	0.606	Yes

$n = 12$; $k = 5$; $df = 4$; $\alpha = 0.050$

Serve (FCFS), Back-of-the-Line (BOL), and the delayed-slot strategy (Delay), at a fixed freeze horizon of 20 minutes.

Fig. 8 shows that the largest differences occur in the stability-related KPIs. BOL results in substantially fewer EAT revisions and sequence position changes compared to FCFS, indicating improved planning stability. The Delay scheduler also reduces instability compared to FCFS, but remains consistently less stable than BOL. Planning effectiveness KPIs show smaller differences across schedulers. BOL introduces a slight increase in mean total delay relative to FCFS, while the Delay scheduler results in the highest delay levels. Vectoring-related KPIs remain relatively clustered, indicating limited sensitivity to the scheduling strategy. Taskload follows the stability pattern, with higher instruction counts observed for both BOL and Delay compared to FCFS.

Table VII confirms these trends and indicates that EAT revisions and sequence position changes are significantly lower for BOL compared to FCFS, while the Delay scheduler shows intermediate performance. Mean total delay is higher for the Delay scheduler, while differences between BOL and FCFS are not statistically significant. Vectoring-related KPIs do not show statistically significant differences.

Table VIII indicates statistically significant effects for the stability KPIs, mean total delay, and instruction count, with the strongest effects observed for the stability-related metrics.

1) Evaluation of Pop-Up Mitigation Strategies With a Further Extended Horizon: The effect of pop-up mitigation strategy was evaluated at a further extended freeze horizon of 25 minutes, comparing the same three levels: First-Come-First-Serve (FCFS), Back-of-the-Line (BOL), and the delayed-slot strategy (Delay).

Fig. 9 shows that the separation between configurations

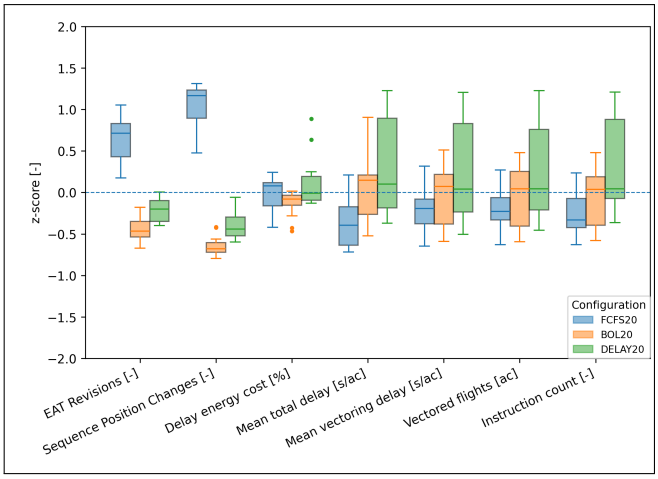

Figure 8. Z-score overview of the scheduler comparison at FH20. KPI values are standardized per metric to show the relative separation between FCFS, BOL, and Delay schedulers across stability, planning effectiveness , and taskload.

Table VII. Comparison of FCFS, BOL, and Delay scheduling at FH20 under full uncertainty.

KPI	Configuration		
Scheduler	FCFS	BOL	Delay
Freeze horizon [min]	20		
EAT Revisions [-]	13.318	4.741*	6.674*
Sequence Position Changes [-]	12.694	2.135*	3.632*
Delay energy cost [%]	0.191	0.160*	0.253
Mean total delay [s/ac]	111.638	117.821*	129.495
Mean vectoring delay [s/ac]	10.361	10.846	13.761
Vectored flights [ac]	14.538	15.193	19.035
Instruction count [-]	147.167	151.693	178.378

a* indicates significant difference versus FCFS (Wilcoxon).
Significance threshold: $\alpha = 0.0167$ (Bonferroni).

increases compared to FH20, particularly for the stability-related KPIs. Both EAT revisions and sequence position changes increase substantially in magnitude, and the differences between FCFS and the alternative strategies become more pronounced.

Planning effectiveness KPIs also show increased separation compared to FH20, although the differences remain smaller than for the stability metrics. Delay-related KPIs indicate that the Delay scheduler results in the highest delay levels, while

Table VIII. Friedman ANOVA results for the scheduler comparison at FH20. Per KPI, the table reports the test statistic, p-value, and effect size for FCFS, BOL, and Delay schedulers.

KPI	χ^2	p	W	Sig.
EAT Revisions [-]	24.000	0.000	1.000	Yes
Sequence Position Changes [-]	24.000	0.000	1.000	Yes
Delay energy cost [%]	4.167	0.125	0.174	No
Mean total delay [s/ac]	8.167	0.017	0.340	Yes
Mean vectoring delay [s/ac]	3.500	0.174	0.146	No
Vectored flights [ac]	3.500	0.174	0.146	No
Instruction count [-]	6.500	0.039	0.271	Yes

$n = 12$; $k = 3$; $df = 2$; $\alpha = 0.050$

BOL remains closer to FCFS. Taskload follows the same pattern as stability, with higher instruction counts for BOL and Delay relative to FCFS. At the same time, the differences between BOL and Delay remain limited for several KPIs.

Table IX confirms these trends and indicates that differences between FCFS and BOL are statistically significant across all KPIs, while differences between FCFS and Delay are significant for most KPIs, except delay energy cost.

Table X show that the scheduling strategy has a strong and consistent effect on stability-related KPIs, while effects on planning effectiveness and taskload are present but generally more moderate.

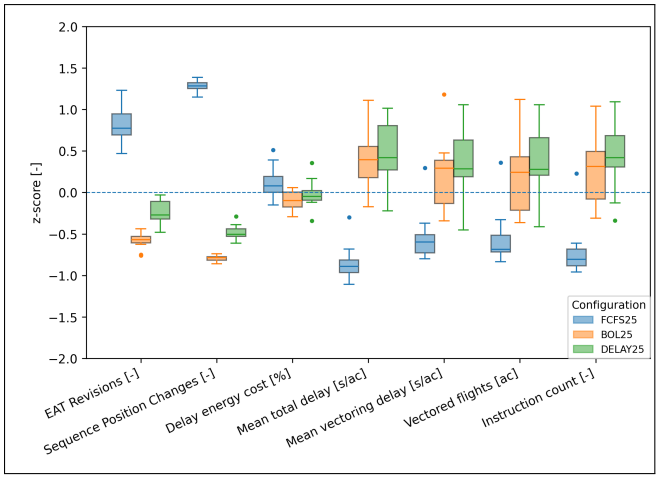

Figure 9. Z-score overview of the scheduler comparison at FH25. KPI values are standardized per metric to show the relative separation between FCFS, BOL, and Delay schedulers across stability, planning effectiveness , and taskload.

Table IX. Comparison of FCFS, BOL, and Delay scheduling at FH25 under full uncertainty.

KPI	**Configuration**		
Scheduler	FCFS	BOL	Delay
Freeze horizon [min]	25		
EAT Revisions [-]	55.092	25.191*	31.868*
Sequence Position Changes [-]	45.057	6.931*	12.175*
Delay energy cost [%]	0.104	0.047*	0.069
Mean total delay [s/ac]	109.088	133.382*	135.068*
Mean vectoring delay [s/ac]	9.771	13.626*	13.853*
Vectored flights [ac]	13.781	18.330*	19.293*
Instruction count [-]	161.720	196.420*	202.946*

a* indicates significant difference versus FCFS (Wilcoxon).
Significance threshold: $\alpha = 0.0167$ (Bonferroni).

2) Comparison to Baseline Operations: The comparison with baseline operations corresponds to the experiment matrix evaluation of scheduling strategy and freeze horizon, comparing FCFS14 with BOL20 and Delay20.

Fig. 10 shows that the clearest differences occur for the stability- and vectoring-related KPIs. Stability metrics increase relative to baseline AMAN, reflecting the earlier planning characteristic of Extended AMAN, although the increase in sequence position changes between FCFS14 and BOL20 is not significant. In addition, BOL20 reduces vectoring delay relative to FCFS14. Across most KPIs, the Delay configuration consistently shows less favorable performance compared to both FCFS and BOL.

Table X. Friedman ANOVA results for the evaluation of pop-up mitigation strategies under a further extended freeze horizon, reporting the Friedman test statistic (χ^2), p-value, and effect size per KPI.

KPI	χ^2	**p**	**W**	**Sig.**
EAT Revisions [-]	24.000	0.000	1.000	Yes
Sequence Position Changes [-]	24.000	0.000	1.000	Yes
Delay energy cost [%]	6.167	0.046	0.257	Yes
Mean total delay [s/ac]	18.000	0.000	0.750	Yes
Mean vectoring delay [s/ac]	12.667	0.002	0.528	Yes
Vectored flights [ac]	12.500	0.002	0.521	Yes
Instruction count [-]	12.667	0.002	0.528	Yes

$n = 12$; $k = 3$; $df = 2$; $\alpha = 0.050$

Table XI confirms these observations and indicates that the reduction in vectoring delay is statistically significant, while mean total delay remains statistically unchanged.

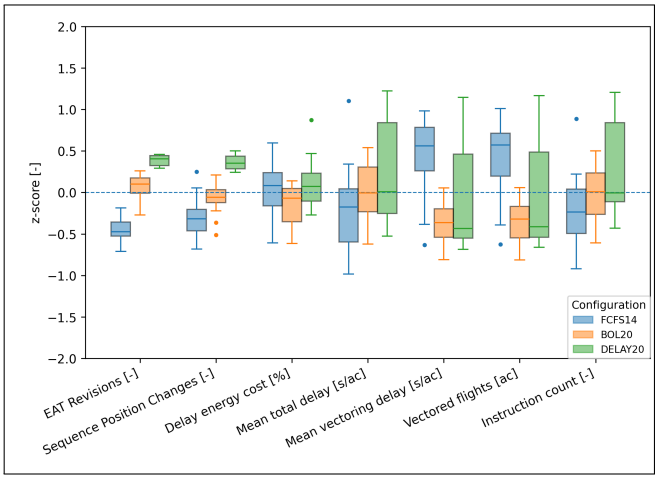

Figure 10. Z-score overview of the comparison between baseline AMAN and the evaluated E-AMAN scheduler configurations. KPI values are standardized per metric to show how the configurations differ across stability, planning effectiveness , and taskload.

Table XI. Comparison of baseline AMAN (FCFS14) and the evaluated E-AMAN scheduler configurations.

KPI	**Configuration**		
Scheduler	FCFS	BOL	Delay
Freeze horizon [min]	14	20	20
EAT Revisions [-]	1.422	4.741*	6.674*
Sequence Position Changes [-]	1.427	2.135	3.632*
Delay energy cost [%]	0.212	0.160	0.253
Mean total delay [s/ac]	116.458	117.821	129.495
Mean vectoring delay [s/ac]	15.209	10.846*	13.761
Vectored flights [ac]	20.660	15.193*	19.035
Instruction count [-]	150.858	151.693	178.378

a* indicates significant difference versus FCFS14 (Wilcoxon).
Significance threshold: $\alpha = 0.0167$ (Bonferroni).

Fig. 11 illustrates the distribution of delay absorption mechanisms. The figure shows that BOL20 shifts delay absorption upstream compared to FCFS14, with a larger share of delay handled through speed control rather than vectoring or holding. Delay20 shows an increase in total delay.

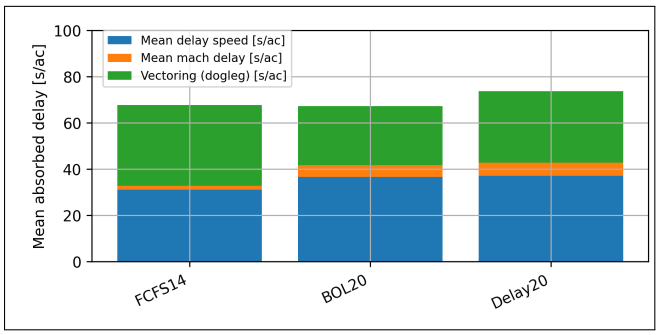

Figure 11. Cumulative delay absorption by mechanism for baseline AMAN and the evaluated E-AMAN scheduler configurations. The figure shows a comparison between schedulers in terms of total delay and delay composition.

D. Effects of Planning at Take-Off and Brussels Pop-Up Traffic

1) Baseline (FCFS) Scheduler: The effect of supplementary mitigation measures was evaluated under FCFS scheduling, considering three configurations: planning at take-off using EFD messages, exclusion of Brussels pop-up flights, and the combination of both measures.

Fig. 12 shows that the largest differences occur in the stability-related KPIs. Planning at take-off reduces EAT revisions and sequence position changes, while excluding Brussels pop-up flights produces a similar reduction in sequence disruption. The combined configuration shows the strongest overall improvement in stability. Planning effectiveness KPIs show more moderate changes. Planning at take-off reduces vectoring-related KPIs, including mean vectoring delay and the number of vectored flights, while excluding Brussels pop-up flights has a smaller effect on these metrics. Taskload follows the stability pattern, with reductions in instruction count associated with improved sequence stability.

Table XII confirms these trends and indicates that planning at take-off significantly reduces both stability- and vectoring-related KPIs. Excluding Brussels pop-up flights significantly improves stability, while effects on other KPIs remain mostly non-significant. The combined configuration shows improvements across most KPIs.

Table XIII indicates statistically significant effects across all KPIs except mean total delay, with the strongest effects observed for the stability-related metrics.

2) Back-of-the-Line Scheduler: The effect of additional measures was evaluated under Back-of-the-Line (BOL) scheduling, considering three configurations: planning at take-off using EFD messages, exclusion of Brussels pop-up flights, and the combination of both measures.

Fig. 13 shows limited separation between configurations for the stability-related KPIs, indicating that neither planning at take-off nor excluding Brussels pop-up flights substantially affects sequence stability under BOL scheduling. Clearer differences are observed for planning effectiveness and taskload. Excluding Brussels pop-up flights reduces delay-related KPIs, while planning at take-off has little visible effect. The combined configuration follows the same pattern.

Table XIV confirms these trends and indicates that planning at take-off does not result in statistically significant

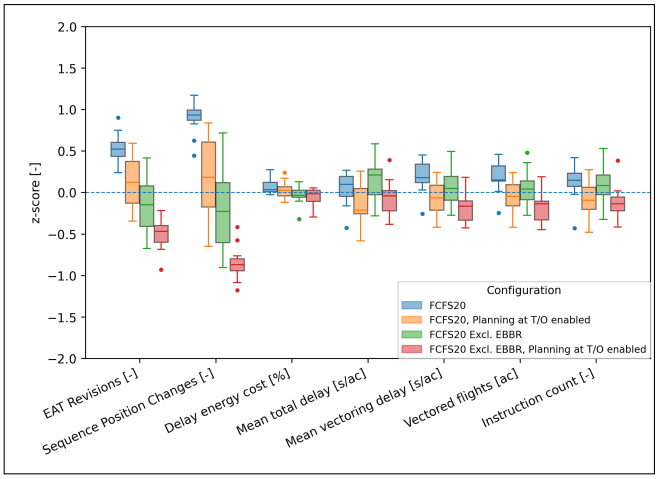

Figure 12. Z-score overview of the FCFS comparison with planning at take-off and Brussels exclusion. KPI values are standardized per metric to show the relative effect of the two supplementary mitigation measures and their combination across stability, planning effectiveness, and taskload.

Table XII. Comparison of planning at take-off and Brussels pop-up exclusion for E-AMAN with FCFS scheduling.

KPI	**Configuration**			
	C1	**C2**	**C3**	**C4**
Scheduler	FCFS			
Freeze horizon [min]	20			
Brussels pop-ups	Included	Included	Excluded	Excluded
Planning at take-off	Disabled	Enabled	Disabled	Enabled
EAT Revisions [-]	13.215	10.280*	8.186*	5.594*
Sequence Position Changes [-]	12.694	8.920*	6.198*	2.858*
Delay energy cost [%]	0.191	0.182	0.167	0.161
Mean total delay [s/ac]	111.682	109.154	112.486	109.793
Mean vectoring delay [s/ac]	10.370	9.677*	9.860	9.096
Vectored flights [ac]	14.552	13.651*	13.842	12.854
Instruction count [-]	147.220	142.627	145.163	140.234

a* indicates significant difference versus C1 (Wilcoxon).
Significance threshold: $\alpha = 0.0083$ (Bonferroni).

Table XIII. Friedman ANOVA results for the FCFS comparison with planning at take-off and Brussels exclusion. Per KPI, the table reports the test statistic, p-value, and effect size for the evaluated configurations.

KPI	χ^2	**p**	**W**	**Sig.**
EAT Revisions [-]	26.700	0.000	0.742	Yes
Sequence Position Changes [-]	31.300	0.000	0.869	Yes
Delay energy cost [%]	8.700	0.034	0.242	Yes
Mean total delay [s/ac]	6.100	0.107	0.169	No
Mean vectoring delay [s/ac]	19.500	0.000	0.542	Yes
Vectored flights [ac]	18.500	0.000	0.514	Yes
Instruction count [-]	13.000	0.005	0.361	Yes

n=12; k=4; df=3; α=0.050

improvements, whereas excluding Brussels pop-up flights significantly reduces delay-related KPIs and instruction count. The combined configuration shows the same behavior.

Table XV indicates statistically significant effects for the delay- and taskload-related KPIs, while no significant overall effect is found for the stability metrics.

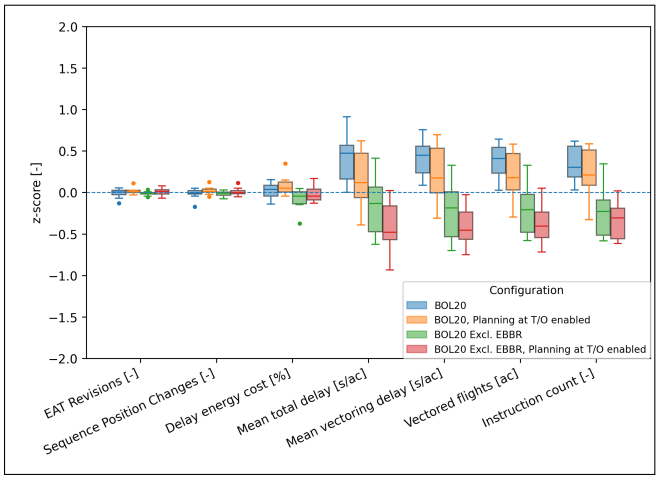

Figure 13. Z-score overview of the BOL comparison with planning at take-off and Brussels exclusion. KPI values are standardized per metric to show how the supplementary mitigation measures affect performance under BOL planning across stability, planning effectiveness , and taskload.

Table XIV. Comparison of planning at take-off and Brussels pop-up exclusion for E-AMAN with BOL scheduling.

KPI	Configuration			
	C1	C2	C3	C4
Scheduler	BOL	BOL	BOL	BOL
Freeze horizon [min]	20			
Brussels pop-ups	Included	Included	Excluded	Excluded
Planning at take-off	Disabled	Enabled	Disabled	Enabled
EAT Revisions [-]	4.816	4.953	4.795	4.910
Sequence Position Changes [-]	2.135	2.236	2.111	2.194
Delay energy cost [%]	0.161	0.173	0.141	0.153
Mean total delay [s/ac]	117.861	115.504	111.361*	109.038*
Mean vectoring delay [s/ac]	10.858	10.443	9.250*	8.832*
Vectored flights [ac]	15.207	14.712	13.000*	12.510*
Instruction count [-]	151.767	150.097	139.816*	138.349*

a* indicates significant difference versus C1 (Wilcoxon).
Significance threshold: $\alpha = 0.0083$ (Bonferroni).

3) Comparison to Baseline: The performance of two mitigation configurations was compared to the baseline E-AMAN: Back-of-the-Line (BOL) scheduling and FCFS with planning at take-off and exclusion of Brussels pop-up flights.

Fig. 14 shows that both configurations lead to substantial improvements in the stability-related KPIs, with clear reductions in EAT revisions and sequence position changes compared to the baseline.

For planning effectiveness, the two mitigation strategies show different behavior. The FCFS-based configuration with planning at take-off and exclusion of Brussels pop-up flights results in lower vectoring-related KPIs, including mean vectoring delay and the number of vectored flights. In contrast, BOL shows more limited improvements in these metrics.

Table XV. Effect of Planning at Take-off and Brussels Pop-up Exclusion for E-AMAN with BOL Planning — Friedman ANOVA Summary

KPI	χ^2	p	W	Sig.
EAT Revisions [-]	5.420	0.143	0.151	No
Sequence Position Changes [-]	2.822	0.420	0.078	No
Delay energy cost [%]	8.700	0.034	0.242	Yes
Mean total delay [s/ac]	18.500	0.000	0.514	Yes
Mean vectoring delay [s/ac]	23.900	0.000	0.664	Yes
Vectored flights [ac]	23.773	0.000	0.660	Yes
Instruction count [-]	21.800	0.000	0.606	Yes

n=12; k=4; df=3; α=0.050

Taskload follows a similar pattern, with the FCFS-based configuration showing lower instruction counts compared to BOL.

Table XVI confirms these trends and indicates that both configurations significantly improve stability compared to the baseline. The FCFS-based configuration additionally improves vectoring-related KPIs, while BOL shows smaller or non-significant differences for these measures.

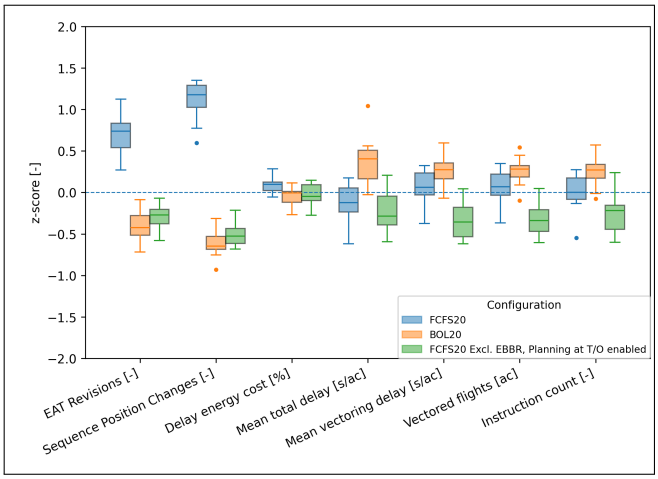

Figure 14. Z-score overview of the comparison between baseline E-AMAN and two alternative mitigation strategies. KPI values are standardized per metric to compare the relative performance of BOL and FCFS with planning at take-off and Brussels exclusion across stability, planning effectiveness , and taskload.

4) Impact of Take-Off Planning on Baseline AMAN Performance: The performance of the current baseline AMAN was compared with two configurations: AMAN with planning at take-off, and an E-AMAN configuration combining planning at take-off and exclusion of Brussels pop-up flights.

Fig. 15 shows that enabling planning at take-off improves stability relative to the baseline AMAN, with clear reductions in EAT revisions and sequence position changes. In contrast, the E-AMAN configuration shows reduced stability compared to the baseline, while improving planning effectiveness. In particular, vectoring-related KPIs, including mean vectoring delay and the number of vectored flights, are lower than in the baseline configuration. Taskload-related KPIs show a similar tendency, with a reduction in instruction count, although this effect is not statistically significant.

Table XVI. Comparison of baseline E-AMAN, BOL20, and FCFS20 with planning at take-off and Brussels pop-up exclusion.

KPI	Configuration		
Scheduler	FCFS	BOL	FCFS
Freeze horizon [min]	20		
Brussels pop-ups	Included	Included	Excluded
Planning at take-off	Disabled	Disabled	Enabled
EAT Revisions [-]	13.215	4.816*	5.594*
Sequence Position Changes [-]	12.694	2.135*	2.858*
Delay energy cost [%]	0.191	0.161*	0.161
Mean total delay [s/ac]	111.682	117.861*	109.793
Mean vectoring delay [s/ac]	10.370	10.858	9.096*
Vectored flights [ac]	14.552	15.207	12.854*
Instruction count [-]	147.220	151.767	140.234

a* indicates significant difference versus baseline (Wilcoxon).
Significance threshold: $\alpha = 0.0167$ (Bonferroni).

Table XVII. Comparison of baseline AMAN, AMAN with planning at take-off, and E-AMAN with planning at take-off and Brussels pop-up exclusion.

KPI	Configuration		
Scheduler	FCFS		
Freeze horizon [min]	14	14	20
Brussels pop-ups	Included	Included	Excluded
Planning at take-off	Disabled	Enabled	Enabled
EAT Revisions [-]	1.441	0.764*	5.594*
Sequence Position Changes [-]	1.427	0.693*	2.858*
Delay energy cost [%]	0.213	0.196	0.161
Mean total delay [s/ac]	116.465	115.449	109.793
Mean vectoring delay [s/ac]	15.208	14.944	9.096*
Vectored flights [ac]	20.653	20.276	12.854*
Instruction count [-]	150.887	148.767	140.234

a* indicates significant difference versus baseline (Wilcoxon).
Significance threshold: $\alpha = 0.0167$ (Bonferroni).

Table XVII confirms these trends and indicates that planning at take-off significantly improves stability compared to the baseline AMAN. The E-AMAN configuration shows significantly lower stability, while achieving significant reductions in vectoring-related KPIs. Improvements in delay energy cost and instruction count are observed but are not statistically significant.

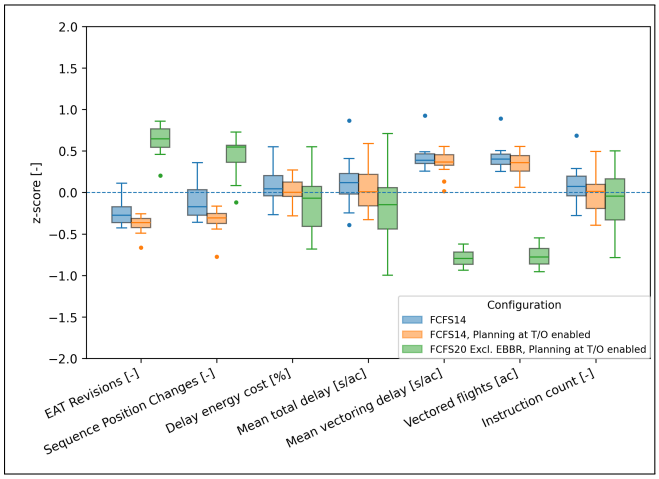

Figure 15. Z-score overview of the comparison with baseline AMAN. KPI values are standardized per metric to compare current AMAN, AMAN with planning at take-off, and E-AMAN with planning at take-off and Brussels exclusion across stability, planning effectiveness , and taskload.

VI. Discussion

Table XVIII provides a qualitative summary of the main findings across all experiments. The table is derived from the statistical results presented in Section V and indicates the direction and significance of differences relative to the reference configuration within each comparison.

A. Effect of Horizon Extension

The freeze horizon extension comparison highlights the fundamental trade-off between earlier planning and increased exposure to uncertainty. Extending the freeze horizon from 14 to 20 minutes produces a measurable increase in planning instability, primarily reflected in higher numbers of EAT revisions and sequence position changes. This behavior is consistent with the operational interpretation of Extended AMAN: earlier slot allocation inevitably incorporates predictions with larger uncertainty envelopes.

Further extension to 25 minutes amplifies this effect. Although the theoretical solution space for delay absorption increases, the longer prediction lookahead causes uncertainty propagation to dominate system behavior. The resulting growth in EAT revisions and sequence position changes indicates diminishing returns of horizon extension. Beyond a certain threshold, additional planning lead time primarily introduces variability rather than stability. These results are consistent with [7].

B. Effect of Stochastic Uncertainties

The results show that stochastic uncertainty affects Extended AMAN performance in a measurable way. Stability KPIs are particularly sensitive, confirming that prediction uncertainty propagates directly into the planning process.

Pop-ups are most strongly associated with sequence instability. Configurations without pop-ups produce substantially fewer sequence position changes, indicating that late insertions caused by pop-ups are a main driver of sequence disruption.

Trajectory prediction uncertainty mainly appears in the form of increased EAT revisions, delay energy cost, and instruction count. Its effect on sequence position changes is smaller, which suggests that this uncertainty primarily causes slot timing variability rather than persistent reordering of the sequence.

For the planning effectiveness related KPIs, the effects are more limited. Although delay energy cost differs significantly, mean vectoring delay does not show a systematic effect. This indicates that uncertainty mainly reduces planning stability rather than strongly affecting the efficiency of delay absorption in this configuration. Overall, the findings show that different uncertainty sources affect different parts of system performance. Extended AMAN instability at Schiphol

Table XVIII. Qualitative summary of experimental results relative to the reference configuration in each comparison.

Reference	Configuration	EAT revisions	Sequence changes	Energy cost	Total delay	Vectoring delay	Vectored flights	Instruction count
A. Effect of freeze horizon extension								
FCFS14	FCFS20	- -	- -	+	+	+ +	+ +	+
FCFS14	FCFS25	- -	- -	+	+	+ +	+ +	-
B. Effect of stochastic uncertainty								
No uncertainty	no pop-ups	- -	- -	- -	-	0	-	- -
No uncertainty	no TP uncertainty	- -	- -	+	+	+	+	+
No uncertainty	no en-route uncertainty	- -	- -	+	+	+	+	+
No uncertainty	all uncertainty	- -	- -	- -	+	+	-	-
C. Evaluation of pop-up mitigation strategies								
FCFS20	BOL20	+ +	+ +	+ +	- -	-	-	-
FCFS20	Delay20	+ +	+ +	-	-	-	-	-
FCFS25	BOL25	+ +	+ +	+ +	- -	- -	- -	- -
FCFS25	Delay25	+ +	+ +	+	- -	- -	- -	- -
FCFS14	BOL20	- -	-	+	0	+ +	+ +	0
FCFS14	Delay20	- -	- -	-	-	+	+	-
D. Effects of planning at take-off and Brussels pop-up traffic								
FCFS20	FCFS20, planning at T/O	+ +	+ +	+	+	+ +	+ +	+
FCFS20	FCFS20 excl. EBBR	+ +	+ +	+	0	+	+	0
FCFS20	FCFS20 excl. EBBR, planning at T/O	+ +	+ +	+	0	+	+	+
BOL20	BOL20, planning at T/O	-	-	-	0	+	+	0
BOL20	BOL20 excl. EBBR	0	0	+	+ +	+ +	+ +	+ +
BOL20	BOL20 excl. EBBR, planning at T/O	0	-	+	+ +	+ +	+ +	+ +
FCFS20	BOL20	+ +	+ +	+ +	- -	-	-	-
FCFS20	FCFS20 excl. EBBR, planning at T/O	+ +	+ +	+	0	+ +	+ +	+
FCFS14	FCFS14, planning at T/O	+ +	+ +	+	0	0	0	0
FCFS14	FCFS20 excl. EBBR, planning at T/O	- -	- -	+	+	+ +	+ +	+

+ + / - -: significant improvement / deterioration relative to reference; + / -: non-significant difference; 0: minimal difference.

should therefore not be interpreted as a purely pop-up-driven phenomenon, but as the combined effect of pop-ups and stochastic trajectory prediction uncertainty.

The Z-score plot shows overlap for mean total delay and mean vectoring delay, which is consistent with the weaker statistical separation found for these KPIs. It is also notable that the no-uncertainty configuration still exhibits a relatively large spread for several KPIs.

In the no-uncertainty configuration, the traffic demand appears to be less evenly distributed over time, with relatively low demand at the very beginning of the scenario followed by a pronounced traffic bunch. As a result, early gaps in the sequence remain unfilled, while delay accumulates later in the scenario when demand increases. When stochastic uncertainty is introduced, the traffic demand becomes more temporally dispersed. This spreading effect allows aircraft to occupy earlier available slots more effectively, leading to a more gradual build-up of delay and reducing peak congestion levels.

However, this behavior is not consistently observed across all simulation runs, which contributes to the large spread in several KPIs. This indicates that the results are sensitive to the presence and timing of traffic peaks. As a consequence, no definitive conclusions can be drawn for some of the KPI's in the no uncertainty configuration. More generally, this highlights a limitation of the simulation setup, as system performance appears to be strongly influenced by scenario-specific demand patterns rather than solely by the modeled uncertainty mechanisms. Some effects between the other configurations are significant.

C. Evaluation of Pop-Up Mitigation Strategies

The scheduler comparison shows clear differences between the evaluated mitigation concepts. At a 20-minute freeze horizon, the back-of-the-line strategy gives the most robust stability performance. By preventing late insertions from disturbing the frozen sequence, BOL limits the propagation of instability. The reduction in EAT revisions indicates that preserving sequence structure is an effective mitigation mechanism under stochastic uncertainty.

The increase in mean total delay under BOL reflects redistribution rather than a deterioration of performance. Delay is shifted toward pop-up flights while stability is preserved for aircraft that were already planned. From an operational perspective, this trade-off may be acceptable, since stability directly affects workload and predictability. It must be noted that, at a 20-minute freeze horizon, only aircraft originating from Brussels (2 flights per scenario, about 1% of the total) experience more than approximately 3 minutes of additional delay due to the BOL scheduler.

The Delay scheduler does not show comparable benefits. Although it is conceptually attractive because it attempts to account for uncertainty before a flight becomes a pop-up, the results show that static delay offsets are not effective at providing an accurate slot. In both FH20 and FH25, the delayed-slot strategy remains less stable than BOL and tends to increase total delay absorption. This suggests that pre-planning based on a fixed statistical buffer is only effective when the underlying uncertainty is sufficiently predictable. Under the uncertainty levels modeled in this study, that condition is not met.

The Z-score plot for the 20-minute freeze horizon indicates that several KPIs still show substantial spread and overlap between configurations. This is consistent with the absence of statistical significance for those measures. At the same time, the figure clearly shows that both BOL and the delayed-slot strategy improve the stability-related KPIs relative to FCFS, with BOL providing the strongest and most consistent effect.

For the 25-minute freeze horizon, instability effects remain dominant across schedulers. While BOL continues to outperform the alternative strategies in terms of stability, the absolute magnitude of EAT revisions and sequence position changes suggests that this horizon length may exceed practical operational limits under the modeled uncertainty conditions. This is also reflected in the corresponding Z-score plot, where the separation from FCFS becomes more pronounced, while BOL and Delay still show considerable overlap for several KPIs.

Taken together, the results indicate that stability-preserving schedulers represent the most viable mitigation strategy within the evaluated uncertainty environment but not for freeze horizons beyond 20 minutes in the Schiphol case. BOL planning provides consistent benefits at FH20, while longer horizons remain fundamentally constrained by uncertainty propagation.

The delay type redistribution is operationally relevant. Compared with FCFS14, BOL20 shifts delay absorption away from vectoring and toward earlier speed- and Mach-based absorption, as illustrated by Fig. 11. Delay is therefore handled earlier in the trajectory and through more predictable control mechanisms, while vectoring-related intervention is reduced. Since the number of instructions does not increase significantly under BOL20, this shift appears beneficial from both an efficiency and workload perspective. The main disadvantage is that the burden of this strategy is concentrated on the pop-up flights themselves. In the evaluated scenarios, these are predominantly flights from Brussels, implying that the strategy effectively prioritizes sequence stability for the main inbound flow at the expense of a small subset of nearby departures.

This raises an important operational consideration. Implementing BOL at a 20-minute horizon would not merely be a technical modification of the scheduler, but also a policy choice regarding how uncertainty should be distributed across traffic flows. Flights for which reliable information becomes available relatively late would systematically be placed further back in the sequence. From the perspective of aggregate E-AMAN performance, this yields clear benefits. From the perspective of equity between traffic flows, however, it implies a structural disadvantage for some nearby departure streams. This trade-off should therefore be considered explicitly in any future implementation study and by LVNL.

D. Effects of Planning at Take-Off and Brussels Pop-Up Traffic

Several observations can be made based on the results presented above. First, planning aircraft at take-off (using EFD messages) improves the performance of E-AMAN under FCFS planning, especially when combined with the exclusion of Brussels pop-up flights. Under this configuration, clear and significant benefits are found for the stability-related KPIs, and additional improvements are visible for the vectoring-related KPIs, though only significant in the case of planning at take-off only. Examining a BOL scheduler and the same comparison, this effect does not occur. Instead, excluding Brussels pop-up flights mainly improves the delay- and taskload-related KPIs, while stability does not improve significantly. The improvement in delay-related performance can be logically explained by the fact that these Brussels pop-up flights are the flights that would otherwise tend to be planned further back in the sequence.

Second, when compared with other mitigation strategies, the combination of planning at take-off and excluding Brussels pop-up flights appears to perform nearly as well as BOL on some metrics, and even better on others. An important qualification, however, is the practical feasibility of these options. Planning at take-off using EFD messages is technically feasible, since these messages are already available and could therefore realistically be integrated into operations. By contrast, BOL would represent a substantial operational change for LVNL, while excluding Brussels pop-up flights would imply a major change for KLM and other stakeholders, likely involving considerable financial interests. A related alternative would be to improve the predictability of Brussels departures sufficiently to enable reliable pre-take-off planning, for example by direct communication with the local control tower. Moreover, excluding Brussels pop-up flights is not a true mitigation strategy for the pop-up problem itself, but rather removes an important source of pop-ups from the system.

Finally, the last comparison indicates that implementing EFD-based planning—that is, planning aircraft at take-off rather than only after radar correlation—can already improve the current AMAN operation in terms of stability. In addition, the E-AMAN configuration combining planning at take-off and excluding Brussels pop-up flights still provides significant improvements relative to the current AMAN in terms of mean vectoring delay and the number of vectored flights, together with a non-significant tendency towards fewer instructions, lower total delay, and lower delay energy cost. These benefits do, however, come at the expense of reduced stability in that comparison. Overall, the results show that the value of these measures depends strongly on the planning strategy in which they are embedded and on the performance dimension that is considered most important.

E. General Discussion

The results confirm that freeze horizon extension is difficult under operational uncertainty. While a longer horizon creates more room for earlier sequencing and upstream delay absorption, it also reduces planning robustness. In the evaluated Schiphol environment, both pop-up flights and trajectory prediction uncertainty contribute to this degradation. Pop-ups mainly affect sequence stability through late insertions, whereas trajectory prediction uncertainty mainly affects temporal stability through repeated EAT revisions.

Among the evaluated scheduler strategies, back-of-the-line sequencing at a 20-minute freeze horizon appears to be the most relevant mitigation option. It shows the strongest improvement in stability while keeping total delay broadly comparable and shifting delay toward earlier and more speed-based instructions. At the same time, the additional experiments show that earlier planning at take-off is also operationally relevant, as it directly reduces the late-visibility mechanism and thereby limits the impact of pop-ups. The Brussels comparison further indicates that a small subset of nearby departures contributes disproportionately to instability. Together, these findings suggest that both stability-preserving scheduling and improved early flight visibility are relevant for making Extended AMAN operationally feasible.

An exploratory analysis was also performed for flights subject to a CTOT or issued slot as a method of uncertainty reduction [20]. In principle, one might expect such flights to exhibit lower uncertainty from the perspective of the AMAN operator, since the departure process is constrained to a narrower operational window and should therefore be more predictable for LVNL. However, a preliminary EFD-based assessment suggests that this uncertainty is not smaller, and may even be larger. At present, the reason for this is not clear. This matters because the delayed-slot concept would become more attractive if slot-regulated departures could be predicted more accurately from the AMAN perspective. Future work should therefore evaluate whether the apparent increase in uncertainty is caused by the underlying operational process, by the way the information is represented in EFD messages, or by selection effects in the sampled traffic.

The en-route uncertainty model used in this study requires careful interpretation. It is based on uncertainty observed in trajectory predictions derived from EFD messages, which are not necessarily representative of the best available prediction accuracy. In addition, the model assumes a linear increase in uncertainty with look-ahead time. This causes uncertainty growth with increasing look-ahead time, making the model particularly sensitive to freeze horizon extension. However, Lubig et al. suggests that this growth may be more moderate under operational conditions [12]. The uncertainty representation used in this study may therefore be conservative, particularly for the en-route component.

This study focuses on high-density traffic under a fixed set of operational conditions, including specific wind settings and runway configurations. The results therefore reflect system behavior within a constrained scenario. Future work should assess the sensitivity of the results to variations in traffic density, wind, and runway use. Evaluating AMAN and E-AMAN performance across a broader range of traffic scenarios would provide insight into the robustness of the observed effects.

More broadly, the results indicate that structural performance improvements could come from uncertainty reduction rather than compensation alone. Scheduler logic can mitigate the operational consequences of uncertainty, but it cannot remove the underlying cause. Promising directions include improved access to intent information through ADS-C [21], wider information exchange through SWIM, and tighter coupling between departure-side planning and arrival-side planning. In particular, integration between DMAN at the departure airport and AMAN at the arrival airport could reduce the amount of uncertainty that propagates into the arrival sequence by making the expected departure sequence known [11]. Further research is needed into how much the uncertainty perceived by the AMAN operator must actually be reduced for specific concepts to become operationally viable. The simulation framework developed in this study is well suited for such threshold and sensitivity analyses.

A further development area is limited tactical intervention before slot freeze. In the current implementation, preplanned flights are not given early delay-absorption instructions because such interventions would alter the ETA used for slot assignment. However, some flights may already clearly require delay absorption before the freeze horizon is reached. A future concept in which planning is based on the original ETA while tactical absorption is initiated earlier may improve both predictability and workload distribution. The simulation framework developed in this study provides a suitable basis for testing such concepts.

Several limitations should be acknowledged. The experiments focus on peak-demand scenarios and do not include pronounced low-demand periods or strongly disrupted weather scenarios. The stochastic uncertainty model assumes independence between uncertainty sources, while in practice some dependencies are likely to exist. Traffic is not tactically deconflicted beyond the modeled AMAN and ATC logic, so the feasibility of certain control solutions may be overestimated. Wake-turbulence effects are not modeled explicitly in the sequencing logic, with a fixed landing interval used instead. Controller behavior is represented by deterministic rule-based logic and does not capture variability in decision-making, inter-center coordination, or pilot compliance. In addition, the study remains a model-based analysis only; before operational implementation, human-in-the-loop simulation or shadow-mode evaluation would be required to assess controller acceptability and practical feasibility. Finally, the results are specific to Schiphol traffic structure and the selected scenario set, even though the methodological framework itself is transferable. Despite these limitations, the results provide a clear indication that practical Extended AMAN implementation at Schiphol will require either reduced uncertainty or stability-oriented scheduling logic.

VII. Conclusion

This study evaluated the influence of stochastic uncertainty mechanisms and pop-up mitigation strategies in an Extended

Arrival Manager (E-AMAN) environment. The results show that freeze horizon extension introduces a clear trade-off between earlier planning and prediction-driven instability. While a longer freeze horizon increases the opportunity for upstream delay absorption, it also increases exposure to different types of uncertainty.

The results further show that pop-ups are a major driver of instability in the evaluated Schiphol environment, especially with respect to sequence disruptions. At the same time, trajectory prediction uncertainty remains an important contributor to E-AMAN performance degradation through repeated changes in Estimated Time Over the Initial Approach Fix (ETO IAF). Extended AMAN performance is therefore driven by the combined effect of late flight visibility and stochastic prediction uncertainty.

Among the evaluated mitigation strategies, back-of-the-line planning at a 20-minute freeze horizon provides the most favorable balance between stability, planning effectiveness and taskload. The strategy effectively constrains instability propagation by protecting the frozen sequence from late insertions, while preserving overall delay performance relative to baseline AMAN operations. In addition, delay absorption shifts toward earlier and more predictable mechanisms, reducing reliance on vectoring without significantly increasing instruction count. Longer freeze horizons produce diminishing returns, as uncertainty propagation effects become increasingly dominant.

Further evaluation shows that earlier planning of nearby departures at take-off, using ETFMS Flight Data (EFD), can also provide meaningful benefits. In particular, planning flights already at take-off reduces the number of pop-ups and significantly improves stability, both in Extended AMAN and in the current AMAN configuration.

The experiments also show that pop-ups from Brussels have a disproportionate effect on E-AMAN instability. This indicates that another route toward feasible Extended AMAN implementation is not only to change the scheduler logic, but also to reduce the amount of late-visible nearby traffic through earlier planning and improved information availability.

Taken together, the results indicate that a promising pathway for Extended AMAN implementation consists of a combination of stability-preserving planning logic and uncertainty-reduction measures. Of the evaluated options, a Back-of-the-Line (BOL) scheduler at a freeze horizon of 20 minutes appears to be the strongest local mitigation strategy. EFD-based planning at take-off appears to be a promising complementary measure that could already improve current operations at a freeze horizon of 14 minutes.

This research establishes a structured framework for stochastic uncertainty modeling and reproducible evaluation of AMAN and Extended AMAN concepts through the BlueSky AMAN Simulator (BAMS). The framework enables systematic assessment of uncertainty sources, scheduling strategies, and delay absorption mechanisms under controlled conditions. This provides LVNL and other ANSPs with a useful basis for evaluating Extended AMAN concepts and potential Common Project 1 (CP1) compliance pathways.

REFERENCES

[1] EUROCONTROL, "Arrival manager implementation guidelines and lessons learned," EUROCONTROL, Tech. Rep. Edition 0.1, Dec. 2010. [Online]. Available: https://skybrary.aero/sites/default/files/bookshelf/2416.pdf

[2] M. Van Horssen, "Cross-border arrival management to reduce traffic bunching at schiphol airport," MSc thesis, Delft University of Technology, Jul. 2017, tU Delft Master Thesis in Aerospace Engineering. [Online]. Available: https://repository.tudelft.nl/islandora/object/uuid:bf30d6f8-4ac3-4cbe-9869-4fe923118e6f

[3] A. Vanwelsenaere, J. Ellerbroek, F. Dijkstra, and E. Westerveld, "The impact of pop-up flights on extended arrival management," in *Proceedings of the 8th International Conference on Research in Air Transportation (ICRAT)*, 2018, extended AMAN / pop-up flight stability analysis.

[4] A. Hollebeek, "Arrival management evaluation: An insight into advanced schiphol arrival planner," BSc thesis, KDC Mainport Schiphol – Centre of Excellence (AUAS & TU Delft), 2019, graduated spring 2019; collaboration between Aviation Academy, Amsterdam University of Applied Sciences and TU Delft.

[5] J. Thipphavong and S. Landry, "Effects of the uncertainty of departures on multi-center traffic management advisor scheduling," in *AIAA 5th ATIO and 16th Lighter-Than-Air Systems and Balloon Systems Conferences*. American Institute of Aeronautics and Astronautics, Sep. 2005. [Online]. Available: https://doi.org/10.2514/6.2005-7301

[6] M. Tielrooij, "Arrival management support in the presence of prediction uncertainty," Ph.D. dissertation, Delft University of Technology, Oct. 2022. [Online]. Available: https://repository.tudelft.nl/record/uuid:a403112b-48a2-40d7-9db3-a5b754f31eee

[7] A. Vanwelsenaere, "Effect of pop-up flights on extended arrival manager," MSc thesis, Delft University of Technology, Jul. 2016. [Online]. Available: https://repository.tudelft.nl/islandora/object/uuid:fe4cb9c1-59a9-45a5-853d-1da2dccc74a1

[8] European Commission, "Cp1 regulation: Commission implementing regulation (eu) 2021/116 on the establishment of the common project one," Official Journal of the European Union, Commission Implementing Regulation, Feb. 2021, entered into force 22 Feb 2021; establishes Common Project One (CP1) supporting the ATM Master Plan. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32021R0116

[9] Luchtverkeersleiding Nederland (LVNL), "Advanced schiphol arrival planner (asap) – ontwerpdocument: Procedures," LVNL – Luchtverkeersleiding Nederland, Technical Report, Mar. 2019, version 2.0; internal design document for AMAN (Arrival Manager) implementation.

[10] K. Sekine, F. Kato, T. Tatsukawa, K. Fujii, and E. Itoh, "Rule design for interpretable en route arrival management via runway-flow and inter-aircraft control," *IEEE Access*, vol. 11, pp. 75 093–75 111, 2023.

[11] S. Rydell, "Arrival and departure manager cooperation for reducing airborne holding times at destination airports," MSc thesis, Cranfield University, Aug. 2011. [Online]. Available: http://dspace.lib.cranfield.ac.uk/handle/1826/8022

[12] D. Lubig, J. Rosenow, N. Peter, and H. Fricke, "From en-route to touchdown: Uncertainty analysis of inbound traffic flows to singapore changi airport," in *Air Traffic Management Research Conference*, 2025.

[13] J. M. Hoekstra and J. Ellerbroek, "Bluesky atc simulator project: An open data and open source approach," in *Proceedings of the 7th International Conference on Research in Air Transportation (ICRAT)*, 2016, pp. 131–138. [Online]. Available: https://resolver.tudelft.nl/uuid:d1131a90-f0ea-4489-a217-ad29987689a1

[14] H. Koolen and I. Coliban, "Flight progress messages document," EUROCONTROL Network Manager, Tech. Rep., Jul. 2020. [Online]. Available: https://www.eurocontrol.int/sites/default/files/2020-06/flight-progress-msg-update-230620.pdf

[15] EUROCONTROL, "Eurocontrol specification for airport collaborative decision making (a-cdm)," EUROCONTROL, Specification, Jul. 2024, draft Edition 1.0 (05 Dec 2024); supports harmonised A-CDM implementation at European airports. [Online]. Available: https://www.eurocontrol.int/sites/default/files/2024-07/eurocontrol-draft-specification-acdm-ed1-0.pdf

[16] ——, "Annex to SESAR DOD G – Operational Scenarios," EUROCONTROL, Brussels, Belgium, Technical report, 2007, e3-WP2-D2.2-050-OS, Version 1.00. [Online]. Available: https://www.eurocontrol.int/sites/default/files/2022-10/E3-WP2-D2.2-050-OS-V1.00-os.pdf

[17] “Qrc8 baancombinaties en liv – capaciteit in asap (apln),” LVNL, Technical Report, 2024, qRC8-rapport over baancombinaties en LIV-capaciteit in het ASAP-systeem (APLN).
[18] EUROCONTROL, “Ddr2 user guide,” EUROCONTROL Network Operations Library, User Guide, 2022, generic user guide for the DDR2 Demand Data Repository service.
[19] A. Field, *Discovering Statistics Using IBM SPSS Statistics*, 3rd ed. Sage, 2009.
[20] SKYbrary, “Calculated take-off time (ctot),” SKYbrary Aviation Safety, 2023. [Online]. Available: https://skybrary.aero/articles/calculated-take-time-ctot
[21] N. M. Guerreiro and M. C. Underwood, “Understanding extended projected profile (epp) trajectory error using a medium-fidelity aircraft simulation,” in *2018 Aviation Technology, Integration, and Operations Conference*. American Institute of Aeronautics and Astronautics. [Online]. Available: https://arc.aiaa.org/doi/10.2514/6.2018-3044

This page is intentionally left blank.

APPENDIX A
CALCULATED TAKE-OFF TIME (CTOT) AND TAKE-OFF PREDICTION UNCERTAINTY

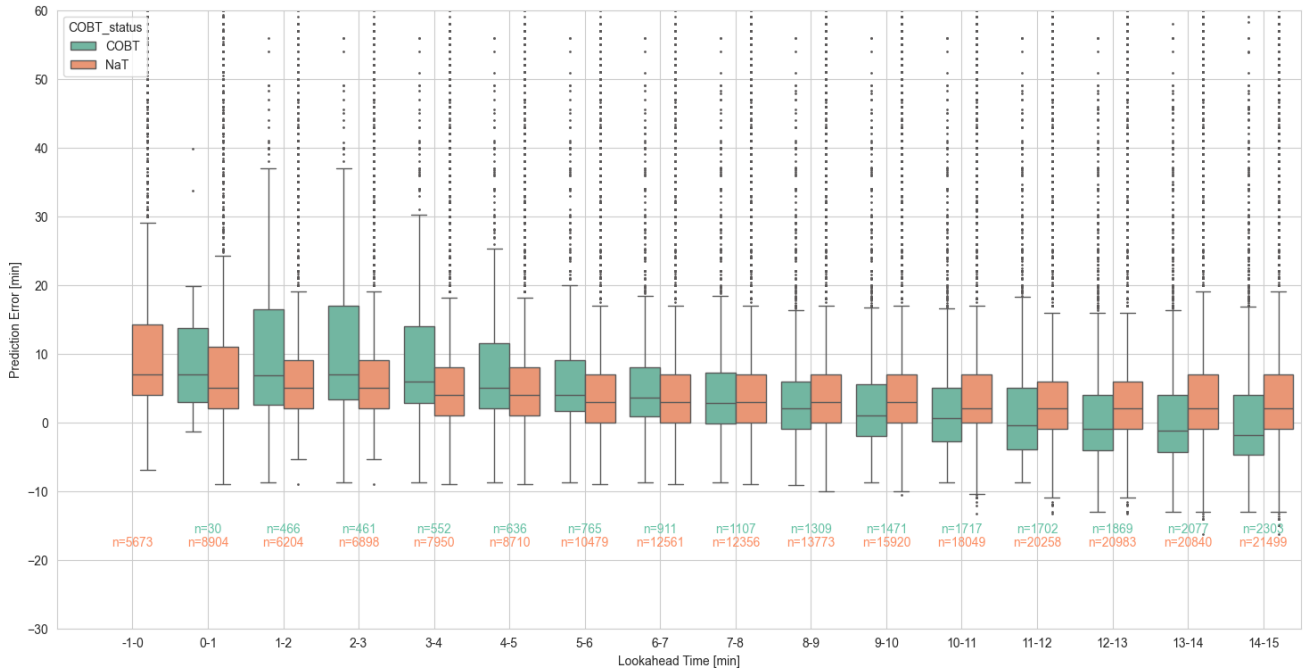

Figure 16. Comparison of take-off prediction uncertainty for flights with and without a CTOT or issued slot. Flights with a slot are shown in green, while flights without a slot are shown in orange.

Fig. 16 compares the take-off prediction uncertainty of flights with and without a CTOT or issued slot. Flights with a slot are shown in green, while flights without a slot are shown in orange. The figure does not show a clear reduction in take-off prediction uncertainty for slot-regulated flights. A similar pattern was found when the analysis was repeated for individual departure airports.

These results suggest that the presence of a CTOT or issued slot does not automatically improve take-off predictability from the perspective of the AMAN operator. Further analysis is required to determine whether this is caused by the selected dataset, by underlying operational effects, or by the fact that slot allocation alone is insufficient to reduce take-off uncertainty.

APPENDIX B
DELAY DISTRIBUTION FIRST-COME-FIRST-SERVE VERSUS BACK OF THE LINE FOR BRUSSELS

Fig. 17 presents a comparison between the FCFS and BOL configurations in terms of the mean total delay per scenario. For both methods, a distinction is made between three categories: all flights, flights originating from Brussels (EBBR), and flights not originating from Brussels.

The boxplots illustrate the distribution of the mean delay across the different scenarios. When examining the subgroups, it becomes clear that under the BOL configuration a large delay is assigned to flights originating from Brussels, while the delay for non-Brussels flights is comparatively reduced. This indicates that the BOL method redistributes delay and concentrates a larger portion of it on Brussels-related traffic, as theorized.

APPENDIX C
SCENARIOS

The 12 scenarios were constructed to represent comparable 3-hour peak periods under different operational conditions. They vary in traffic day, runway configuration, wind profile, and time shift, while keeping the underlying demand level within a narrow range. Together, they provide a set of semi-steady-state peak scenarios for evaluating the relative effect of uncertainty and scheduler logic under paired Monte-Carlo conditions. The applied shift hours indicate the temporal offset used when generating the scenario from the original traffic day, allowing multiple peak-period realizations to be derived from the same source dataset.

In case of the exclusion of the Brussels pop-up flights, these were replaced with flights with a similar ETA. The rest of the scenario was kept identical.

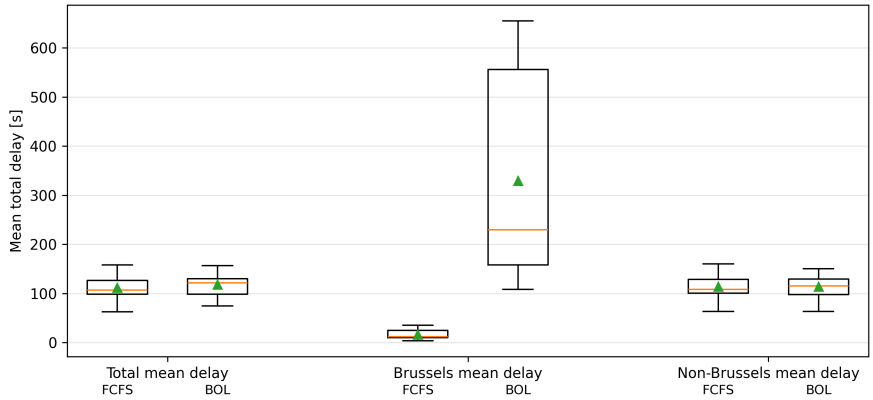

Figure 17. Comparison of mean total delay per scenario for FCFS and BOL, split into all flights, Brussels-origin flights, and non-Brussels flights.

Table XIX. Overview of the 12 generated traffic scenarios used in the simulation experiments. The table lists the traffic day, runway configuration, wind setting, runway rate, and applied time shift for each generated scenario. Each scenario covers a 3-hour traffic window.

Scenario	Traffic day	Runway configuration (RIVER, SUGOL, ARTIP)	Runway rate	Shift hours	Wind profile (alt [ft], direction, magnitude [kts])
sc1	2024-05-22	[18R, 18R, 18C]	34	0	0/60/15, 4000/60/30, 10000/60/50
sc2	2024-05-22	[36C, 36C, 36R]	34	3	0/180/15, 4000/110/30, 10000/60/50
sc3	2024-05-22	[18R, 18R, 27]	34	6	0/60/15, 4000/90/30, 10000/90/50
sc4	2024-05-22	[18R, 18R, 18C]	34	9	0/310/15, 4000/290/30, 10000/290/50
sc5	2024-02-09	[18R, 18R, 18C]	34	0	0/180/15, 4000/225/30, 10000/180/30
sc6	2024-02-09	[36C, 36C, 36R]	34	3	0/180/15, 4000/110/30, 10000/60/50
sc7	2024-02-09	[36C, 36C, 36R]	32	3	0/180/15, 4000/110/30, 10000/60/50
sc8	2024-02-09	[18R, 18R, 27]	34	6	0/60/15, 4000/90/30, 10000/90/50
sc9	2024-01-26	[18R, 18R, 18C]	34	0	0/60/15, 4000/60/30, 10000/60/50
sc10	2024-01-26	[36C, 36C, 36R]	34	3	0/180/15, 4000/110/30, 10000/60/50
sc11	2024-01-26	[18R, 18R, 27]	34	6	0/60/15, 4000/90/30, 10000/90/50
sc12	2024-01-26	[18R, 18R, 18C]	34	9	0/310/15, 4000/290/30, 10000/290/50

APPENDIX D
ERROR SAMPLING AND INTRODUCTION

This appendix section describes how the stochastic uncertainty terms are sampled and how they are introduced into the simulation. The uncertainty model distinguishes between four error components: take-off planning uncertainty, departure-route uncertainty, en-route trajectory prediction uncertainty outside the FIR, and trajectory prediction uncertainty within the FIR.

A. Sampling of Error Terms

For each aircraft, the uncertainty terms are sampled once and then kept fixed throughout the simulation. This ensures that each flight is assigned a consistent stochastic realization over all prediction updates. A deterministic pseudo-random number generator is used based on a combination of the global simulation seed and the aircraft callsign. As a result, the same aircraft receives the same error realization when the same scenario seed is used, which is essential for the paired Monte-Carlo comparison between configurations.

The take-off planning uncertainty is sampled from airport- and lookahead-dependent Johnson-SU distributions derived from the EFD analysis. For a given aircraft, the relevant lookahead is determined from the difference between the freeze horizon and the predicted flight time to the IAF. If this lookahead becomes negative, the error is zero. In addition to the random sample itself, the corresponding percentile value used by the delayed-slot scheduler is also extracted from the same take-off distribution.

Departure-route uncertainty is sampled from an airport-specific Johnson-SU distribution. This uncertainty is represented as a relative timing error over the departure segment rather than as an absolute time shift. En-route trajectory prediction uncertainty outside the FIR is sampled from a single Johnson-SU distribution representing the normalized error up to the coordination point. Within the FIR, trajectory prediction uncertainty is sampled from a normal distribution with a standard deviation of 1.5%, reflecting the smaller spread associated with known local intent.

After sampling, each component can be scaled independently using the error multipliers defined in the AMAN settings. This allows individual uncertainty sources to be enabled, disabled, or amplified in the experiment design without changing the underlying distributions.

B. Introduction of Take-off Uncertainty

Take-off planning uncertainty is introduced as a shift in the actual take-off time relative to the predicted take-off time. For aircraft that are planned before take-off, the sampled take-off error is converted to a time shift and applied to the aircraft creation time in the simulation. As a result, the aircraft appears later or earlier than originally predicted, which directly affects whether it becomes visible to the AMAN in time to be preplanned or instead appears as a pop-up.

This mechanism is only relevant for aircraft that are still on the ground when the uncertainty is applied. It therefore primarily affects nearby departures and experiments involving planning at take-off or delayed-slot scheduling.

C. Introduction of Airborne Uncertainty

For the airborne phases, uncertainty is not introduced as a direct time offset, but through a modification of the effective groundspeed. Each sampled error term is interpreted as a relative timing error over a specific trajectory segment. During simulation, this is converted into a groundspeed factor,

$$V_{\mathrm{ground,eff}} = V_{\mathrm{ground,nom}} \cdot (1 - E_{\mathrm{segment}}), \tag{11}$$

where E_{segment} is the sampled normalized error for the active phase.

The predictions are updated as the simulation runs. The resulting timing error is therefore not constant, but gradually decreases from a maximum as the aircraft approaches the end of the affected segment.

D. Phase Logic During Flight

The active uncertainty component depends on the phase of flight. Before the end of the departure segment, the departure-route error is applied. After that, the en-route error outside the FIR becomes active. Once the aircraft enters the FIR, the model transitions to the within-FIR uncertainty representation. In addition, if an adjacent-center instruction has already been issued, the simulation also switches from the outside-FIR en-route error to the within-FIR error representation, because the aircraft intent is then assumed to be known more accurately.

The uncertainty is only activated once the aircraft reaches the relevant planning window. In the present implementation, this window starts 5 minutes before the freeze horizon relative to the predicted ETO at the IAF. This avoids introducing drift too early, when the aircraft is still operationally irrelevant for AMAN planning.

E. Effect on AMAN Predictions

The cumulative time drift caused by the active uncertainty terms is stored as the aircraft's instantaneous prediction error. This drift is added to the nominal trajectory predictor output to obtain the effective ETO at the IAF and the corresponding ETA observed by the AMAN. The AMAN therefore does not receive a direct stochastic perturbation, but rather a gradually evolving prediction error that follows from the realized uncertainty in take-off timing and effective airborne progress.

This approach makes it possible to distinguish between two different mechanisms of instability. First, take-off uncertainty and visibility logic determine whether aircraft become late-visible pop-ups. Second, departure-route and en-route uncertainty continuously perturb the arrival-time predictions of aircraft that are already known to the AMAN. The simulation can therefore separately represent instability caused by late visibility and instability caused by prediction drift.

APPENDIX E
SCHEDULERS

A. First-Come, First-Served (FCFS)

The FCFS scheduler represents the reference planning logic and is closest to current operational practice. Aircraft are inserted into the sequence primarily based on their predicted arrival time at the IAF. When a newly visible aircraft enters the planning process, it is assigned the first feasible slot in sequence while preserving the required separation to the aircraft ahead. If this aircraft is inserted before already planned or frozen aircraft, the flights behind it may need to be shifted, which can lead to EAT revisions and sequence position changes. This principle is illustrated on the left-hand side of Fig. 18.

B. Back-of-the-Line (BOL)

The Back-of-the-Line (BOL) scheduler is a stability-preserving mitigation strategy for pop-up flights. Instead of inserting a newly visible aircraft at its FCFS-based position in the existing sequence, the aircraft is appended to the back of the currently frozen sequence on its assigned runway. In other words, the pop-up flight is not allowed to disrupt the order of aircraft that have already been frozen. This logic is illustrated on the right-hand side of Fig. 18.

C. *Delayed-slot scheduler (Delay)*

The delayed-slot scheduler assigns an expected delay already during pre-planning. For each flight, this delay is based on the 75th percentile of the take-off uncertainty distribution corresponding to the departure airport and the relevant prediction lookahead. The delayed ETA is then obtained by adding this percentile delay to the nominal ETA. Slot assignment is subsequently based on this delayed ETA rather than on the nominal ETA. This mechanism is illustrated in Fig. 19.

In the present implementation, this scheduler assumes that aircraft are visible from take-off onward rather than only after reaching a specific altitude threshold. This assumption is necessary because the scheduler aims to account for uncertainty before the aircraft becomes a pop-up. The concept is therefore based on the idea that additional uncertainty can be absorbed through a statistical buffer in the planning phase.

However, if an aircraft is already airborne and still has not been frozen by the time it enters the freeze horizon, the delayed-slot logic is no longer applied. In that case, the aircraft is handled using Back-of-the-Line logic and is appended to the end of the frozen sequence. The Delay scheduler can therefore be interpreted as a hybrid concept: it first attempts to mitigate late disturbances through conservative pre-planning, but falls back to BOL behavior once a flight has effectively become a late-visible pop-up.

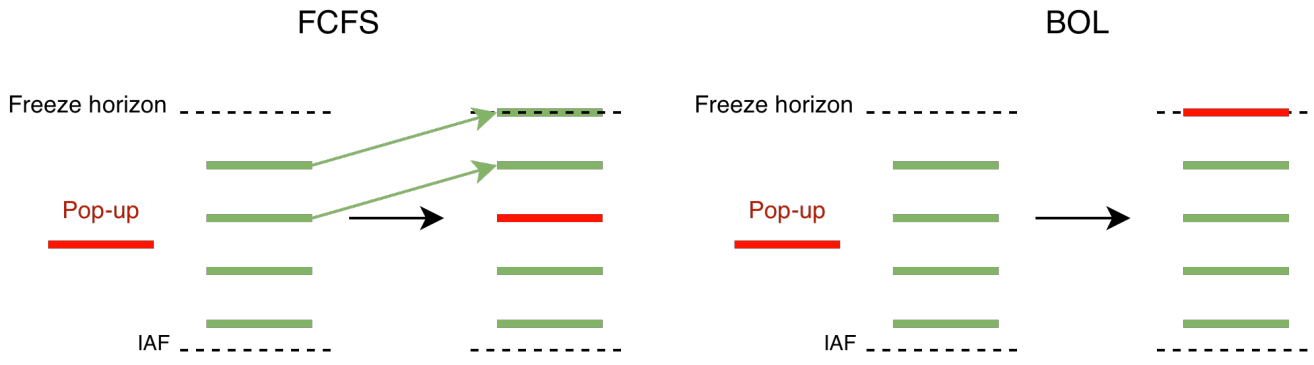

Figure 18. Schematic overview of the evaluated FCFS and BOL scheduler concepts.

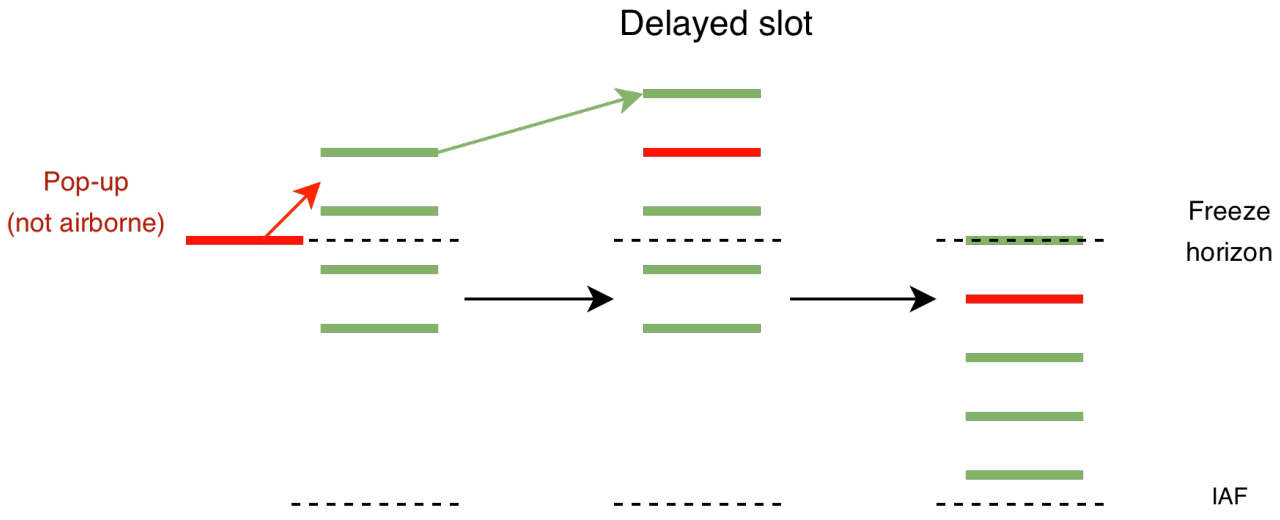

Figure 19. Illustration of the delayed-slot scheduling concept.

APPENDIX F
SETTINGS

This appendix section summarizes the most important AMAN settings used in the simulation environment. Only the settings that directly affect the planning logic, pop-up behavior, and tactical solution space are included here. Together, these parameters define how aircraft become visible to the AMAN, how slots are assigned and frozen, and how tactical delay absorption is performed.

A *standard early planning bias* of 60 seconds is applied, meaning that aircraft are planned slightly earlier than their nominal ETA when no preceding slot constraint is active. The current approach window is defined as ± 120 seconds around the EAT.

The available tactical solution space is defined through limits on speed control, vectoring, and holding. Aircraft may absorb delay using up to 50 knots of speed reduction, while time recovery is limited to 25 knots of speed increase. For adjacent-center coordination, a standard Mach reduction of 0.04 is used when in-sector speed reduction is insufficient. The maximum dogleg extension is constrained through a dogleg ratio, and the minimum usable speed outside the TMA is set to 190 knots. In addition, adjacent-center coordination is only considered above FL260, which acts as the handover altitude between local and non-local control.

Table XX summarizes the most relevant AMAN settings used in the simulation experiments.

Table XX. Main AMAN settings used in the simulation experiments.

Setting	Value	Description
Freeze horizon	14, 20, or 25 min	Time before IAF at which slots become frozen
Planning horizon	Freeze horizon + 15 min	Earliest point at which aircraft can be preplanned
Capacity	38 ac/h/runway	Determines nominal slot spacing
Nominal separation	≈ 95 s	Calculated from runway capacity
Standard early bias	60 s	Early slot bias in unconstrained planning
Approach window	± 120 s	Allowed deviation around EAT at IAF
Instruction margin	20 s	Minimum deviation before ATC intervenes
Replan pull-forward	90 s	Maximum earlier pull during replanning
Standard visible altitude	13,000 ft	Default radar visibility threshold
Airport-specific visibility	e.g. EBBR 1,000 ft, EDDL 2,000 ft	Lower thresholds for nearby airports
Max slowdown	50 kt	Maximum tactical speed reduction
Max speedup	25 kt	Maximum tactical speed increase
Mach reduction	0.04	Standard Mach reduction used for adjacent-center delay absorption
Minimum speed	190 kt	Minimum usable speed outside TMA
Dogleg ratio	1.3	Maximum dogleg extension relative to nominal route
Handover altitude	FL260	Boundary between local and adjacent-center control
Expected delay percentile	75%	Buffer used by delayed-slot scheduling
EAT update threshold	10 s	Minimum slot change counted as an EAT revision
Error multiplier	$(1, 1, 1, 1)$	Scaling factors for the four uncertainty phases
Within-FIR TP std. dev.	1.5%	Standard deviation of local trajectory prediction error

APPENDIX G
KPI DEFINITIONS

This appendix section summarizes the definitions of the main KPIs used.

A. EAT Revisions

revisions count the number of relevant slot updates after an aircraft is no longer in the preplanned or ground state. A revision is only counted when the assigned slot changes by more than 10 seconds. At scenario level, the KPI is the total number of such updates over all aircraft.

B. Sequence Position Changes

Sequence position changes are represented by the `swaps` variable. This KPI counts how often aircraft already in the frozen sequence are displaced by pop-up insertions or by subsequent sequence compression after replanning. At scenario level, the KPI is the total sum of swaps over all aircraft.

C. Delay Energy Cost

Delay energy cost represents the additional trajectory-related work caused by tactical delay absorption relative to the nominal trajectory. It is computed at scenario level as the ratio between total extra work and total work using:

$$\text{Delay energy cost} = \frac{\sum \text{extrawork}}{\sum \text{totalwork}} \times 100\%.$$

D. Instruction Count

Instruction count is the total number of tactical ATC instructions issued in a scenario. Speed-based instructions count as one instruction, a dogleg as two, and a holding pattern as four. The KPI is used as a proxy for controller taskload.

E. Nonzero Vectoring Delay

Nonzero vectoring delay counts the number of aircraft for which delay is absorbed through lateral path extension. It therefore measures how often vectoring-based delay absorption is used, rather than how much delay it absorbs.

APPENDIX H
OVERVIEW OF THE BLUESKY AMAN SIMULATOR (BAMS)

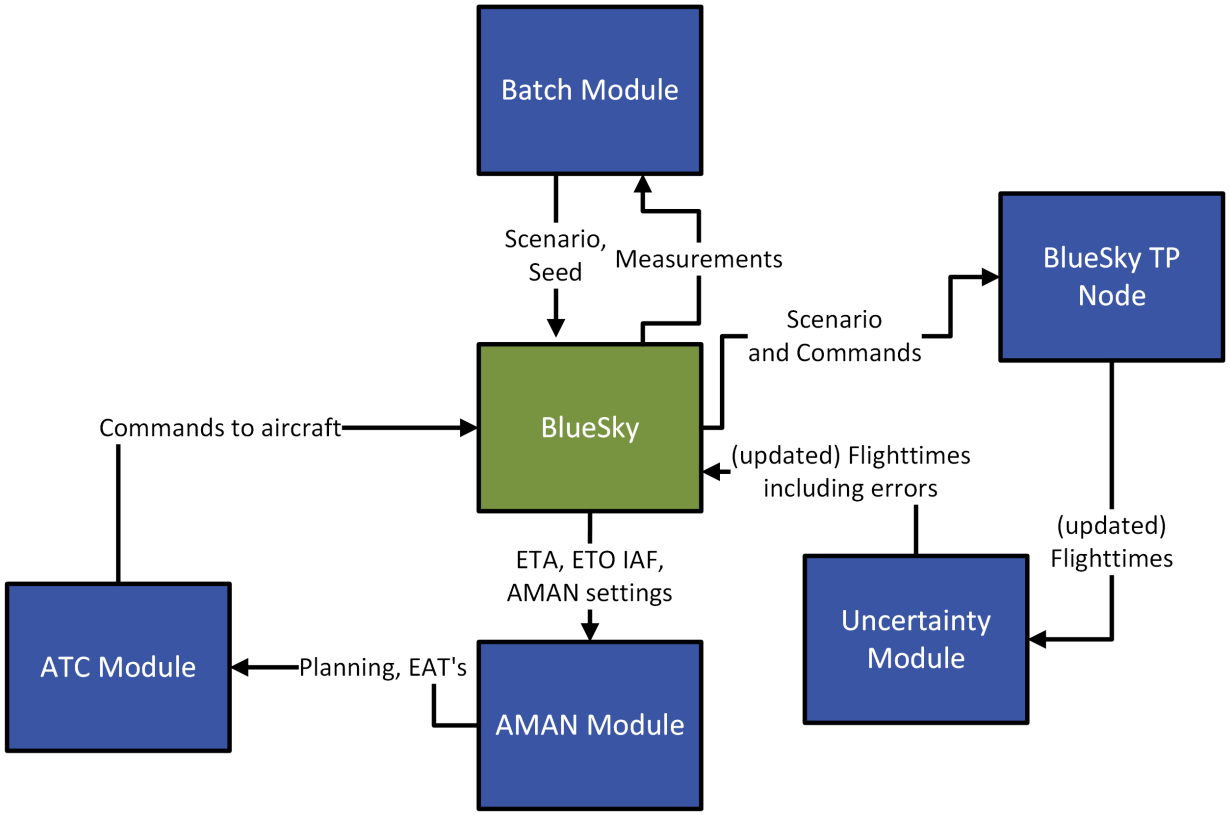

Figure 20. Overview of the main modules in the BlueSky AMAN Simulator (BAMS) and their interactions.

The BlueSky AMAN Simulator (BAMS) consists of several interacting modules that together represent the core functions of an arrival-management environment under uncertainty. The architecture is centered around the AMAN logic itself, but is extended with dedicated components for trajectory prediction, uncertainty generation, prediction handling, and tactical ATC intervention. Fig. 20 shows the main simulator modules and their information flow.

The core of the simulator is formed by the `AMAN` module. This module maintains the main flight DataFrame and is responsible for the AMAN planning logic. It determines aircraft planning states, performs preplanning, identifies pop-up flights, assigns slots, freezes aircraft at the configured freeze horizon, and stores the resulting planning variables such as slot, EAT, sequence relations, and bookkeeping variables. It also contains the logic for replanning and slot compression when late-visible aircraft disturb the frozen sequence.

Trajectory information is provided by the trajectory predictor module (`trajectory_predictor`). This module generates predicted waypoint crossing times, ETOs, and ETAs based on the BlueSky route and aircraft state. It can operate through a separate child node and communicates predictions back to the main simulation using BlueSky's network interface. In this way, the AMAN continuously receives updated time predictions for the IAF, runway, FIR entry, and other relevant reference points if an instruction is issued.

Tactical control actions are modeled in the `ATC` module. This module monitors frozen aircraft and determines whether instructions are needed based on the current time-to-lose-or-gain relative to the EAT window. It selects between speed control, dogleg vectoring, holding, direct-to clearances, and replanning, depending on aircraft state and available solution space. The ATC module also processes trajectory predictor updates after an instruction has been issued and checks whether the intended correction has been achieved or whether a new action is required.

Uncertainty is introduced through the dedicated error-generation functionality that is called from within the AMAN and prediction-handling workflow. Four uncertainty phases are represented: take-off planning uncertainty, departure-route uncertainty, en-route uncertainty outside the FIR, and en-route uncertainty within the FIR. These sampled errors are stored per flight and

then translated into take-off shifts or groundspeed factors, such that the AMAN observes gradually evolving time errors rather than a single static perturbation.

Together, these modules form a closed-loop simulation architecture. The trajectory predictor supplies time estimates, the AMAN uses these estimates to plan and freeze traffic, the ATC module intervenes when needed to maintain adherence to the plan, and new predictor updates then modify the estimated arrival times again. This closed interaction makes it possible to evaluate not only static planning performance, but also the dynamic propagation of uncertainty and the operational effect of different pop-up mitigation strategies.

APPENDIX I
MONTE-CARLO CONTROL

A dedicated `montecarlo` or batch plugin is used to control the simulation experiments. This plugin acts as the central experiment manager of BAMS. The user can start a complete batch of runs with a single scenario command, after which the plugin automatically loads the requested scenario and configuration, assigns seeds, distributes runs over multiple BlueSky nodes, and collects the results.

The plugin stores the run definitions and results in a central batch DataFrame. Each row corresponds to one scenario–seed combination and contains the execution status, assigned node, runtime information, and returned KPIs. By controlling the seeds centrally, the plugin ensures paired stochastic conditions between configurations. It therefore forms the backbone of the Monte-Carlo methodology used in this study.

APPENDIX J
MACH CROSSOVER AND SPEED PROFILE LOGIC

The `Mach_Crossover` plugin manages the nominal speed profile of each aircraft during climb, cruise, and descent. It ensures a realistic transition between calibrated airspeed (CAS) and Mach number as altitude changes, while still allowing tactical ATC speed interventions.

For each flight, the plugin uses speed-profile parameters such as climb, cruise, and descent CAS and Mach values. These are initialized with default values, but can also be set per flight.

Above the transition altitude, the logic determines whether the aircraft is climbing, cruising, or descending and applies the corresponding CAS or Mach target. Below the transition altitude, aircraft revert to a fixed low-altitude descent speed.

An important feature is that the plugin respects manual speed commands. Once a user-issued speed command is given, the plugin no longer overrides the selected speed. When VNAV speed control is restored, the automatic crossover logic becomes active again.

APPENDIX K
HOLDING PLUGIN

The `holding` plugin is used to model tactical holding patterns for aircraft that cannot absorb sufficient delay through speed control or vectoring alone. The plugin allows a holding pattern to be defined at specific waypoints, including the inbound radial, altitude limits, and maximum holding speed.

The plugin supports standard entry procedures by distinguishing between direct, parallel, and teardrop entry, based on the relative geometry between the inbound track and the radial to the holding fix. Once the aircraft reaches the holding waypoint, the plugin issues the required heading and direct-to commands through the BlueSky stack in order to fly successive holding legs.

Holding can be either indefinite or time-based. For time-based holding, the remaining required delay is monitored and translated into an appropriate leg length. If only a small amount of delay remains, the hold is cancelled and the aircraft continues its route. If the delay is still substantial, additional holding legs are generated recursively each time the aircraft passes the holding fix.

Preliminary Thesis (Previously Graded)

Abstract

The extension of the Arrival Manager freeze horizon, as mandated by the Single European Sky initiative, is intended to facilitate earlier delay absorption which decreases fuel burn and improves arrival flow predictability. However, extending the planning horizon also increases exposure to uncertainty in takeoff planning and trajectory prediction. These uncertainties cause so-called pop-up flights: aircraft that enter the Arrival Manager horizon after the sequence has been frozen, resulting in an increase in workload and fuel burn.

This thesis aims to support the implementation of an Extended Arrival Manager at Schiphol Airport by investigating the effects of takeoff and trajectory prediction uncertainty on the performance of the Extended Arrival Manager, and investigating and evaluating mitigation strategies. This preliminary thesis aims to investigate the Arrival Manager currently in use, the previous research conducted, and the characteristics and magnitude of the different uncertainties.

A quantitative analysis of real-world ETFMS Flight Data message is performed to characterize uncertainty in the pre-departure and en-route phases. The results show that both takeoff and trajectory prediction errors are significant, and that a longer freeze horizon significantly increases the likelihood of pop-up flights.

A high-fidelity, modular simulation environment is to be developed using the BlueSky air traffic simulator, integrating a trajectory predictor, Arrival Manager module, and ATC logic. The simulation supports stochastic modeling of uncertainties and enables Monte-Carlo experiments across multiple configurations. The mitigation strategies to be evaluated are a delayed slot strategy and optimized rescheduling.

List of Figures

List of Tables

Acronyms

Glossary

AAA Amsterdam Advanced Air Traffic Control System

AAH Active Advisory Horizon

AAT Advanced ATC Tower

ABI Advanced Boundary Information

ACC Area Control Center

ACT Activation Message

ADS-C Automatic Dependent Surveillance-Contract

AIP Aeronautical Information Package

AMAN Arrival Manager

ANSP Air Navigation Service Provider

APP Approach Control

ARSIM AMAN Research Simulator

ASAP Advanced Schiphol Arrival Planner

ATC Air Traffic Control

ATOT Actual Takeoff Time

ATS Air Traffic Service

B2B Business to Business messages

CBAS Cross Border Airspace

CCIS Closed Circuit Information System

CDM Collaborative Decision Making

COP Coordination Point

CP1 Common Project 1

CPR Correlated Position Report

CTA Control Area

CTOT Calculated Take-Off Time

CTR Control Zone

DDR Demand Data Repository

DMAN Departure Manager

DPI Departure Planning Information

E-AMAN Extended Arrival Manager

EAT Estimated Approach Time

EDD Electronic Data Display

EFD ETFMS Flight Data message

EH Eligibility Horizon

EPP Extended Projected Profile

ETA Estimated Time of Arrival

ETO Estimated Time Over

ETOT Estimated Takeoff Time

EU European Union

EUROCONTROL European Organisation for the Safety of Air Navigation

FCFS First-Come First-Serve

FH Freeze Horizon

FIR Flight Information Region

FL Flight Level

FMS Flight Management System

FUM Flight Update Message

HMI Human Machine Interface

IAF Initial Approach Fix

IFR Instrument Flight Rules

IPS Inbound Priority Sequencer

KNMI Koninklijk Nederlands Meteorologisch Instituut

KPA Key Performance Areas

KPI Key Performance Indicator

LIV Landing Interval

LLDA Low-Level Delay Absorption

LVNL Luchtverkeersleiding Nederland

MUAC Maastricht Upper Area Control

NM Network Manager

NM Nautical Mile

PDF Probability Density Function

QRC Quick Reference Chart

RECAT Recategorization of Wake Turbulence Categories

RETD Revised Estimated Time of Departure

RTA Required Time At

SARA Speed and Route Advisor

SES Single European Sky

SESAR Single European Sky ATM Research Joint Undertaking

SID Standard Instrument Departure

SSR Secondary Surveillance Radar

STA Scheduled Time of Arrival

STAR Standard Terminal Arrival Route

SWIM System Wide Information Management

TBO Trajectory Based Operations

TBS Time-Based Separation

TMA Terminal Control Area

TOBT Target Off-Block Time

TOD Top Of Descent

TP Trajectory Predictor

TTA Target Time of Arrival

TTG Time To Gain

TTL Time To Lose

TTLG Time To Lose or Gain

TTO Target Time Over

TWR Air Traffic Control Tower

UAC Upper Area Control

UCO Under Control

UTA Upper Control Area

UTC Coordinated Universal Time

WTC Wake Turbulence Category

XMAN Cross-Border Arrival Manager

1

Introduction and Research Proposal

Motivation

The airspace around major airports is busy and complex. To help manage the arrival stream into airports, Arrival Manager (AMAN) systems are in use at many major airports. The AMAN creates a planning sequence, a time at which aircraft may enter the Terminal Control Area (TMA) and an estimated Time To Lose (TTL) or Time To Gain (TTG). The goal is to reduce the workload of Approach Control (APP) by limiting the number of aircraft in TMA, reduce the workload of Area Control Center (ACC) by constructing a stable sequence and reduce the fuel burn by allowing for earlier delay absorption.

According to legislation and the Single European Sky (SES) program, the AMAN must increase the horizon at which it freezes its planning. The goal of the horizon extension is to reduce fuel burn by allowing instructions to aircraft to adhere to the Estimated Approach Time (EAT) to take place earlier, and thus require fewer holdings or a smaller amount of speed deviation, both reducing fuel burn. Another goal is to reduce pilot and controller workload. However, previous research has shown that the horizon extension could have the opposite effect on fuel burn and pilot and controller workload, when aircraft within the freeze horizon do not adhere to their takeoff planning. These aircraft then pop-up in the planning, requiring the controller to change the sequence to accommodate this aircraft. This increases the controller and pilot workload, as well as fuel burn.

Previous research has been conducted on mitigating the negative effects of pop-up aircraft. One of these studies investigated scheduling these aircraft at their expected arrival time, although the prediction of the expected arrival time needs to be quite accurate to reduce the impact of the pop-up aircraft. Suggestions have been made to investigate a Target Time Over (TTO) Initial Approach Fix (IAF) for these pop-up aircraft, to be able to schedule them in advance while they could adhere to this planning.

Previous research has conducted simulations that feature some assumptions that could have quite a substantial impact on the performance of the AMAN and Extended Arrival Manager (E-AMAN), such as the omission of Trajectory Predictor (TP) inaccuracy or the simulation of only one type of aircraft. Furthermore, the inaccuracy in the takeoff prediction was a set amount instead of a stochastic model. Quantification of inaccuracies relevant to the AMAN have been conducted, although studies have mainly focused on inaccuracy at the Coordination Point (COP).

To move forward with the extension of the horizon, further investigation is required into dealing with pop-up aircraft since the current solutions would not achieve sufficient mitigation of the negative effects if the horizon would be extended. To conduct the research with a more accurate model of the different prediction inaccuracies, a data analysis of the prediction inaccuracies must first be conducted.

Research Objective

This research aims to contribute to the development of Extended Arrival Management, specifically in aiding the extension of the freeze horizon at Schiphol airport considering uncertainty. This shall be achieved by developing and evaluating solutions to the negative impact of uncertainty on the arrival

management process, using currently in-use systems and data. To be able to evaluate the solutions to the negative impact of uncertainty, a high-fidelity simulation shall be developed, evaluating using parameters representing stability, accuracy, capacity and taskload.

Research Questions

Following the research objective presented above, the following main and sub research questions were formulated. The answers to the subquestions together aim to answer the main questions.

Main Research Question:

- How can the negative effects of uncertainty in takeoff and trajectory prediction on the performance of the extended arrival manager be mitigated?

Sub-Questions:

1. How can the uncertainty in takeoff and trajectory prediction be quantified?
2. What is the impact of takeoff uncertainty on the performance of an extended arrival manager?
3. What is the impact of trajectory prediction uncertainty on the performance of an extended arrival manager?
4. What mitigation strategies are available to reduce the negative effects of takeoff and trajectory prediction uncertainties on the extended arrival manager?
5. Are these mitigation strategies effective enough in improving the performance of the extended arrival manager considering uncertainty?
6. If the considered mitigation strategies are not effective enough in improving the performance of the extended arrival manager, which conditions concerning uncertainty should me met to effectively implement an extended arrival manager?

Preliminary Questions: Before diving into the research question and sub-questions presented above, the questions below shall be answered by a literature study and preliminary research.

1. What are the main concepts of an (extended) arrival manager?
2. How can uncertainty be measured in a relevant manner for the arrival manager?
3. How can the performance of an extended arrival manager be quantified?
4. What have previous investigations into the (extended) arrival manager found concerning the effect of pop-up aircraft and mitigation of these effects?
5. What simplifications and assumptions can increase the fidelity of the simulation while maintaining reasonable development costs?

Approach

The following figure presents a schematic overview of the steps taken in this research. The research is split into three phases: preliminary, development and experiment. The preliminary contains the background and literature studies, as well as the study into the uncertainty of flight planning. The development mainly entails the development of the AMAN simulator as well as the Air Traffic Control (ATC) module mimicking the AMAN users. Verification and validation will be done throughout this phase. Finally, the experiment phase consists of conducting the experiments and the further processing of the results.

The investigation into the uncertainty of flight planning will be carried out using data sources existing at Luchtverkeersleiding Nederland (LVNL). The AMAN simulator will be made using the BlueSky Air Traffic Simulator [1]. The results from the investigation into the uncertainty will be used in the scenarios, and used in the development of the operational concepts to reduce the impact of the uncertainty. chapter 5 elaborates upon the development required to perform the simulations.

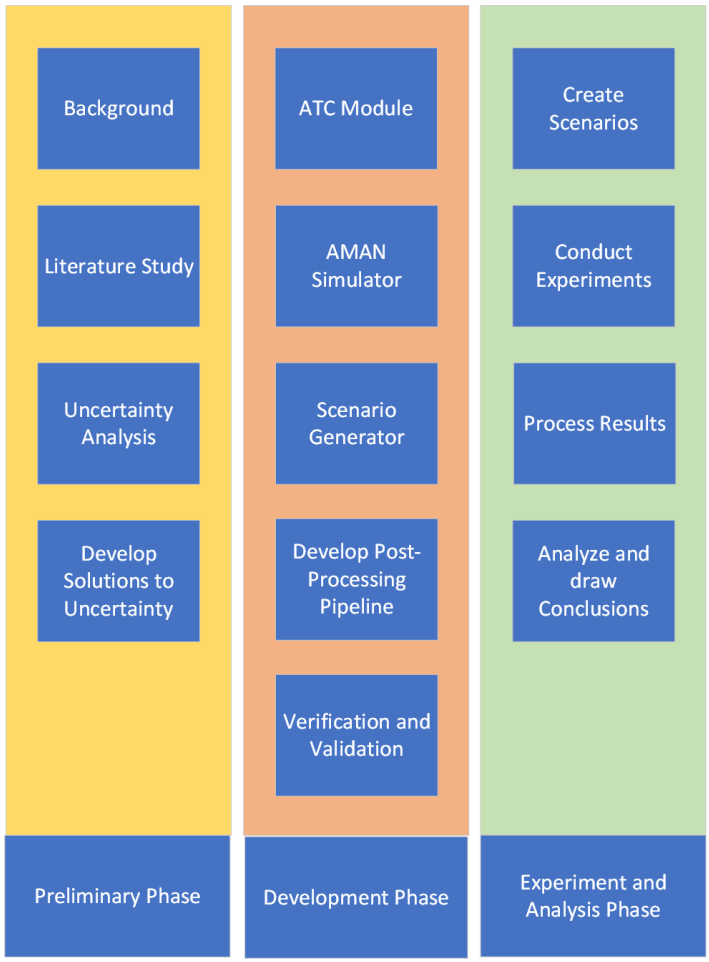

Figure 1.1: Research Phases and Steps

2

Background

This chapter provides the reader with the necessary background information to the following chapters. Air Traffic Control, the messages sent during the arrival process, the Arrival Manager and the systems within are explained in this chapter, beginning with airspace design in section 2.1. The next section elaborates upon the messages sent among different organizations. An introduction to the Arrival Manager, its use and future is given in section 2.3. Finally, pop-up aircraft are explained in section 2.4.

2.1. Airspace Design and Arrival Routes

Air traffic control is the coordination of aircraft during their flight. The responsibility of this coordination of traffic is usually taken up in individual countries, which have one or a few Air Navigation Service Provider (ANSP). An ANSP manages a Flight Information Region (FIR), which is the airspace above a large part of- or an entire country. In the Netherlands, the Amsterdam FIR is managed by Luchtverkeersleiding Nederland (LVNL). In Europe, since there is a large degree of fragmentation of the airspace as countries are quite small, parts of the higher elevation airspace has been combined and are managed by European Organisation for the Safety of Air Navigation (EUROCONTROL). The airspace is divided into many different categories, ranging from uncontrolled to requiring Air Traffic Control (ATC) guidance. In general, the latter is near large airports, while the former is at low levels further away from airports.

The control zones around large airports, such as Schiphol, are divided further into several parts generally following the different phases of flight. A short overview is given below [2].

- The upper control area is used by en-route traffic, and is managed by Upper Area Control (UAC). In the Netherlands and surrounding countries, Maastricht Upper Area Control (MUAC) is responsible for the coordination of air traffic. The Upper Control Area (UTA) ranges from 24.500 feet up to 66.000 feet [3].
- The Control Area (CTA) is the airspace between approximately 10.000 feet and 24.500 feet. Most aircraft in this airspace departed an airport within the control area, or are headed towards one. The control area is managed by Area Control Center (ACC). Aircraft enter the Control Zone (CTR) at Coordination Point (COP)'s and exit at the Initial Approach Fix (IAF), when they are handed over to Approach Control (APP). ACC must ensure separation between aircraft, while also managing any delays required by the limited capacity in the Terminal Control Area (TMA). The instructions of the Arrival Manager (AMAN) are used to coordinate the flow of aircraft from the CTA into the TMA.
- The TMA is managed by Approach Control (APP) and is used by controllers to bring the flow of traffic from different IAF's together onto final approach. The airspace allows for vectoring aircraft to obtain the correct separation. The airspace ranges from 3.000 up to 10.000 feet. The traffic consists of several streams of high density traffic, which must be merged safely. An AMAN is used at many major airports to help manage the traffic flows, and divide slots among the incoming aircraft. The supervisor of APP operates the AMAN according to the current traffic situation. This information is then showed to the controllers at ACC.

- The airspace directly around the airport, the CTR, is managed by the Air Traffic Control Tower (TWR) controllers. Ensuring separation at takeoff and landing is the responsibility of the tower controllers. The tower controllers control aircraft from being established on the approach up to exiting the runway, and from entering the runway up to an altitude of approximately 3.000 feet. Only a few nautical miles around the airport are controlled by the tower.
- Finally, the ground controllers coordinate pushback, startup and taxiing aircraft.

These different types of airspaces are visually represented in Figure 2.1. Note that the different airspaces are often split up into sectors. For example, ACC in the Netherlands is split up into five sectors, which then can be combined into larger combinations.

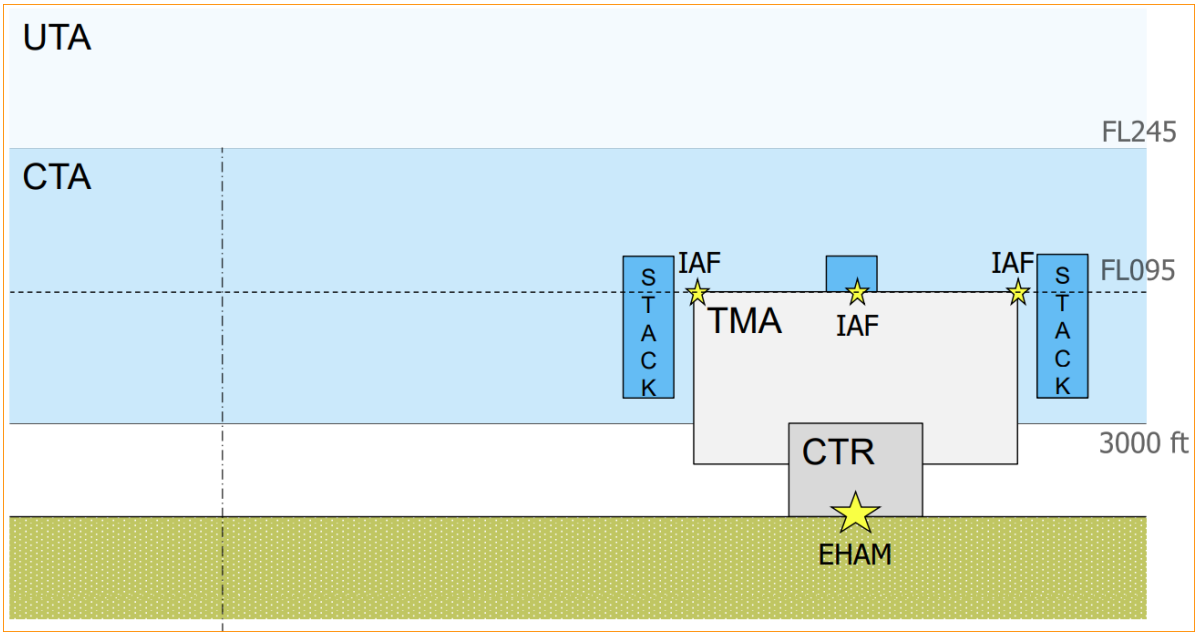

Figure 2.1: Vertical Schematic of the airspace structure [2]

The aircraft are handed over to ACC by an adjacent center at the COP. Controllers can deviate from the standard route if they see reason to, in collaboration with the adjacent center. In the Dutch airspace, aircraft follow standard routes in the CTR, thus from the COP up to the IAF. They are used to simplify the separation of traffic by ACC, while also allowing for the merging of streams of traffic at the three different IAF's. Controllers are allowed to instruct aircraft to deviate from these Standard Terminal Arrival Route (STAR)'s to maintain separation or increase efficiency. For example, a waypoint can be skipped to shorten the flighttime, or a detour can be assigned if a delay is required. Figure 2.2 shows the chart of the STAR's used for aircraft inbound to Schiphol.

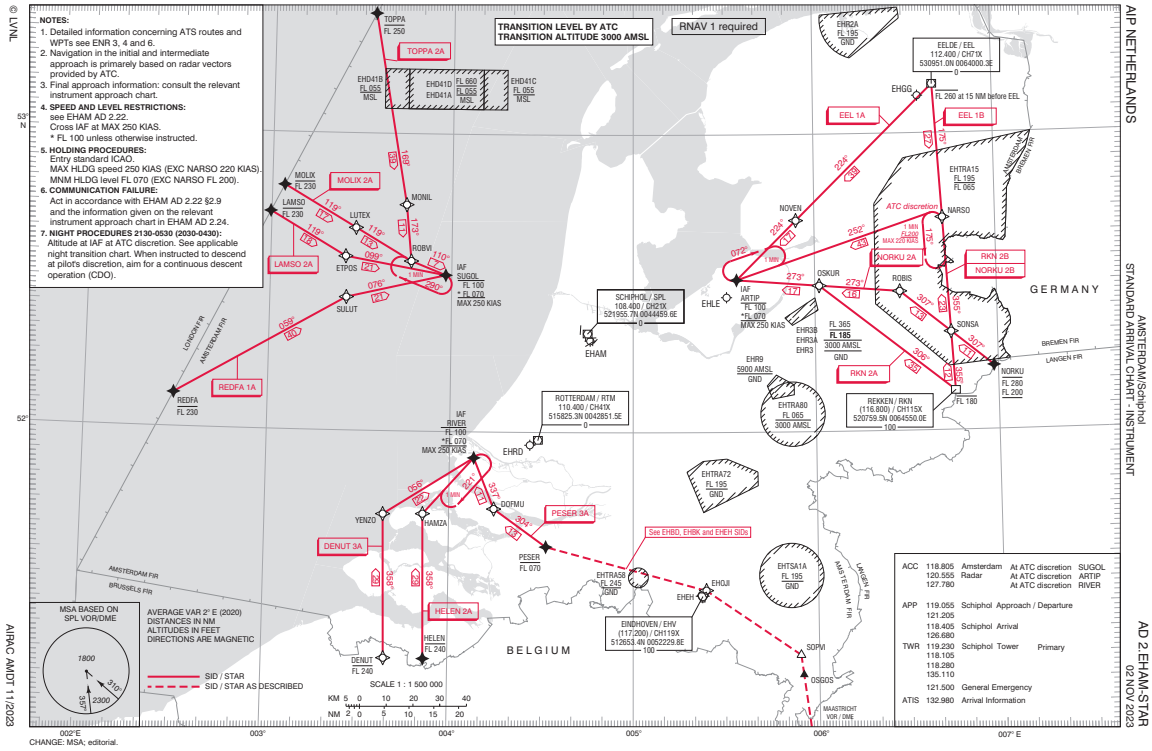

Figure 2.2: STAR chart for Schiphol airport [4]

The three IAF's are RIVER, SUGOL and ARTIP. In principle, all inbound Instrument Flight Rules (IFR) traffic is assigned an IAF, based on their scheduled route. If required, holding patterns or "stacks" are used to delay aircraft before they enter the TMA. The holding patterns are situated at the IAF's, with a minimum altitude of Flight Level (FL) 070. There is an additional holding pattern at the NARSO waypoint at FL 200. The latter is used infrequently. The holding patterns at the IAF's all feature 1 minute legs and turns, meaning an addition of 4 minutes to the flighttime of an aircraft if a single holding pattern is traversed. As previously noted, when an aircraft passes the assigned IAF they are transferred from ACC to APP.

2.2. Arrival Process Messages

This section provides an overview of the different messages exchanged between LVNL, Schiphol, Eurocontrol and the adjacent centers.

Eurocontrol is the Network Manager (NM) which entails multiple responsibilities such as balancing capacity and demand, and being a central coordinator of communication as a Business to Business messages (B2B) service. The latter entails the distribution of flightplans, slot information, Aeronautical Information Package (AIP) services, and flight update messages. Next to that, Eurocontrol also maintains multiple other databases for other purposes, such as the Demand Data Repository (DDR) database.

Currently LVNL receives the flight update messages through the ETFMS Flight Data message (EFD) format [5]. The EFD format has quite some datafields, in the following paragraphs the most important fields are discussed. Each EFD message concerns a single flight, and for each flight multiple updates are given. Each datafield is repeated in each message.

2.2.1. General Information and Route

In every EFD message some general information is given, such as the callsign, the flightplan ID, aircraft registration, aircraft type and message timestamp. The route is also included, and contains the departure and arrival airports, the estimated off-block time and the planned route which consists of waypoints, altitude and planned time. The aircraft can deviate from this route at ATC discretion.

2.2.2. Updates

EFD messages are a constant stream of data of all flights inbound, outbound and overflying the Dutch airspace. Multiple messages are published per flight. The first message is published approximately three hours before departure. Updates are published every 60 minutes up to 2 hours before Target Off-Block Time (TOBT), after which every 15 minutes an update is published until the flight is completed. Updates are also published if the departure planning changes, if the planned route is changed or if the estimated time along the route is changed by 5 minutes or more.

Flight State (Code)	Description
Filed (FI)	Flight plan filed (≈ 3 h before departure)
Slot Allocated (FS)	Regulation placed, no slot yet assigned
Slot Issued (SI)	Regulation placed and slot assigned
Tactical Activated (TA)	Assumed airborne by Network Manager (post-EOBT)
Suspended (SU)	No known planning, flight not cancelled
Cancelled (CA)	Flight cancelled
ATC Activated (AA)	Confirmed airborne by ATC
Terminated (TE)	Flight completed

Table 2.1: Key ATFM flight states and their meanings

2.2.3. Airport Collaborative Decision Making

Multiple airports near to Schiphol have implemented Collaborative Decision Making (CDM). The goal of CDM is to improve efficiency, predictability and punctuality of the air traffic in Europe [6]. The CDM process is based on milestones and the communication of the achievement of these milestones, per flight. The milestones are for example the acceptance of the flightplan or takeoff from the aircraft from an outstation. In theory, the predicted departure times by airports that have implemented CDM are more accurate than airports without the CDM process.

The milestone communication to LVNL is done via EFD messages, with several message types. At each update, the prediction of the departure time should be improved. Since the Departure Planning Information (DPI) is part of the EFD, if a new planning is known, an update will be published. Multiple messages with the same state could be given. The following different messages and accompanying CDM states are given, in chronological order:

Milestone	CDM State	Trigger / Timing
Predicted Departure Planning Information (P-DPI)	DPIEXPECTED	Flight-plan submission (≥ 3 h before EOBT)
Early Departure Planning Information (E-DPI)	ESTIMATED	Flight-plan, flight and slot correlated
Target DPI – Targeted (T-DPI-t)	TARGETED	2 h before TOBT or departure from outstation
Target DPI – Sequenced (T-DPI-s)	PRESEQUENCED	40 min before TOBT, based on refined TSAT
ATC Departure Planning Information (ATC-DPI)	ACTUALOFFBLOCK	Push-back under ATC control

Table 2.2: Key CDM departure-planning milestones and their triggers

CDM airports are all major airports in Europe, with an overview given in [7]. There is also a less sophisticated type of airport: the Advanced ATC Tower (AAT). These airports only provide the Predicted Departure Planning Information (DPIEXPECTED) and the ATC Departure Planning Information (ACTUALOFFBLOCK) messages.

For reference, in chapter 4 it is shown that from all inbound traffic to Schiphol, 1.41% departs within 180 NM from a CDM airport, 0.81% from an AAT airport, and 1.50% from a standard airport. This results in a total of 3.72% of inbound flights originating within 180 NM. The 120 NM and 180 NM horizons are relevant horizons to the AMAN, which are further elaborated upon in section 2.3.

The figure below in Figure 2.3 shows which airports have CDM and AAT capabilities implemented, together with a 120NM and 180NM radius around Schiphol. The 120NM radius is an approximation of the current freeze horizon, which will be explained later in subsection 2.3.3.

2.2.4. Correlated Position Report

The EFD messages also contain a datafield in which the eventtype of the message is given. An example is the periodic transmission (PTX). For the data analysis done in chapter 4, the Correlated Position Report (CPR) is also relevant. This event shows that the aircraft position is correlated with EUROCONTROL systems, and the flown route and the timestamps at the waypoints in the route have been updated [8]. The CPR does not contain the current coordinates of the aircraft.

2.2.5. Handover Messages

Next to the EFD messages, there are also two types of messages received from the adjacent centers: Advanced Boundary Information (ABI) and Activation Message (ACT) messages. The ABI message is sent approximately 30 minutes before handover at the COP. The ACT message is sent a few minutes before handover. Both messages contain information on the planned route and current location of the flight. Both messages also allow the correlation of Secondary Surveillance Radar (SSR) to the flightplans, and lets Amsterdam Advanced Air Traffic Control System (AAA) and the AMAN at LVNL show the aircraft on the relevant displays.

2.3. Arrival Manager

The AMAN is a tool used by air traffic controllers to optimize the flow of arriving traffic while avoiding a traffic overload in the TMA. It provides a planning of aircraft for each runway on a short horizon and a planning of TMA entry. Eurocontrol defines the AMAN as the following: ”The general objective of an Arrival Manager (AMAN) is to provide electronic assistance in the management of the flow of arriving traffic in a particular airspace, to particular points, such as runway thresholds or metering points.”[9]. AMAN systems generally strive to combine and balance safety, capacity, efficiency and the environment as areas of improvement by functioning as a support tool, at which they are overall successful at. The system provides a sequence with a scheduled or expected time, and it is often controlled by a designated controller, with information shown to different sectors [9].

This section will provide more insight into the implementation of the AMAN system at LVNL for inbound aircraft to Schiphol airport. Advanced Schiphol Arrival Planner (ASAP) consists of three subsystems: the Human Machine Interface (HMI), which is integrated into the AAA HMI, the Trajectory Predictor (TP), and the scheduler.

2.3.1. AMAN goals and benefits

As previously stated, the AMAN has the goal of combining safety, capacity, efficiency and the environment performance areas of the arrival process, while keeping them balanced, in this case at Schiphol airport.

The main priority behind safety is managing workload for APP controllers, who are managing the TMA. The AMAN provides a regulation on the amount of traffic in the TMA, reducing workload of the APP controllers. This can be seen in the primary goal of ASAP: ”to limit the amount of inbound traffic in the TMA, and plan such that the inbound aircraft in the TMA can be handled as a stable traffic flow” [10]. Effectively, it keeps aircraft delay and aircraft count within the TMA to a minimum by requiring the delay absorption to take place outside the TMA. The ACC controllers working with the system need it to be trustworthy and provide viable plannings, resulting in the secondary goal of ”planning realistic and stable Estimated Approach Time (EAT)’s, as early as possible” [10]. This requirement makes adherence to the planning by the ACC controllers viable in terms of workload.

As the stream of arriving aircraft is regulated, the arrival capacity of the airport is also impacted. If the

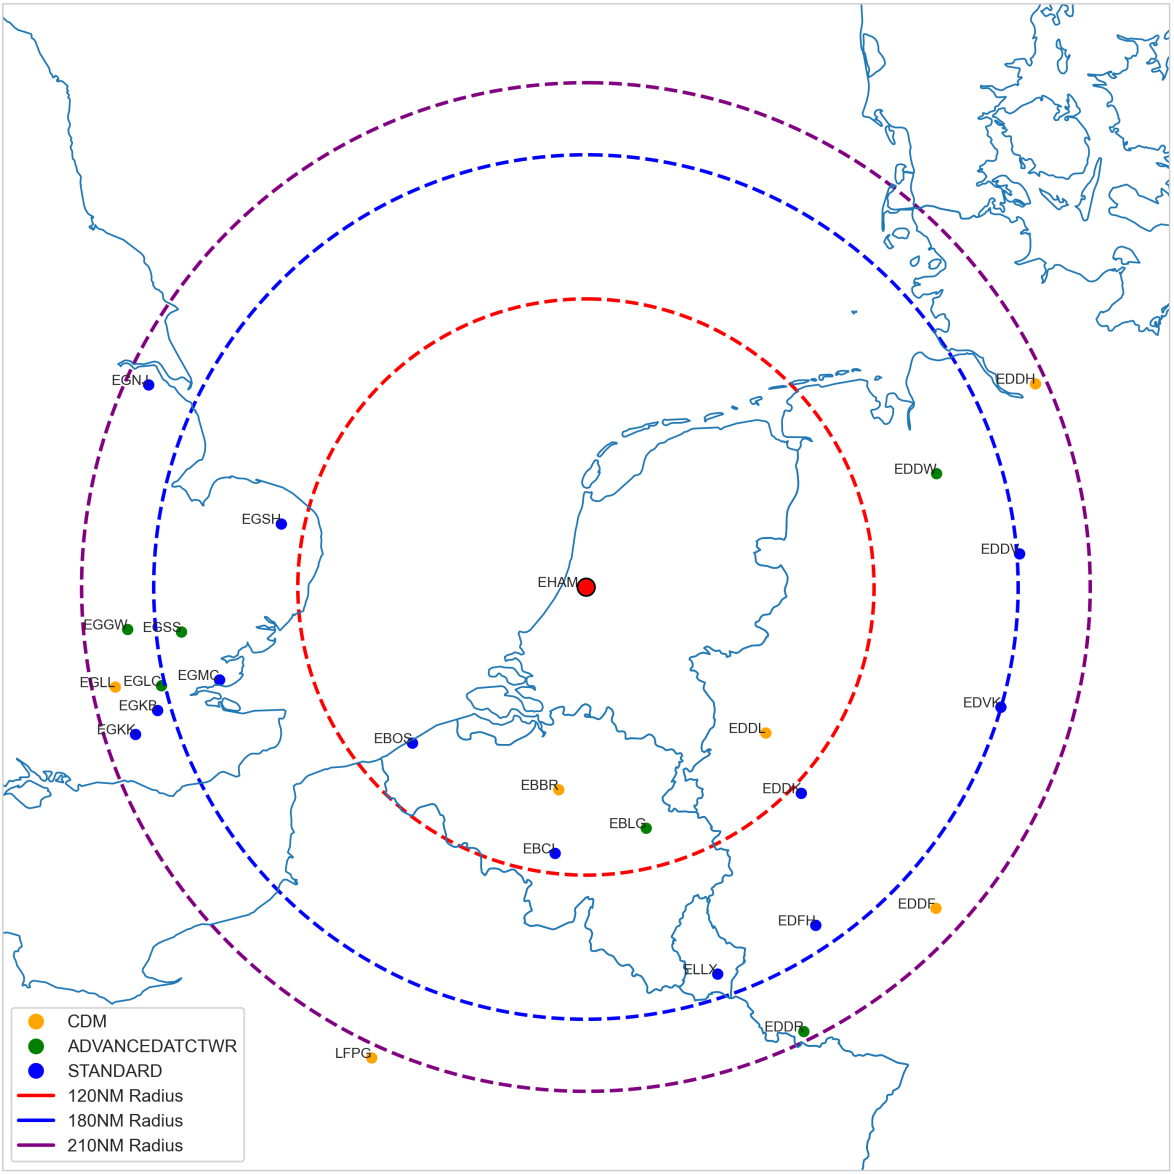

Figure 2.3: Map of the current (120NM), extended (180NM), and near-horizon (210NM) horizons from Schiphol with pop-up contributing airports

implementation is done properly, the inter-arrival time could be reduced at the runway threshold as the arrival stream is more consistent. This would increase capacity. As the AMAN plans 14 minutes ahead [11], (part of) the delay can be absorbed earlier than when no AMAN would be in place. This improves the efficiency and environmental impact of the arriving traffic by reducing the amount of Low-Level Delay Absorption (LLDA).

2.3.2. Scheduler

The scheduler provides two outputs per aircraft: the landing slot and the EAT. The landing slot is when the aircraft is planned to touch down, while the EAT is the time at which the aircraft is allowed to enter the TMA. The landing slot sequence scheduler in ASAP is relatively simple, it is currently First-Come First-Serve (FCFS), based on the Estimated Time of Arrival (ETA) calculated by the ASAP TP. The calculation of the EAT requires the Estimated Time Over (ETO) IAF, assigned runway and the Landing Interval (LIV). The LIV is based on current weather conditions and Wake Turbulence Category (WTC), and is used in combination with the ground speed to provide a time-based LIV. The equations used is given below in Equation 2.3. A table is given below in Table 2.3. The entries without numbers correspond to a radar separation minimum of 2.5NM.

Table 2.3: RECAT-EU Wake Turbulence Separation Minima (in NM) [12]

Leader ↓ / Follower →	**A** Super	**B** Upper H	**C** Lower H	**D** Upper M	**E** Lower M	**F** Light
A (Super)	3	4	5	5	6	8
B (Upper H)	–	3	4	4	5	7
C (Lower H)	–	–	3	3	4	6
D (Upper M)	–	–	–	–		5
E (Lower M)	–	–	–	–	–	4
F (Light)	–	–	–	–	–	3

The scheduler plans the aircraft landing sequence and provides slots for aircraft based on the ETA. The basis is a FCFS schedule, with a LIV added to maintain separation according to WTC and wind, which can be seen in Equation 2.1. If the ETA of the aircraft is later than the slottime, the slottime is changed to the ETA. The LIV is distance based, and needs to be converted to a time-based LIV, the equation can be seen in Equation 2.3. It is based on the ground speeds, and it must be satisfied in all parts of final approach.

The EAT is simply the slottime minus the predicted flighttime from the IAF to the runway. Currently, the approach margin is two minutes, allowing aircraft to be handed over from ACC to APP plus or minus two minutes from the EAT. This allows for a bit less rigid planning in ACC airspace. From the EAT the Time To Lose (TTL) or Time To Gain (TTG) is calculated. This is shown at ACC to help with planning to hand over the aircraft at the correct time. Finally, the approach manager is able to change the schedule if required. In subsection 2.3.7 these changes will be elaborated upon.

$$T_{slot} = T_{previous} + LIV_{timebased} \tag{2.1}$$

$$\text{if ETA} > T_{slot} : T_{slot} = ETA \tag{2.2}$$

$$LIV_{timebased} = \frac{LIV_{distance}}{\text{Calculated Ground Speed}} \tag{2.3}$$

$$EAT = T_{slot} - T_{\text{flighttime IAF-Runway}} \tag{2.4}$$

$$EAT_{window} = T_{slot} + / - 2minutes \tag{2.5}$$

$$TTL/TTG = EAT - ETO\ IAF \tag{2.6}$$

The scheduler also takes into account the runway configuration at Schiphol. When two runways are in use, the aircraft from two IAFs are planned on one runway, and the aircraft from the other IAF are planned on the second runway. This configuration, and which runways are in use are at the discretion of the approach planner.

2.3.3. Horizons

There are three main horizons considering the AMAN [13].

- The Eligibility Horizon
- The Freeze Horizon
- The Active Advisory Horizon

The Eligibility Horizon (EH) is the horizon at which the AMAN starts calculating the ETA, and thus the slot and corresponding sequence and EAT. The EH is currently limited by the SSR range and the ABI message from the adjacent centers. Upon crossing the Freeze Horizon (FH), the sequence and EAT is fixed for that aircraft. The current FH is at 14 minutes prior to ETO IAF for each aircraft. The Active Advisory Horizon (AAH) is where the aircraft start receiving instruction from the ATC to adhere to the EAT planning. Currently, this is effectively at the FIR boundary, when ACC comes into contact with the aircraft inbound to Schiphol. Occasionally, communication with an adjacent center occurs which moves the effective AAH further away. This occurs when aircraft need to absorb a lot of delay.

It can be noted that the horizons can be viewed in terms of time or distance. The regulations presented in [14] clearly state a distance. Though, since the systems of LVNL work with a time based horizon [11], it is likely that a horizon extension will consist of a time-horizon extension.

2.3.4. Trajectory Predictor

A TP predicts when an aircraft is at each waypoint and at the runway. Especially the estimated time at the IAF, or ETO IAF is relevant to ASAP. The TP of ASAP has two parts: a planning TP and an executive TP. The planning TP uses standard speeds and trajectories of aircraft to predict when an aircraft is ETO stack and the ETA of each aircraft. If speed instructions have been given to aircraft, the executive TP uses the instructed speed to calculate the section towards the IAF.

The model of the TP used falls somewhere in between the nominal and intent based state projection method found in [15]. Nominal speeds and altitudes are used, together with a nominal route in the TMA, depending on the IAF and runway, and a nominal route up to the IAF. If the air traffic controller gives a speed instruction, it will be used to calculate the arrival time at the IAF, by switching to the executive TP, partially based on intent. The aircraft performance is based on a large amount of radar data, and divided according to aircraft type and operator [10]. This is stored in an XML database.

2.3.5. Human Machine Interface

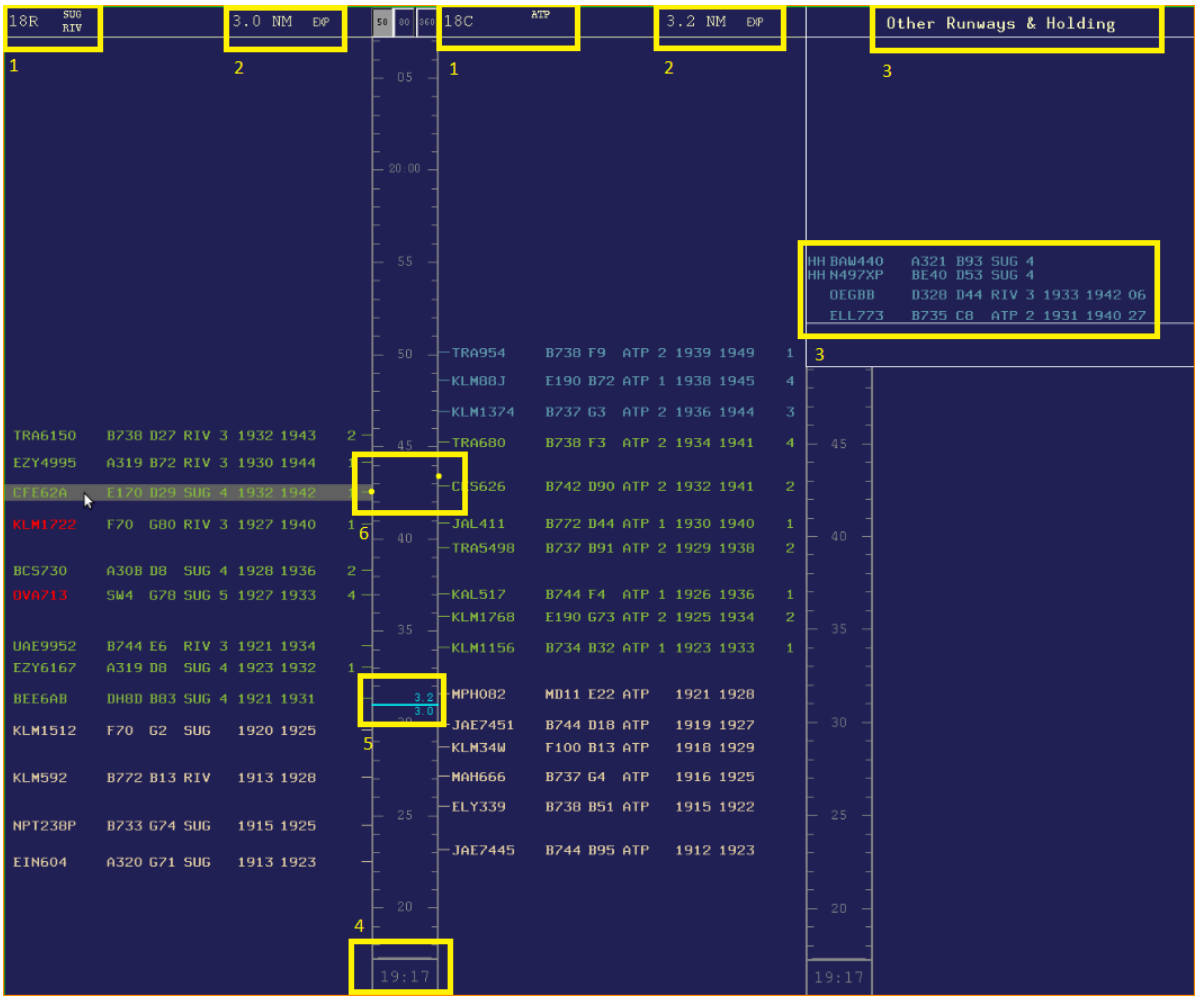

Figure 2.4: ASAP display example [10]

The HMI of ASAP is shown in two formats. Most importantly, the format for the approach planner. An example of this display is shown in Figure 2.4, with two primary runways used for inbound aircraft. The so-called Electronic Data Display (EDD) is split up in three columns: two for the primary runways, and one for the phased-out runway together with other runways and holding. The aircraft are represented in a similar fashion to the obsolete paper ”flight-strips”, with callsign, type, gate, IAF, sector (1-5), EAT, ETA and TTL or TTG in minutes. In the other runways bay, the assigned runway is also shown, and if the aircraft does not pass an IAF, the departure airport is shown. The latter happens when a flight departs from Rotterdam for example.

The color of the aircraft shows the status. If it is blue, the aircraft is pre-planned. A green indication shows that the aircraft has been planned, and the planning has been frozen. White indicates that the aircraft is Under Control (UCO) at APP, which also removes the delay indication. The red aircraft labels indicate aircraft that have entered the stable planning after it has been frozen, indicating pop-up aircraft. The pre-planned aircraft in blue are not given commands to absorb their indicated delay, as their planning is not fixed yet. These are shown to indicate the future demand to the approach planner [10].

1. Main runways and respective IAF’s
2. Static Landing Interval
3. Other runways and holding bay
4. Current Coordinated Universal Time (UTC)
5. Approximate handover point
6. Landing slot change if allocated runway is switched

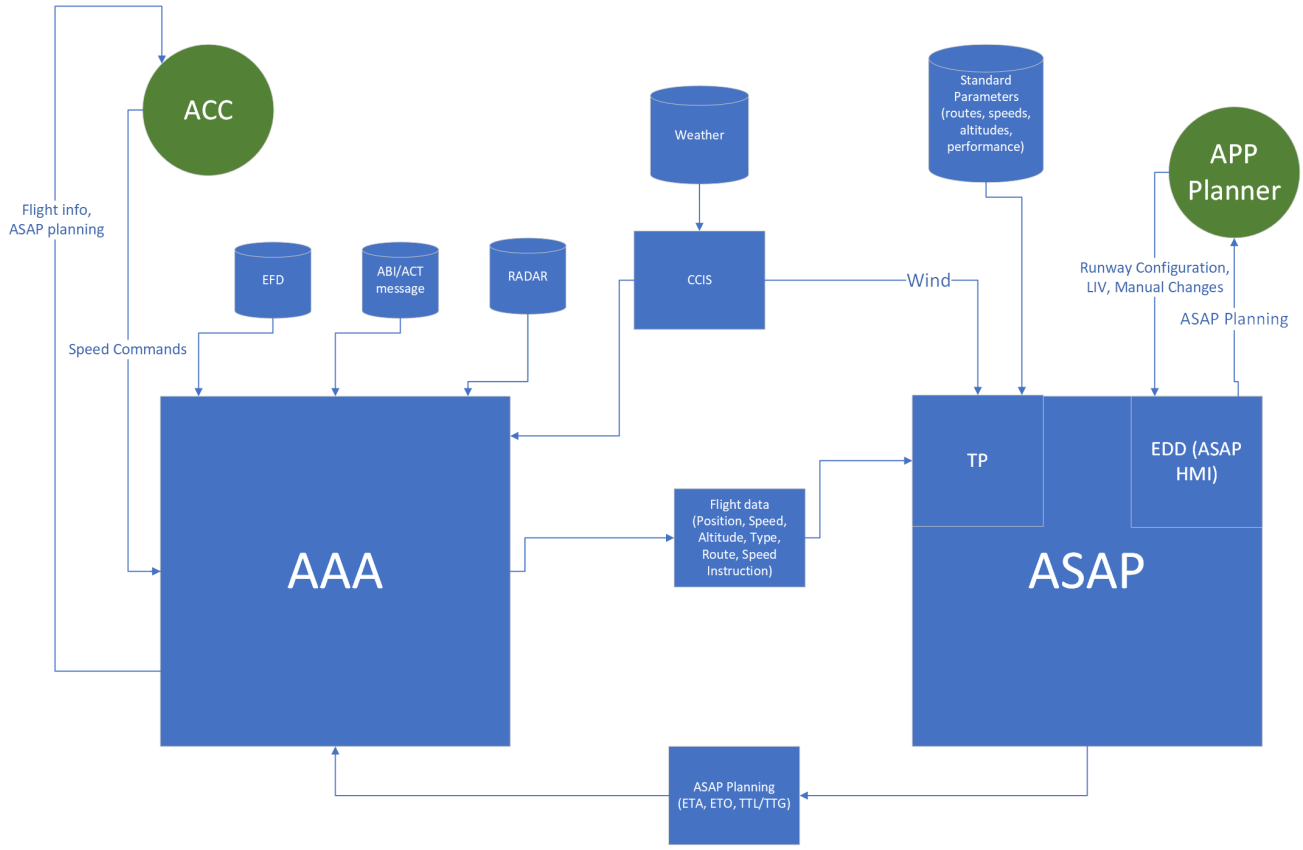

Figure 2.5: Schematic overview of ASAP

Note that Figure 2.4 shows the ASAP HMI before the implementation of Recategorization of Wake Turbulence Categories (RECAT) Time-Based Separation (TBS). The implementation of this recategorization has an effect on ASAP, on average allowing a higher capacity in windy conditions, as the amount of time separation is decreased at the runway.

The second HMI type is for the ACC controllers. The information shown is much more concise, as their role is to only guide aircraft according to the planning in their sector, not to influence the planning itself. The data is sent to the AAA system, to be shown on the main radar display. It includes the callsign, aircraft type, EAT and the time to lose or time to gain, in minutes and seconds. Flights that are not in the sector of the controller, pre-planned or handed over to APP are not visible.

2.3.6. Inputs

Figure 2.5 gives a general overview of how AAA and ASAP work together. AAA is the main system of LVNL to handle traffic, and works in collaboration with ASAP. AAA shows the radar screen used by controllers, while also being the main system for the processing of messages. The output of ASAP is integrated in AAA for the ACC controllers. The approach planner has a separate screen, called the EDD. This is further elaborated in subsection 2.3.5.

ASAP takes inputs from several datasources to create a realistic and accurate planning. First of all, the radar correlation data is fed from AAA to the TP to use accurate aircraft location to calculate the predicted ETA. The TP also uses aircraft and airline specific performance information, with information such as cruise mach, descent speed, and different speeds on final approach. It also requires routes and expected flight levels.

The SSR data is correlated with the filed flight plans, and the ABI and ACT messages. In the case of departures in the TMA, the ABI is replaced by the Revised Estimated Time of Departure (RETD) message, after coordination has been done between the controllers. The flightplan also contains the

WTC to determine the landing interval, and the STAR from the FIR to the IAF.

The TP also requires wind and temperature information, being given in five levels per sector and five for the TMA, with the source being Koninklijk Nederlands Meteorologisch Instituut (KNMI) via Closed Circuit Information System (CCIS).

According to the wind conditions, preferential runway use, as well as departure and arrival demand, the approach planner chooses a runway configuration. This is done directly in ASAP. The approach planner also determines the standard LIV, which is mainly dependent on the wind and visibility conditions. This is done using a Quick Reference Chart (QRC), [16]

Finally, commands made by ACC to aircraft are also taken into account in the TP. These are processed via the AAA system.

2.3.7. Practical Use

In real life, ASAP is mostly used by the approach planner and ACC. ASAP creates the planning, while the approach planner monitors and optimizes. ACC instructs traffic such that the planning is adhered to.

The current window to hand over aircraft is +/- 2 minutes [10]. Though, most ACC controllers aim to hand over aircraft as early as allowed without continuous speed instructions. Effectively, they aim for handover to APP at 1.5 minutes before EAT. The TTL or TTG which the aircraft still has when crossing the IAF is absorbed in the TMA using instructions from APP. The 2 minutes margin is mostly important to the ACC controllers, as they are required to hand over each aircraft within this margin. They have a few ways to obtain this accuracy. First, if the aircraft would arrive at the IAF too early, the aircraft can be put into a holding pattern, usually at the IAF, entering at a higher altitude and possibly descending to FL070, as more aircraft enter the holding stack. This is quite fuel inefficient and increases the taskload of the pilots and controllers. If the possibility is there, the air traffic controller instructs the aircraft to slow down, and possibly decrease the altitude and thus ”linearly” absorb the delay. The downside to this is that there is not a lot of delay that can be absorbed in this way, if the freeze horizon is only 14 minutes. A third method is vectoring the aircraft, increasing the path length. To decrease the flight time, and thus gain some time, the controllers are able to speed up the aircraft or to let the aircraft fly a directer route. These actions are usually done when aircraft are in their control, and occasionally when an inbound aircraft is still under control at the adjacent center.

The approach planner controls ASAP. The system automatically plans the aircraft in the available slots and calculates the respective delay required. The approach planner can then manually intervene. This could be done if one runway is facing delay, while the other runway is not: the approach planner then switches the slot of the aircraft, preferably as soon as the planning has been frozen. Furthermore, the approach planner also arranges that Pop-up flights and flights from airports within the TMA receive a slot and are included in the planning. This requires insight and communication with adjacent centers or control towers. These flights appear in the ”other” bay in the EDD.

Front loading also occasionally occurs. If the approach planner decides to front load, aircraft are planned before their ETO IAF, and thus have to speed up or shorten their route. If the approach planner decides to front load multipe aircraft, the function ”expedite margin” is used. All selected aircraft receive a slot 1 minute before their current slot. This is done if delay in the future is expected, while there currently is capacity available. Specific aircraft can be manually changed as well, and expedited up to 4 minutes before its’ initial slot time. This can be achieved by issuing clearances to aircraft to bypass certain waypoints in their route. All of these actions are based on the experience of the controller, and are not calculated by the system. They also often require coordination with adjacent centers, or ACC.

The approach planner is also responsible for selecting the runway use. The expected demand is shown in time blocks in terms of the amount of runways required, for inbound and outbound aircraft. The approach planner then takes the weather information and determines which runways are suitable and available for use, and then selects which runways are in use for landings. For planning purposes, runways can be switched, phased in or phased out. When the change is selected by the approach planner, the switch takes place after the current freeze horizon. The flights planned on a runway that

is phased out move towards the ”phase out” bay on the EDD.

Finally, the approach planner is responsible for managing and monitoring the general planning in ASAP. They should make sure ACC and APP use the planning appropriately, while also solving sequence mistakes, and optimizing the planning overall.

2.3.8. Future of the AMAN

The current AMAN, ASAP, has been operational for a few years. In the near future, changes to the planning horizon are required to maintain compliant with European law with the goal of improving the the efficiency of the arrival process across multiple large airports in Europe. The following quote is part of the CP1 requirements, which are part of the Single European Sky (SES) program. ”Extended AMAN systems must provide arrival sequence time information and associated advisories into en-route ATC systems to a minimum of 180 nautical miles from the arrival airport as well as into ATC systems of airports impacted by the extended AMAN horizon, unless a shorter distance is recommended in the deployment programme” [14]. The implementation must be completed before the 31st of December 2024. More insight in the regulation is provided in [17], concerning airports that lie within the freeze horizon. In short, if there are airports that significantly impact the AMAN because they lie within the planning horizon, they can be excluded from coordination for the Extended Arrival Manager (E-AMAN) if they have not implemented CDM or AAT. If they are using the CDM or AAT concepts, the Air Traffic Service (ATS) should display relevant and up-to date DPI, and if possible include the information in the arrival planning. The latter is currently being implemented at LVNL.

The required horizon extension will cause some problems, since more pop-up aircraft will occur. Previously, extending the horizon to 16 minutes has been tested, although this was deemed unsuccessful. A further elaboration on pop-up flights is given below in section 2.4.

In the future, this horizon extension is a step towards a cross-border arrival manager, which in turn is a step towards Trajectory Based Operations (TBO). It all has the goals of increasing fuel efficiency, reducing safety risks, increasing capacity, and futureproofing the air traffic management system. [11]

Extending the AMAN horizon introduces some issues. Four main issues can be identified [15]:

- Pop-up flights due to the inaccuracy of estimated takeoff times.
- Inaccuracy in estimated flighttimes due to a longer prediction horizon of the trajectory predictor.
- The division of airspace increasing the amount of communication required to facilitate earlier speed instructions, together with the increased amount of non-adjacent centers to communicate with.

In the next section, an introduction is given to pop-up flights. The inaccuracy of the trajectory predictor is discussed in the next chapter, since quite some statistical research has been done into this issue. The division of airspace can be quite a difficult problem to solve. The airspace around the Netherlands is divided into the upper and lower airspace. The upper airspace is controlled by MUAC, also including Belgium, Luxembourg, and a part of Germany. Increasing the horizon while also providing active advisories to aircraft within the horizon would create difficulties, since it would include French airspace, not controlled by MUAC or another adjacent center. The coordination workload of the controller would increase, together with the requirement that any coordination between Reims UAC would need to be set up.

The radar range is also an issue. Currently, LVNL only uses radar data from their own sources. This limits the information to the actual radar range. An approximate SSR range map is shown in Figure 2.6. The range is not large enough to provide coverage of the required 180Nautical Mile (NM) radius required if the AMAN freeze horizon is extended. The current approximate range is 150NM. In the future, System Wide Information Management (SWIM) is expected to be implemented. This would at least ensure that LVNL receives the SSR data from other ANSPs. Another option would be to rely on EFD messages from aircraft that are in-air, but not within radar range. Currently, LVNL is in the process of using the EFD messages in ASAP to improve the predictability of the amount of arrival traffic. These two solutions would allow for the extension of the freeze horizon in terms of coverage of in-flight aircraft.

The SWIM program, which aims to enable longer range AMAN through the sharing of data required for

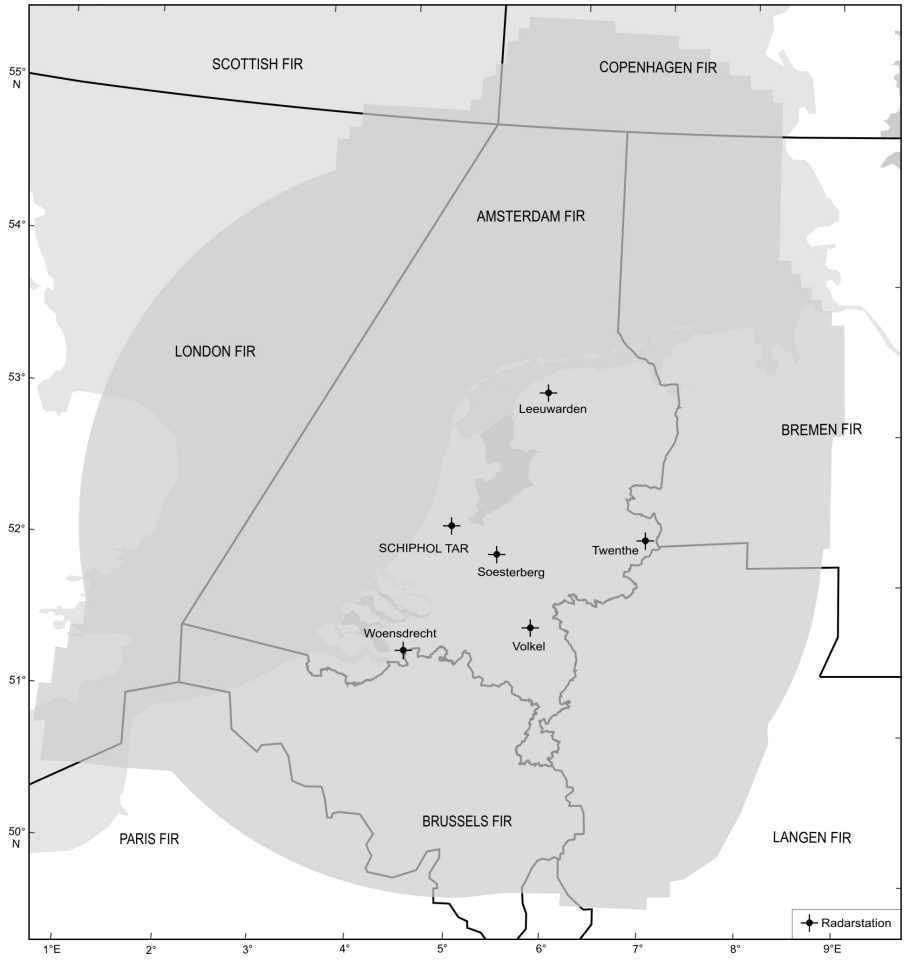

Figure 2.6: Radar range at FL245, approximately 150NM [4]

arrival aircraft, could also help in this limiting factor. “SWIM consists of standards, infrastructure and governance enabling the management of ATM information and its exchange between qualified parties via interoperable services”[18]. The information shared includes route and location of the aircraft, as well as AMAN advisories. This information sharing should allow ANSPs to effectively communicate different speed instructions to other ANSPs such that they can communicate with aircraft to provide earlier instructions to adhere to their EAT.

2.4. Pop-Up flights

Pop-up flights occur when an aircraft is not accounted for in the AMAN planning, and only becomes visible when the planning has already been frozen. This is often due to the pop-up aircraft taking off from a nearby airport just outside the FIR, within the freeze horizon. If this happens, the approach planner is responsible for manually adjusting the sequence of the arriving aircraft to provide an approach and a landing slot to this pop-up aircraft. The slot is assigned based on the insight of the controller, but where possible FCFS. The result is that the controller and pilot workloads increase due to the sequence change and the changed Scheduled Time of Arrival (STA), while the overall fuel use increases due to additional delay [19]. When the AMAN horizon is extended, the amount of pop-up aircraft increases, causing the problems associated with them to become more pronounced. This will be discussed in section 3.3. Figure 2.7 shows a schematic representation in the AMAN process for a normal flight, while Figure 2.8 shows a schematic representation for a pop-up flight. The Freeze horizon stays the same between both flights, though, the departure airport is nearer to the destination.

Currently, 1.78% of scheduled flights are pop-up flights. If the horizon would be extended to 180NM, this would increase to 3.72% if only the freeze horizon is considered. A quite substantial amount of

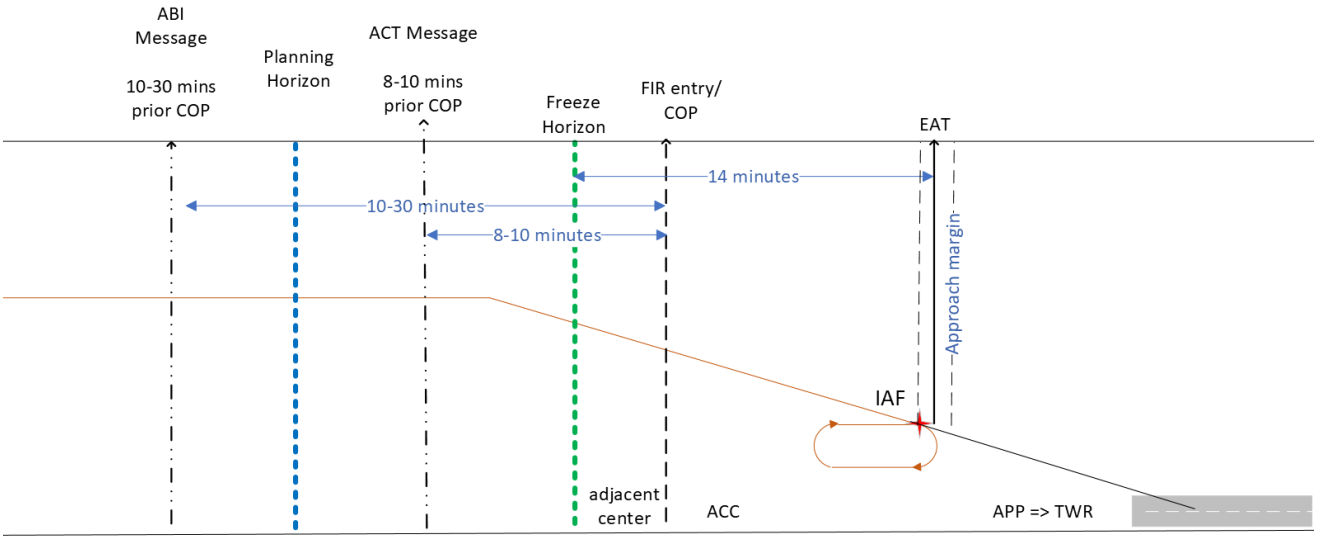

Figure 2.7: Diagram of normal AMAN sequence

airports are barely outside of the 180NM radius, which could give issues with the error in the Standard Instrument Departure (SID) prediction. A closer look is taken at this in chapter 4.

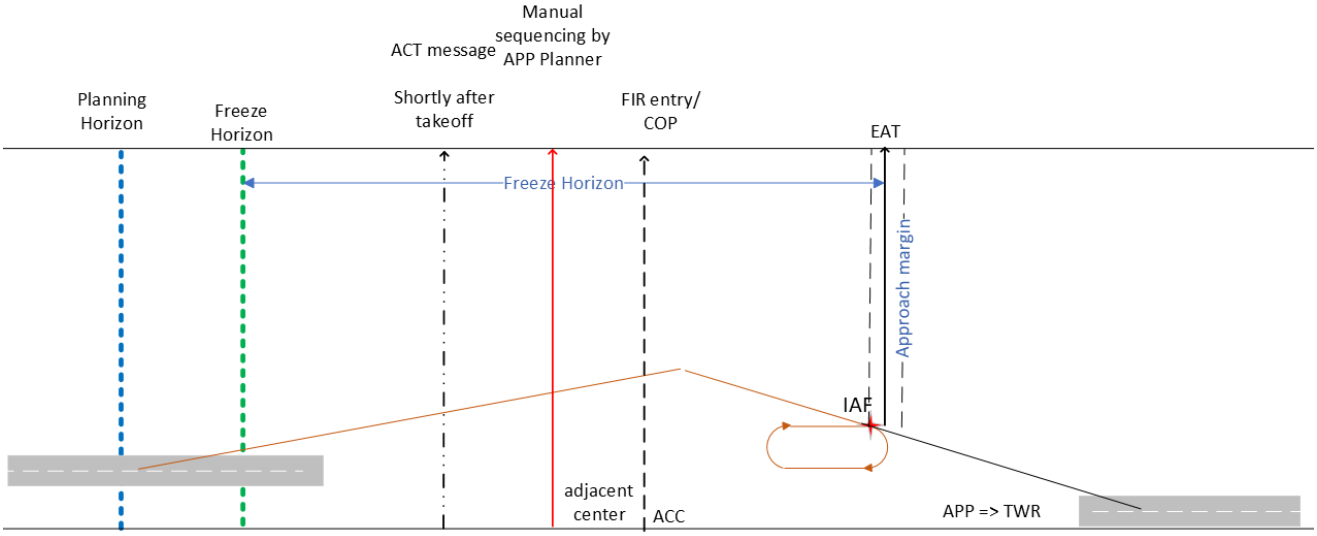

Figure 2.8: Diagram of pop-up flight AMAN sequence

The Pop-up flights are not caused by the aircraft taking off within the freeze horizon, but rather by the inaccuracy in the prediction of the ETA. The inaccuracy of this prediction causes the ANSP to disregard the aircraft in the AMAN planning until after takeoff. A further investigation into this prediction is done in chapter 4.

3

Previous Research

This chapter aims to answer the remaining preliminary research questions, and provide an academic background for further research. First, the previous research into uncertainty relevant to the AMAN is discussed. section 3.2 discusses previously used methods of measuring the systems' performance. A legislative, academic and real world background is then given in section 3.3, after which section 3.4 discusses possible mitigation strategies. The information in the next section is more pragmatic, as it discusses specifics on previous AMAN simulations and how they compare to the actual AMAN in section 3.5. Finally, section 3.6 discusses previous research into the possibilities of absorbing delay en-route.

3.1. Uncertainty

As explained previously in section 2.4, pop-up aircraft are primarily caused by uncertainty in the arrival time at the IAF. The process before the departure of an aircraft has a large influence on this arrival time. Furthermore, the TP accuracy influences the prediction accuracy at the IAF, which in turn is heavily influenced by pilot and controller intent.

Previous research has predominantly focused on the uncertainty of an aircraft at the COP. From the COP, the aircraft becomes under control of the ANSP operating the AMAN. The uncertainty between the COP and IAF is much lower, as the controller intent is under control. The selection of the IAF captures the uncertainty relevant to the AMAN. To be able to quantify the different inaccuracies in a representative manner, the uncertainty in the predictions is to be split up into three parts: the departure planning uncertainty, the departure route uncertainty, and the trajectory prediction uncertainty. This will allow to measure the influence of the different sources of uncertainty on the performance of the (extended) AMAN.

3.1.1. Uncertainty in Arrival Messages

Previous research has examined the uncertainty relevant to the AMAN while studying different stages of the arrival process. The research in [20] has investigated the uncertainty at the COP using Flight Update Message (FUM) messages. FUM messages are only updated if the ETA deviates by more than five minutes from the previous estimation, resulting in a significant level of uncertainty even shortly before arrival at the COP, with an update bias. However, a valid conclusion can still be drawn: estimates on arrival times of airborne aircraft are much more accurate than aircraft not yet airborne with the same lookahead time. This indicates that the uncertainty caused by processes on the ground has a larger effect on the total uncertainty than the trajectory prediction uncertainty. Time of day has a minor influence on prediction error, while flight state, as explained in section 2.2, has a larger influence.

Currently, EFD messages are used to update ANSPs, along with B2B messages provided by EUROCONTROL. These messages are also updated if there is a prediction deviation of more than five minutes from the previous message, as well as every 15 minutes while the aircraft is airborne. The DPI is also included in this datastream, if available. This results in many more updates given per flight

[15]. When compared to the FUM messages, the frequency of these updates has a significant impact on the uncertainty of arrival times when the aircraft is airborne, as there is no deviation threshold for each periodical update. These messages were utilized in the investigation of uncertainty at the COP presented in [15] to select uncertainty features from in order to investigate a probabilistic debuncher. The EFD messages were used in [15], since these were available at LVNL. The analysis showed that the amount of uncertainty is differentiated by three features: prediction horizon, flight state, and CDM state. The most accurate prediction is when the aircraft is airborne, with a short prediction horizon to the COP. Tactical Activated, and Flight Plan Filed flight states both introduce more uncertainty when compared to airborne flights. The uncertainty increases if the prediction lead time is increased in each flight states. CDM equipped airports provide more accurate predictions than standard airports, which is due to the distinction between different CDM state while still on ground. The prediction accuracy data is used to generate probability density curves as a basis for further research. It is noted that the data includes a portion of the flight, up to the COP, and thus includes ATC and pilot intent.

When comparing the prediction error graphs from [20] and [15], it seems that the predictions in [15] seem more accurate, which would be due to the use of EFD messages. Furthermore, the separation of CDM versus non-CDM airports is made, which in case of CDM airports also seem to enable more accurate predictions of COP arrival times.

To conclude, the errors at the COP are substantial, largely due to departure planning uncertainty, while other research can show that it has a major effect on AMAN performance. However, previous research has only investigated the combined inaccuracies of departure planning and the TP up until the COP. The individual effect of the departure planning, and the individual effect of the TP is unknown.

3.1.2. Trajectory Prediction

To be able to quantify the accuracy of a trajectory predictor, a closer look is taken towards the working of a trajectory predictor. First of all, the state of the aircraft must be known. This includes location, time, speed and aircraft characteristics such as mass [20]. The intent of the air traffic control and the aircraft and pilots (often through the autopilot) determine the variations in the state. Finally, atmospheric conditions and especially wind can influence the state of the aircraft. If these variables are not exactly known to the trajectory predictor, an error is propagated.

The current in use TP at LVNL has access to accurate information about the location of an aircraft through SSR. The aircraft characteristics are estimated. The intent of the aircraft is partially known through the flight plan, and updates through the EFD messages. Currently, the main issue with the trajectory predictor of ASAP is the use of standard values. These values are based on historical data, divided into different aircraft types and airlines, and updated every couple of years. The TP uses these standard speeds throughout the prediction, including the cruise phase in another FIR. It is very likely that the aircraft speed is not the same as the standard speed due to pilot or controller intent. Only if a LVNL controller executes a speed instruction, the instructed speed is used in the TP. The TP can be improved significantly if the intent of the aircraft or other controller is known. First, the actual speed of the aircraft should be taken into account, which will likely improve the TP accuracy. The intent of the aircraft during descent, such as Top Of Descent (TOD) and descent speed could also be transmitted to the ground using Automatic Dependent Surveillance-Contract (ADS-C) Extended Projected Profile (EPP). This would greatly reduce the inaccuracy during the cruise and descent phase if this intent is used in the TP [21] [22]. Other variables that influence TP accuracy are wind and aircraft model accuracy.

The research in [15] provides some data on actual TP accuracy between the COP and FIR. The maximum allowable error is 15.9 seconds, which would be quite a low margin of error. In conversation, [23] and [24] elaborated that 1.5% error is considered when using ASAP.

It is known that the departure route performance is quite difficult to estimate. This is due to the intent of the controllers and aircraft, and the difference in performance. The intent has a major influence on the accuracy of the TP in the departure phase, since minor deviations to the route, speed and altitude have larger effects in terms of flighttime than during cruise. Furthermore, the performance of the aircraft during climb is known a lot less accurately than during cruise. This is due to a larger variation in aircraft state (especially mass), and operating procedures. Since quite a lot of airports are very near the E-

AMAN horizon, the departure route performance uncertainty is quite relevant to the E-AMAN planning, as aircraft which are still climbing on the departure route might have to be planned, introducing a larger error. If at least the aircraft intent is known during this phase of flight, the uncertainty can be decreased. It could be possible to use ADS-C EPP data for this, if it is available.

When looking at [15] into the error at the COP when an aircraft is confirmed airborne, with an ATC Activated flightstate and a prediction horizon of 20-25 minutes, an approximation of the error can be found. The boxplot shown shows whiskers at approximately minus and plus 3 minutes, with the 25 % and 75% quartiles lying at approximately minus 1, and plus 1.5 minutes, with a median of zero. This gives some sort of reference numbers for the later investigation into the inaccuracy of the TP due to differences in intent and performance.

In previous research into the effect on the AMAN by the TP has not been extensive. [19] and [25] all reduce the error of this prediction as much as possible. However, the TP module in their research does not eliminate all error. [13] does include some TP error to increase simulation fidelity, by using a median and standard deviation in terms of aircraft speed. This was based on flighttime differences in the DDR database. However, no investigation has been done on the effect of this uncertainty. The result is that there is no information on the effect of the prediction error on the AMAN, while the research into the departure uncertainty effects and solutions has been done with a nonzero TP error.

3.2. Measuring AMAN Performance

Evaluating the performance of a Arrival Manager (AMAN) system requires a structured approach using Key Performance Areas (KPA) and Key Performance Indicator (KPI). KPA represent high-level operational goals such as capacity, efficiency, predictability, environment, and safety, while KPI are the measurable quantities used to assess how well the system meets these goals.

In[26], the implementation of an AMAN system at Changsha Airport in China is analyzed using a structured KPI selection methodology. The study first defines KPAs—capacity, efficiency, environment, and predictability—and then selects KPIs that align with these. The selected KPIs include peak arrival throughput, capacity utilization, controller workload, TMA dwell time, final approach interval, and trajectory coherence. Three scenarios are compared: no AMAN, a newly implemented AMAN, and a matured AMAN. The KPI represent the AMAN performance but also incloudes human operational behavior. While useful for evaluating real-world implementations, these indicators are less suited for simulation-based studies.

Similarly, [27] evaluates AMAN and airspace re-categorization using comparable KPA: capacity, cost, efficiency, and environment. The selected KPI include flight time in the TMA, airport acceptance rate, final approach interval, and taxi-in time. The analysis is again based on operational data and therefore measures combined system and human performance.

In contrast, studies like [19] and [13] are based on simulated implementations of AMAN. In [19] the performance is evaluated using KPIs more specific to system efficiency and controller taskload. These include:

- Average LLDA, number of aircraft requiring LLDA, and delay energy cost
- Number of disturbed descents (where the STA is modified during descent)
- Interarrival spacing for assessing capacity
- Sequence stability measured by position changes
- STA revisions, average STA differences, and affected aircraft for assessing controller workload

LLDA is defined as the portion of delay that cannot be absorbed linearly before the IAF and thus requires instructions such as doglegs or holding. Disturbed descents and frequent STA revisions indicate increased workload and sequence instability. Position changes are counted per slot shift, with a move of two slots counted as two changes. These KPIs are particularly sensitive to pop-up aircraft and planning stability within the freeze horizon.

In [13] similar KPIs are used, adding a breakdown of speed changes by location (pre- or post-TOD) and magnitude (≥1 knot or ≥5 knots).

The research presented in [15] evaluates a debuncher for arriving traffic using metrics such as number of delayed aircraft, total delay, average time from COP to IAF, and changes in AMAN delay. While more specific to the debuncher context, some metrics could support generalized AMAN performance evaluation.

Finally, [28] simulates an E-AMAN-like system and reports only on total delay reduction. However, no workload or sequence stability KPIs are used.

3.3. AMAN Horizon Extension

As stated in subsection 2.3.8, the horizon of the AMAN must be extended to comply with European Union (EU) regulations. This requirement is formalized under Regulation (EU) 2021/116, also known as Common Project One (CP1), which mandates that the AMAN planning horizon should reach at least 180 nautical miles from the arrival airport. This extension aims to improve the efficiency and predictability of arrival sequencing by allowing earlier delay absorption.

Single European Sky ATM Research Joint Undertaking (SESAR) supports this extension through the implementation guidelines [17]. First of all, SWIM shall be used when it becomes available, until then the currently available information sharing systems shall be used. In case that the horizon extension up to 180NM from the arrival airport does not provide benefits due to in-horizon departures, the requirements stipulate that evidence or a justification is submitted to allow a shorter FH to be used. However, this reduced horizon can only be used in selected ACC sectors, not in general. The guidelines conclude that when available, and deemed usable, the DPI shall be used to integrate in-horizon departures into the AMAN. A flow chart as shown in [17] is shown in Figure 3.1.

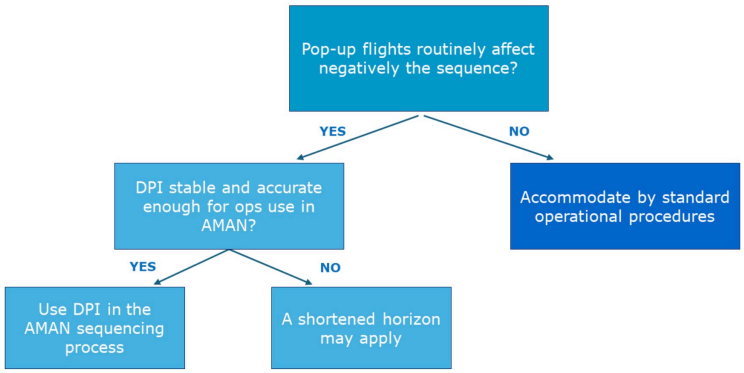

Figure 3.1: In-horizon departures approach flow chart[17]

Previous simulation-based studies confirm both the potential benefits and operational challenges associated with horizon extension. [19] specifically evaluated the extension of the AMAN freeze horizon and analyzed the impact of pop-up flights. Using the BlueSky air traffic simulation environment, the researcher implements a custom AMAN module (AMAN Research Simulator (ARSIM)) and simulates twelve three-hour peak traffic samples in fast time. Two planning horizons are compared: the current 120NM and an extended 180NM horizon, in combination with three scheduling strategies. These include a baseline method, a strategy using FUM messages to anticipate off-block times, and a heuristic method that estimates the probability of pop-up flights based on airport CDM type (CDM, Advanced Tower, or standard).

The results indicate that horizon extension significantly increases the proportion of pop-up traffic—aircraft entering the freeze horizon after the sequence has been frozen, from 1.8% under a standard horizon to 10% under an extended one. While this increases controller workload and destabilizes the arrival sequence, with STA and EAT revisions increasing 27-fold and position changes by a factor of six. However, there is a notable reduction in LLDA and total delay cost by 17% and 22% respectively. Three different scheduling strategies were tested: one using the current methodology, one using FUM messages, and one using a heuristic relation between airport type and pop-up probability. The results

show that if the number of pop-ups is low, or if adherence to scheduled times is good, the extended horizon performs well in case of planning using FUM messages. However, in the presence of high uncertainty, the performance gains are offset by unacceptable increases in taskload. Furthermore, the heuristic scheduling approach did not prove beneficial.

The feasibility of implementing cross-border arrival management to reduce traffic bunching near the IAF is researched in [13]. The research focuses on two operational adjustments: the extension of the AMAN planning horizon and the use of a tighter approach margin, specifically reduced to 30 seconds.

The study evaluates performance using LLDA, sequence stability, the extent of traffic bunching, and the number of disturbed descents. The results show that extending the planning horizon leads to a clear reduction in disturbed descents and a decrease in LLDA, enabling more effective delay absorption at higher altitudes. However, these benefits come at the cost of increased sequence instability, as reflected by a higher number of sequence changes and STA revisions. These are caused by pop-up flights departing from in-horizon airports. The effect of the reduction in disturbed descents and LLDA and increase in sequence changes and STA revisions is more pronounced in the 350NM horizon than the 250NM horizon. The 450NM horizon further increased sequenced changes and STA revisions, without a significant decrease in disturbed descents and LLDA

The impact of a tighter approach margin proved insignificant. While the reduction to 30 seconds did not notably improve sequence stability or delay metrics, it did result in a higher number of tactical instructions, indicating increased workload for controllers.

Both [13] and [19] support further research to reduce the impact of the horizon extension on sequence stability, as no satisfactory mitigation strategy to pop-up flights has been found.

Similarly, [28] shows that the departure time uncertainty on an extended arrival manager has a detrimental effect on the performance. This research has been conducted around Philadelphia, using simulated air traffic including departure uncertainty. The horizon used was set at 400NM, with two schedulers being evaluated. The first is scheduling aircraft using their predicted departure time, while the second is to only schedule pop-up aircraft once they become airborne. Similar to previously discussed research, the conclusion is that preplanning pop-up flights is detrimental to the performance of the E-AMAN.

A human in the loop simulation is presented in [29] to investigate the operational feasibility of an En Route Arrival Manager (En Route AMAN) at Tokio Haneda. This system aims to support upstream controllers by issuing speed control and runway assignment instructions far before aircraft enter the TMA, essentially being an E-AMAN. An approximate horizon of 180NM was used in this research.

The system was evaluated using the ESCAPE Light simulator, under realistic traffic and weather conditions, including north-wind configurations at Tokyo Haneda (RJTT). Automatic runway assignment, and early speed control were evaluated. The early speed control in this case consisted of a fixed deceleration (-0.03 Mach) at an entry point.

The results show that the proposed system could lead to a reduction in average flight duration of 25.6 seconds in high-traffic scenarios. Furthermore, the number of instructions issued by controllers decreased by 20%. Questionnaires confirmed that the workload had decreased when compared to the regular horizon. This study confirms that extending the AMAN horizon does have positive impact on the performance when humans are using the system. It is noted that this study did not include pop-up traffic.

London Heathrow currently employs an extended-horizon, cross-border arrival management Cross-Border Arrival Manager (XMAN). The system is designed to absorb arrival delay while aircraft are still en-route, minimizing holding and reducing fuel burn [30].

When predicted holding delays at Heathrow exceed seven minutes, upstream air traffic control centers, such as MUAC are requested to issue a speed reduction, usually between 0.02 and 0.04 Mach reduction. The instructions occur at 350NM from London Heathrow [31]. The delay absorbed is transferred from the holding stacks, to the cruise phase. Early operational results show that significant gains are made by the system, reducing the fuel burn by approximately 3700 tonnes annually [31].

It must be noted that the procedures at London Heathrow are much more dependent on holding stacks to absorb delay. As much as 36% of flights join the stack [32], with aircraft spending quite a large

amount of time in the stack when they do, approximately 6.5 minutes [33]. This makes the procedure easier to implement, versus the approach LVNL uses at Schiphol. The stacks are used as little as possible, delay is proactively reduced, and two landing runways are available at inbound peaks. This causes the delay per aircraft to stay much lower, thus not providing the early knowledge of the delay when inbound. Furthermore, not all aircraft are part of this XMAN, only the aircraft that are airborne before the 350NM horizon. Due to the 7 minute threshold, it is certain that the concerned flight needs to absorb delay.

3.4. Mitigation Strategies for Pop-Up Flights

Currently, pop-up flights are accommodated by the operator of the AMAN, who manually inserts the flight into the sequence. Often, this causes a STA revision for multiple flights. If the horizon increases, the amount of pop-up flights also increases, thus requiring a mitigation strategy.

Scheduling pop-up traffic using their predicted departure times has previously been investigated by [19], [13] and [28]. In presence of uncertainty in that prediction, the results from all three show that this strategy to mitigate the negative effects of pop-up traffic was not effective.

The research by [25] researches a dynamic Inbound Priority Sequencer (IPS) model to optimize inbound flight sequencing from a network airline's perspective. The research does not focus on pop-ups, though the solution it provides could apply as a mitigation strategy.

To mitigate the disruption caused by such late-appearing flights, the IPS model incorporates a mechanism known as the Calculated Take-Off Time (CTOT). For selected departure airports the CTOT allows the airline to assign a preferred take-off time to the flight, enabling it to be integrated into the arrival sequence before takeoff. The research notes that the actual takeoff time must be within two minutes of the used CTOT to be able to fit in the sequence accurately.

The dynamic IPS model recalculates the arrival sequence at fixed intervals (every 30 seconds), making it robust to changes in flight status, including updates to departure times and new entries into the action horizon. This rolling update mechanism contrasts with traditional AMAN systems, where the sequence is often frozen well before such flights become visible. Results from the case study at Schiphol Airport show that the IPS model reduced overall fuel cost, missed connections, and holding patterns compared to a baseline FCFS approach.

A probabilistic debunching concept is proposed in [15], aiming to reduce the effect of the uncertainty on the performance of an E-AMAN. The debuncher would operate on available EFD messages, determining a probability of a traffic bunch arriving at the IAF. A traffic bunch is defined as at least two aircraft arriving at the IAF within each others' separation minimum. The debuncher then would instruct aircraft based on this bunching probability, to absorb delay early. The effects on the AMAN were generally slightly positive, though it would not provide enough improvement to be able to mitigate the effects of pop-ups to a satisfactory extent.

Previous research at Stockholm using an Required Time At (RTA) to determine a takeoff window for pop-up flights from regional nearby airports has proven to be beneficial [**rydell_arrival_nodate**]. The actual flights could absorb delay at the gate, while calculating a takeoff window to be able to arrive at the IAF on the RTA. However, a significant increase in communication between the airports was required to accommodate this. The thesis by [**rydell_arrival_nodate**] further investigates the implementation of RTA in an extended AMAN context. The RTA is assigned when crossing the freeze horizon, or prior to takeoff for pop-ups. In simulations using historical traffic data, RTA assignments were shown to reduce the deviation from scheduled sequences and limit the amount of instructions. Pop-ups could still be sequenced effectively by assigning a feasible RTA prior to departure. An Flight Management System (FMS) calculated departure window could make sure that major airports can accommodate the flight departing at a time in which it can adhere to the RTA. The main advantage of the RTA mitigation strategy is the possibility to absorb delay on-ground, minimizing fuel burn. However, it must be noted that it would require quite some additional communication between ANSPs, unless a complete AMAN-Departure Manager (DMAN) integration is implemented. Furthermore, aircraft behavior would change according to FMS instructions, which would reduce the predictability which is required for controllers to maintain separation. Similarly, [19] has also recommended research into a Target Time of Arrival

(TTA), which is very similar in theory to the RTA.

3.5. AMAN simulations

As noted in previous sections, research into the AMAN where simulations have been made have been conducted multiple times, most often in BlueSky [1]. Each research evaluated different aspects of the AMAN, or performed research where a functional AMAN was required for research into other areas. The commonality of the presented research is that they all feature simplifications of the AMAN in their simulation, and there are recommendations made to improve the simulation. This section provides an overview per relevant research of the simplifications and their possible implications, and the recommendations made relevant to this research.

3.5.1. Trajectory Prediction

The research in [19] and [25] minimized the Trajectory Prediction error in their simulations. The TP is based on the TP used in ASAP [19]. When compared to the aircraft performance in the simulation used, there may be quite a significant error. [19] notes that there might be an effect on the E-AMAN in combination with pop-up flights. [25] uses a mean speed offset and deviation to include some error in the TP, although only outside of the Dutch FIR. This method is likely employed because of its' simplicity, and mimics controller actions, wind and other uncertainties. The result of this uncertainty is not investigated individually, only as part of the total research.

As noted in subsection 3.1.2 the approximate error used by controllers is approximately 1.5%, while [15] has calculated that the worst case TP error in the Dutch FIR was 15.9 seconds. It can be noted that the previous research has focused mainly on en-route TP inaccuracy, while the climb phase might introduce the main errors in the predictions. To conclude, the effects of the inaccuracies in the TP module of the (extended) AMAN are unknown, but assumed to be much smaller than pop-up aircraft.

Another major TP related simplification is the assumption of zero wind. The effect on the TP is quite large, but constant and similar for all aircraft, and the accuracy of the wind model is of larger influence than the magnitude of the wind. Therefore, impact on AMAN performance is deemed likely to be small by [19], [13] and [25], although it does have an effect on both the TP accuracy and the traffic scenario. The TP accuracy is reduced if wind is present, since the TP does not include a perfect wind model. Furthermore, the aircraft behavior especially with respect to ground speed and tracks is altered dependent on wind.

3.5.2. Traffic Scenario

The traffic scenarios in previous AMAN simulations focus on inbound peak traffic at Amsterdam Airport Schiphol. The scenarios in [19] are comprised of 12 three-hour periods based on the actual traffic of 6 days in the summer of 2015. The off-peak periods are not implemented in the scenarios. The scenarios in [13] are essentially the same, while [25] has only used smaller batches of westerly arrivals. This is done since the presented research is much more computationally intensive, and requires much more preparation. [19] notes that investigation into off-peak periods could be relevant for the performance of the E-AMAN. Since the AMAN is most beneficial when the arrival flow is very near or above landing capacity, the off-peak peak periods saturating a single runway are most interesting to the performance of the AMAN.

Another area of simplification has been the aircraft types: [19] and [13] both only use Boeing 747-400 in their scenario's, with the associated aircraft performance. This is also the case for pop-up flights. The effect is likely not very large, although it reduces the fidelity of the simulation. [25] has used different aircraft types, though no investigation is done into the effects of these aircraft types. All of the discussed simulations used a fixed LIV of 100 seconds, which would be realistic with the single aircraft type. Currently, a dynamic LIV is used, as described in section 2.3.

One of the more extensive simplifications is the omission of departing traffic from Schiphol. All three previously cited sources did not take departures into account in their scenarios. Including the departing traffic would result in a much more complicated scenario, while the performance of the AMAN would likely be nearly identical. This would be because of the arrival and departure routes being separated from each other outside of the TMA without requiring ATC instructions. The departure routes do limit

the options of delay absorption in ACC airspace when vectoring, but this can be accounted for in simulations.

In contrary to the real-life situation as indicated in chapter 2, the aircraft are not vectored in the TMA to ensure separation at the runway in the simulations done in [13] and [19]. The reasoning behind this simplification is that the implementation would be time extensive, and no impact on the AMAN would occur. Both have used simple fixed routes to the runway. [25] went as far as not simulating the TMA portion of the scenario. If the TMA would be simulated, it is likely that the effect would be noted in the numerical results, as any changes are diluted by the enlarged sample. Simulating and measuring the TMA operations might not be useful in terms of obtaining numerical results. Additionally, as noted by [19], the fixed routes towards the runway are often already conflict free, if the aircraft pass the IAF within a 30-second approach margin. The implementation of these fixed approach routes in the real life airspace is also a future goal of LVNL in the airspace redesign program.

Finally, in [19] the uncertainty in takeoff times of pop-up aircraft was modeled as a single, set delay, which would be varied to study the impact of this delay. Since this delay is not based on a statistical model, studying the effects of the variable is simplified, while reducing similarity to the real life scenario of a stochastic process. The effects of this simplification on AMAN performance are unknown, although speculation would suggest that including negative delay might have a slightly less negative effect on performance in similar circumstances.

3.5.3. Simulation

A few relatively minor simplifications have been made to shorten development time or reduce computational cost of previous simulations. For example in [19], the simulated aircraft only enter the simulation at a radius of 500NM, to reduce the computational cost. Furthermore, the simulation frequency is reduced from 20Hz to 5Hz, with only very minor effect on the simulation and thus on the AMAN performance.

3.5.4. ATC Instructions

ATC instructions are quite a complicated process to model in a simulation. The solutions implemented on a daily basis by controllers are creative and difficult to model. However, the solutions mainly focus on ensuring separation between aircraft, both in the same phase of flight and between different streams of traffic. One of the simplifications made in [19], [25] and [13] is to not deconflict aircraft. This is done due to the implementation being difficult if an accurate representation of controller instructions should be included. Furthermore, the effect of the controller instructions on AMAN performance is likely to be small. The simulations do provide separation at the IAF if the approach margin is set to 30 seconds, and would provide for separation at the runway.

Other simplifications are made in the solutions in the AMAN process. The simulations previously discussed, all do not allow so-called frontloading or ad-hoc runway changes (scharrelen). The result of these simplifications is that the AMAN in all simulations performs worse overall, in AMAN and E-AMAN scenarios with and without pop-ups. However, there might be different effects according to the different scenarios. Speculating would suggest that in the case of many pop-ups, with these actions available to the controller, there could be more opportunities to efficiently include pop-up traffic in the arrival sequence.

Finally, [19] assumed that a speed reduction of 10% should be possible for all aircraft in cruise and descent, in terms of planned speed. This assumption slightly differs from the conclusion in [34], which used speeds based on aircraft performance and their maximum range cruise speed. The latter would provide for a more ideal solution especially in upper airspace, although as noted before, a larger speed reduction might be more beneficial in a scenario in which an E-AMAN is operational.

3.5.5. Recommendations

[19] recommends to further investigate E-AMAN performance during off-peak periods, as well as to increase the actions available to the simulated controller, such as frontloading, expedite margin, and the ad-hoc changing of runways. These recommendations are also made by [13], together with the suggestion of investigating a delay threshold to reduce the amount of speed changes made by aircraft, and applying discrete speed instructions instead of continuous. The recommendations of the delay

threshold and discrete speed instructions are essentially used in the current AMAN operation. It will likely be worthwile to investigate the implementation into an E-AMAN, with the goal of reducing the amount of simulated controller actions required.

It is also recommended to try to reduce the amount of radio transmissions to be made by UAC, for which the previous recommendations might provide a significant reduction. [19] also recommends the simulation to be adapted to other airports, since the performance might be very different at different hubs. Finally, [19] recommends to research the assignment of Target Time Over (TTO) to pop-up aircraft to solve the performance degradation associated with these aircraft. In this concept, pop-up aircraft are planned before their takeoff and assigned a slot in the sequence, but their slot time is not revised if they encounter a positive or negative delay. This solution could significantly improve the performance of the E-AMAN, though a closer look must be taken at aircraft that cannot make their assigned slot due to being delayed too much.

Furthermore, [19] also recommends including TP errors in future research, as well as investigating the effects of these errors.

3.6. Delay Absorption Possibilities

[34] has investigated the possibility of delay absorption in upper airspace, governed by MUAC, on aircraft inbound to Schiphol. The experimental part of the research has been done in the BlueSky simulator. This research approached the arrival management problem from the upper airspace, without considering the requirements on the rest of the airspace to be able to implement the management in upper airspace. However, it provides insight into the delay possibilities in upper airspace, and is thus very relevant to extending the AMAN horizon. The options proposed in this research are to decrease the speed of the aircraft, to decrease the altitude by 2000 feet, to add track miles by adding a detour, or a combination of all three. These are named linear delay, dropping, detouring, and turtling respectively. The reduction of altitude in the research is set at 2000 ft, while the airspeed reduction is limited by the maximum cruise range mach of the specific aircraft. The relationships obtained to calculate the average maximum delay per method are the following: Linear delay amounts to 0.4 seconds per NM, dropping to 0.7 seconds per NM, detouring to 8 seconds per additional NM plus 0.4 seconds per NM, and finally turtling amounts to 8 seconds per additional NM plus 0.7 seconds per NM. The current estimation of possible delay absorption used by controllers is 1 minute per 100 NM, or 0.6 seconds per NM.

If the FH is set at 200 NM, it can be estimated very roughly that aircraft could be delayed approximately 80 seconds if linear delay is used, 140 seconds if dropping is applied and 120 seconds if the rule of thumb of the controllers is used. Turtling or detouring could add more delay, but would increase fuel consumption.

It can be noted that the maximum range speed is used in this research. This number might be a more realistic to achieve for aircraft, but the actual realistic limit might be lower. The maximum endurance speed is nearly always lower than the maximum range speed. If the maximum endurance speed would be used such that the aircraft does not require holding or a longer path, the fuel burn could decrease.

The research also evaluates the impact of implementing these delays in the MUAC airspace in terms of workload and fuel efficiency. There is some reduction in fuel burn in MUAC airspace shown in the presented research for the dropping and turtling delay absorption methods. This is caused by part of the descent taking place at an earlier stage. The linear holding causes no change in fuel use, while detouring causes a slight increase. The author notes that: ”en-route delay is not inherently more fuel efficient than delay absorption measures taken in the lower airspace” [34]. However, it can be noted that the presented research only investigated the results in the upper airspace. If delay absorption in the upper airspace is implemented, it could result in fewer holding patterns being flown which would reduce total fuel burn. It shows that further research into the possible reduction in fuel burn if aircraft are delayed en-route versus lower level delay is required.

The controller workload increases slightly overall, usually remaining within 3% of the baseline workload except for low season inbound peaks. The workload could then be up to 5-10% higher for a controller. These figures show that the workload of controllers could increase overall, or shift towards different air traffic control centers, making it relevant to research further when en-route delay is applied in an AMAN

application.

In conversation with ACC, it was noted that for aircraft from Brussels a delay of six minutes could be given to an aircraft using speed and vector instructions. A negative delay of approximately 2 minutes could also be ensured through giving aircraft directer clearances. Note that aircraft from Brussels have the shortest flightpath in ACC airspace before arriving at the IAF, thus larger delays could be possible from other directions.

4

Preliminary Findings

4.1. Goal and Approach of Preliminary Research

The goal of this chapter is to answer the preliminary questions regarding uncertainty, to be able to accurately represent the different time uncertainties in an E-AMAN simulator, and to present the possible solutions to the pop-up aircraft. A pop-up aircraft is an inbound flight that becomes visible to the Arrival Manager after the freeze horizon has passed. Previous research has investigated FUM[19][20] messages, which show quite a large prediction error at FIR boundary, with similar research in [15] showing that EFD data has smaller prediction errors. Though, this research has focused on the error at the FIR boundary or COP, while the main error is observed in the pre-departure phase. Furthermore, the error introduced by a deviation from the planned departure route is unknown, though might be relevant for the AMAN performance due to near-horizon departures. Finally, there is an aspect of TP uncertainty, especially due to unknown and uncontrollable pilot or ATC intent outside the FIR. To be able to accurately quantify the error introduced in these three flight phases, the EFD data is investigated in the following sections. Due to the three distinctly different flight phases, the analysis of prediction errors is split up in a pre-departure phase, a departure route phase, and an en-route phase. The en-route phase is again split up in a outside-FIR and inside-FIR phase. EFD is relevant as it is already used in ASAP and other LVNL systems, and thus well known within the organization. Furthermore, it is available in the currently used systems, which makes using it in an E-AMAN context easier.

In section 4.3 it is shown that many aircraft depart near the horizon of the E-AMAN. This could implicate that there could be airports at which aircraft take off moments before the freeze horizon. This implicates that the uncertainty in the departure route could influence the performance of the E-AMAN. Splitting the uncertainty at the FIR boundary into multiple aspects allows for this uncertainty to be represented accurately in the simulation. Equation 4.1 shows the simple formula for the total error at the FIR boundary, while Figure 4.1 shows a schematic representation of the prediction errors.

$$E_{\text{IAF}} = E_{\text{TP in FIR}} + E_{\text{TP outside FIR}} + E_{\text{Departure Route}} + E_{\text{Takeoff}} \tag{4.1}$$

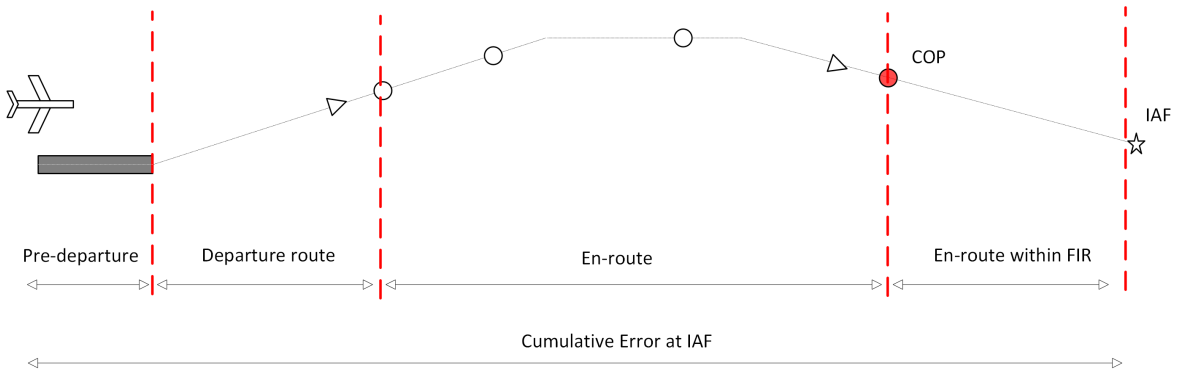

Figure 4.1: Schematic representation of prediction error in different phases of flight

First, section 4.2 gives an introduction about the data selection, processing and filtering. section 4.3 dives deeper into the pop-up occurrence when the horizon is varied, as well as the CDM type of the respective departure airport. Section 4.5 elaborates upon the uncertainties in departure planning to be able to quantify this accurately in further simulations. Adding to this, section 4.6 elaborates on a similar investigation into the departure route, while section 4.7 sets out the analysis of the uncertainty in trajectory prediction of the en-route phase. Finally, section 4.8 and section 4.9 set out the options and space that controllers have to absorb delay, and which mitigation strategies to reduce the impact of uncertainty and resulting prediction errors on the E-AMAN performance are to be considered.

4.2. Data Selection, Processing and Filtering

The following sections are all based on an analysis of EFD data provided by LVNL. The data spans from 2023-02-27 up to 2024-04-26. A pipeline was set up to process and compress the data, since the raw data is more than 500 gigabytes and takes a long time to analyze. It also contains a lot of data of flights not relevant to the arrival manager, such as flights simply overflying the Dutch FIR. As discussed before in section 2.2, EFD data is divided into messages, and multiple messages are given per flight. An example message can be found in **??**. The EFD data is chosen since it is available at LVNL, both for this research, but also in operational context. It also contains all of the data required to conduct the uncertainty analysis of the takeoff times, departure route time error, and trajectory prediction time error. Other available data would have less availability or would not be as accurate as EFD data.

In the pipeline, each message is read, and if the aircraft is inbound to Schiphol, the message is added to the respective flight. After processing a month of data, these flights are stored as an object on a hard drive. After processing the entire dataset, there is a flight object for each flight inbound to Schiphol airport. These can then be processed further for analysis. The flights have multiple attributes relevant for a statistical analysis, such as actual takeoff time, or time at Cross Border Airspace (CBAS). Each flight is processed again to calculate prediction errors, lookahead time, and all other values, for each individual message, and filtered again to only store relevant messages. These individual messages are then stored in a dataframe for easy plotting. The errors in the expected departure route, and the time to CBAS are also calculated, and stored in a separate dataframe. These dataframes are then all stored in pickle format for easy access. Below, a schematic representation is given of the process.

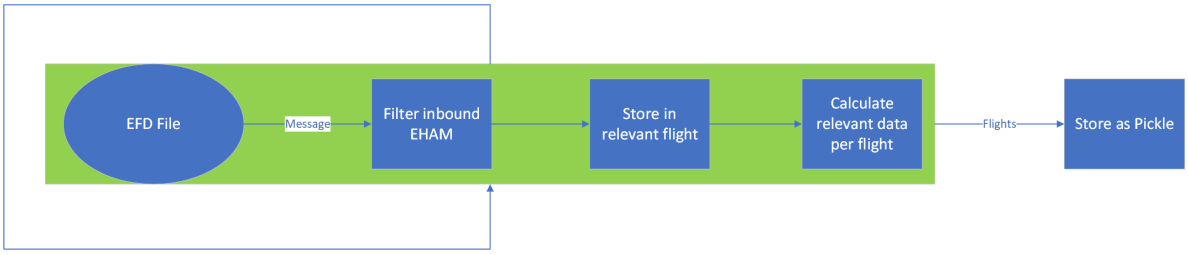

Figure 4.2: EFD processing pipeline

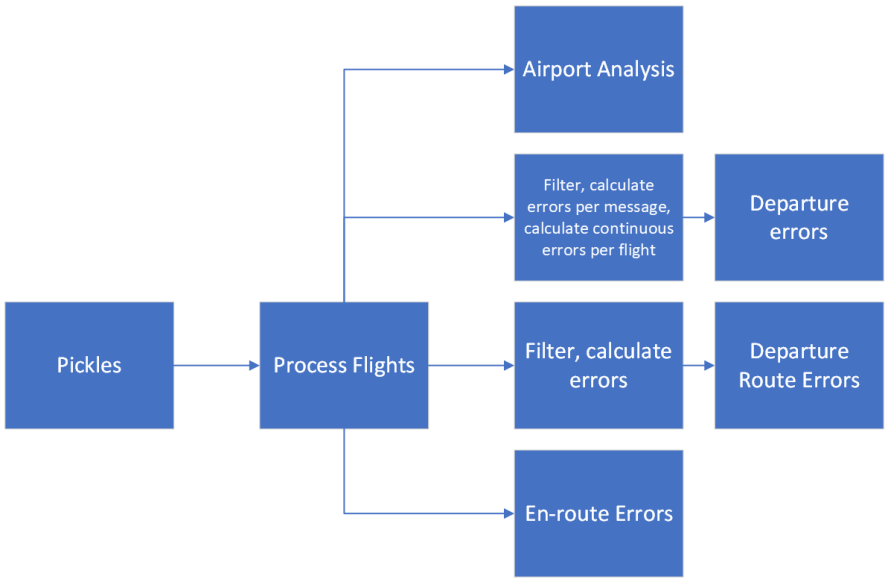

Figure 4.3: Flight data processing pipeline

A few filters are used on the data to ensure that the relevant data is used. A short overview is given below, divided per analysis. The results of this data analysis are shown in the coming sections.

1. All flights considered are inbound to Schiphol
2. Filters for the departure planning time error analysis:
   - 2.1 The planned flight duration is less than 2 hours
   - 2.2 There is an Actual Takeoff Time (ATOT) specified in a CPR message
   - 2.3 The flight has not been canceled
   - 2.4 The Estimated Takeoff Time (ETOT) lies within 1 hour of the message
3. Filters for the departure trajectory prediction time error analysis:
   - 3.1 The planned flight duration is less than 2 hours
   - 3.2 There is one common point above FL180 and below FL300 in both the planned departure, and flown departure route
   - 3.3 There is one common point above FL100 and below FL300 in both the planned departure, and flown departure route for flights from Brussels
4. Filters for the cruise trajectory prediction time error analysis:
   - 4.1 There is a termination message for the flight
   - 4.2 The planned route and flown route both cross a COP

Since this research is focused on the AMAN at Schiphol, the choice is made to only include flights inbound to Schiphol. Furthermore, pop-up flights depart near Schiphol, thus a choice is made to only include aircraft with a planned flighttime of less than two hours for the analysis of the takeoff time and departure route errors analysis. The other filters ensure that the relevant time error can be calculated. The common point is chosen above FL180 and FL100 for Brussels, since these altitudes usually contain the majority of the departure route. If a higher altitude were to be chosen for the Brussels departures, there usually would not be a match between the planned and flown route.

4.3. Pop-up Occurrence with varying Freeze Horizon

To give an indication of the impact of increasing the freeze horizon of the arrival manager, this section shows the amount of impacted flights if the freeze horizon would be extended from approximately

120NM, to 180NM. It is noted that this analysis is done in distance, while the AMAN works on a time based system. Though, the goal of this section is to show the relevance of developing mitigation strategies for pop-up flights in an E-AMAN environment.

Figure 4.4 shows a map with the current 120NM radius, the proposed extended 180NM radius and a radius of 210NM plotted. The latter is for reference to which airports could be considered near-horizon airports. Furthermore, the relevant airports are plotted, together with the CDM type.

Table 4.1 summarizes the major airports that currently lie within the 120NM freeze horizon and contribute to pop-up traffic. In this table, it is clear that Dusseldorf and Brussels are currently the main pop-up causing airports, with multiple flights daily, which results in a total pop-up percentage of 1.78%. Note that Dutch airports, especially Rotterdam, coordinate departures with Schiphol approach, so they are not considered pop-ups.

Table 4.1: Departures from current in-horizon airports (120NM)

Airport	Airport Name	# of Flights	% of Flights	Distance[NM]	Airport Type
EDDL	Dusseldorf	1851	0.72	96	CDM
EBBR	Brussels	1781	0.69	85	CDM
EHRD	Rotterdam	219	0.08	24	STANDARD
EHEH	Eindhoven	202	0.08	56	STANDARD
EHGG	Eelde	115	0.04	82	STANDARD
EBAW	Antwerp	77	0.03	68	STANDARD
EHBK	Maastricht	62	0.02	92	STANDARD
EBLG	Liege	50	0.02	104	ADVANCEDATCTWR
Other	21 airports	267	0.10		STANDARD
Total		4624	1.78		

When the horizon is extended to 180NM, additional airports begin contributing pop-up flights. Table 4.2 shows the inbound flights from the airports that lie between the current and extended horizon, which would all be considered pop-up aircraft if the horizon were to be extended. It is clear that the amount of pop-up flights doubles, which greatly increases the negative impact on the performance of the E-AMAN as shown by [19]. The total amount of pop-ups would more than double, up to 3.72% of all flights inbound to Schiphol.

Table 4.2: Additional departures from extended AMAN in-horizon airports (120-180NM)

Airport	Airport Name	# of Flights	% of Flights	Distance[NM]	Airport Type
ELLX	Luxembourg	1350	0.52	170	STANDARD
EGSH	Norwich	1238	0.48	129	STANDARD
EDDW	Bremen	1161	0.45	153	ADVANCEDATCTWR
EGSS	Stansted	832	0.32	169	ADVANCEDATCTWR
EGMC	Southend	161	0.06	157	STANDARD
EDDK	Cologne	82	0.03	124	STANDARD
Other	22 airports	136	0.05		STANDARD
Total additional		4960	1.91		
Total pop-ups		9584	3.72		

Table 4.3 lays out the departures within 30 miles from the extended horizon. This table shows that there are many flights taking off just outside of the freeze horizon which are inbound to Schiphol. These flights could influence the E-AMAN performance in two ways: the TP is less accurate in the departure route phase, and thus increases the uncertainty. Futhermore, if the horizon would be extended further than the proposed 180NM, the amount of pop-ups would drastically increase.

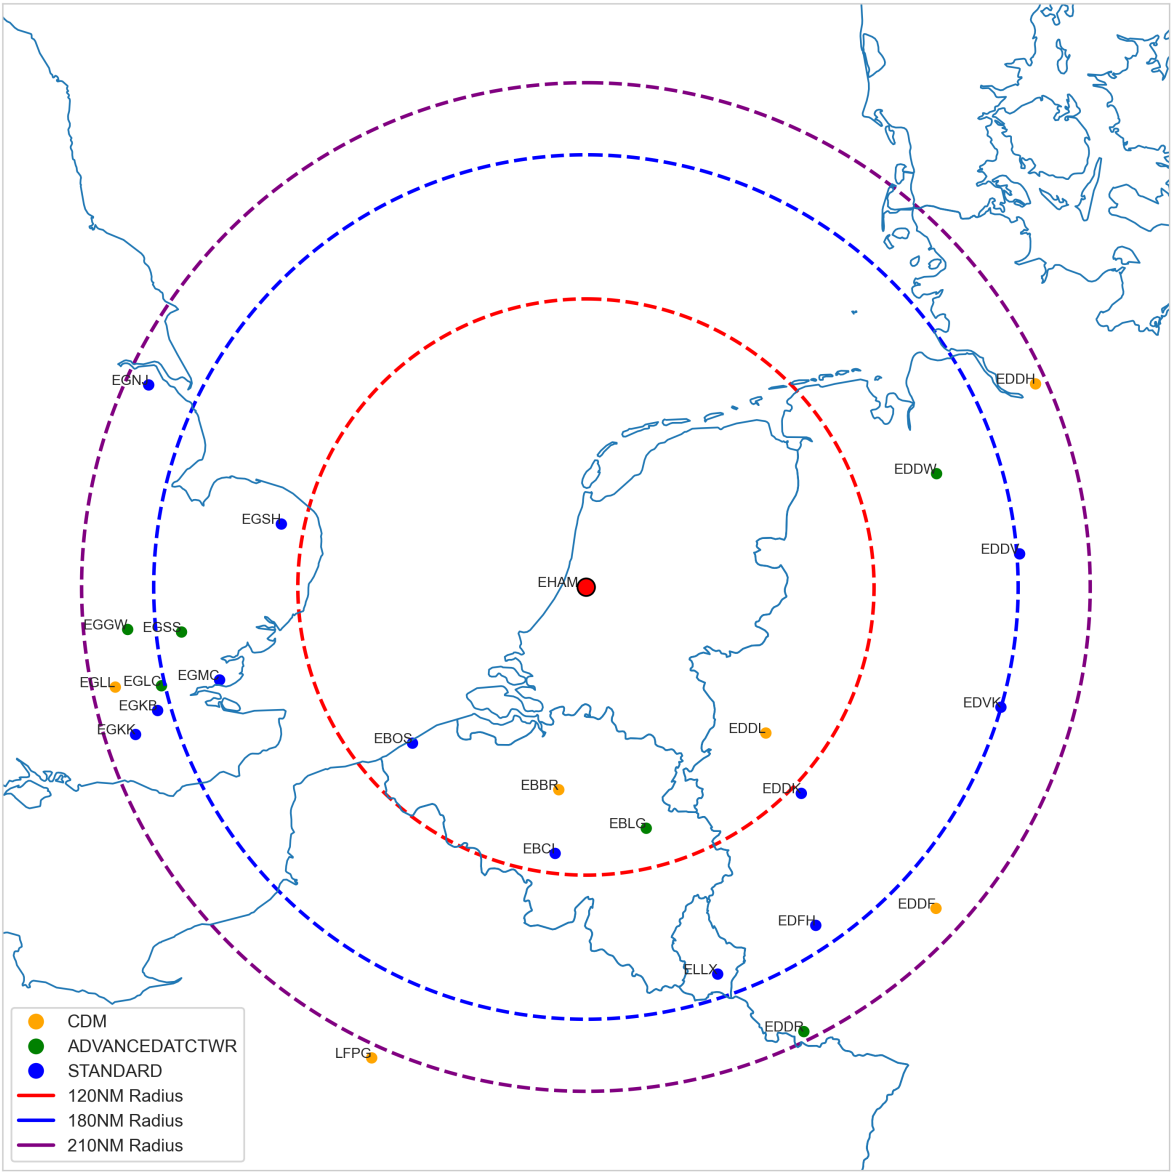

Figure 4.4: Map of the current (120NM), extended (180NM), and near-horizon (210NM) horizons from Schiphol with pop-up contributing airports

Table 4.3: Departures from near-extended horizon airports (180-210NM)

Airport	Airport Name	# of Flights	% of Flights	Distance[NM]	Airport Type
EGLL	Heathrow	6109	2.37	200	CDM
EGLC	London City	4195	1.63	181	ADVANCEDATCTWR
EDDF	Frankfurt	4102	1.59	197	CDM
EGKK	Luton	3224	1.25	197	STANDARD
EGGW	Gatwick	2506	0.97	191	ADVANCEDATCTWR
EDDH	Hamburg	2182	0.85	205	CDM
EDDV	Hannover	1387	0.54	181	STANDARD
EGNJ	Humberside	950	0.37	200	STANDARD
EGKB	London Biggin Hill	255	0.10	185	STANDARD
EGWU	RAF Northolt	106	0.04	197	STANDARD
Other	12 airports	77	0.03		STANDARD
Total additional		25093	9.74		
Total pop-ups		34.677	13.45		

An overview of the percentage and CDM composition of inbound pop-ups is given in Table 4.4. This table shows that the amount of CDM and AAT airports is quite significant, and that the amount of pop-up flights greatly increases if the horizon is extended to 210 NM. The composition with respect to the CDM and AAT airport types is of interest, due to the DPI messages given which theoretically increase the departure planning accuracy. A closer look is taken in section 4.5.

Table 4.4: Flights within distance ranges: share of total and share of pop-up flights per airport type

Airport Type	0–120NM	Normalized	0–180NM	Normalized	0–210NM	Normalized
Total %	1.79%	100%	3.72%	100%	13.45%	100%
CDM	1.41%	79%	1.41%	38%	6.21%	46%
ADVANCEDATCTWR	0.03%	2%	0.81%	22%	3.41%	25%
STANDARD	0.35%	20%	1.50%	40%	3.83%	28%

Finally, a graphic representation of the share of pop-up flights is given as a function of the FH in Figure 4.5. The figure shows the major increase after 180NM, due to some major European airports being at that distance.

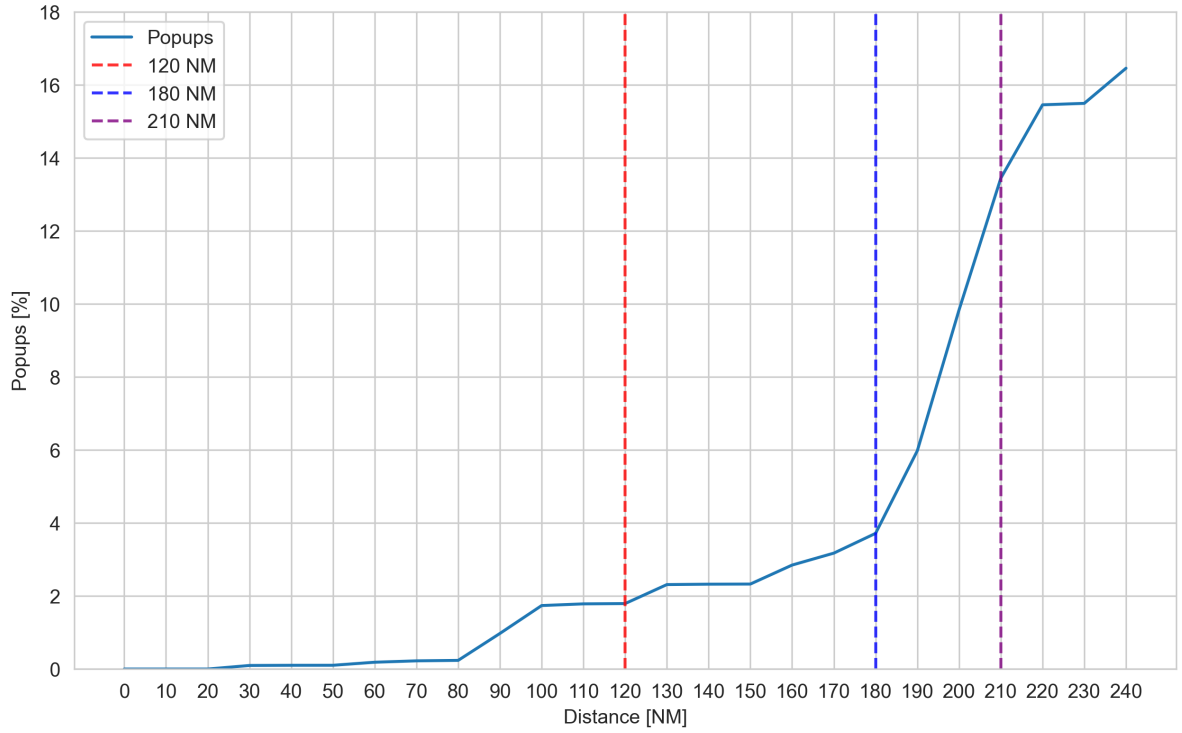

Figure 4.5: Pop-ups as a function of Freeze Horizon

4.4. Feature Selection

To realistically represent prediction uncertainty in an AMAN context, the relevant uncertainty sources need to be represented using a probability density function, based on distinct features. The goal of this section is to select relevant features that influence the statistics of the prediction errors of the different flight phases.

Previous research has evaluated the TP accuracy at the COP [15] [20]. These studies identified several features of interest in EFD and FUM data. These were prediction lookahead time, CDM airport classification, flight state, CDM state and departure airport.

The flightstate feature is of major influence on the accuracy of the prediction. It also represents the flight phase, but not directly. Since this research aims to represent the uncertainty as accurately as possible, the decision has been made to split up the uncertainty analysis in the different flight phases, but not exactly along the flightstate feature. Flightstate does not have a value for the departure route, and some inaccuracies occur especially around the flightstate of ”Tactically Activated” [15].

The prediction lookahead time, defined as the time between the prediction timestamp and the predicted event, is of major influence on the accuracy of the predictions. The lookahead is especially important in the takeoff uncertainty, as events with major influence on the takeoff planning can occur much more easily, such as pushback delays. A horizon of 15 minutes is chosen, as this would capture the uncertainty at the estimated FH for every airport with departing pop-up aircraft. The resolution is chosen to be 1 minute. The lookahead time is not used as a feature for modeling the departure route uncertainty. The deviation from the filed departure route is mainly driven by ATC procedures and instructions, not the time at which the prediction is made. For en-route predictions, however, the lookahead time is again a relevant feature, as the lookahead time is directly influencing the amount of flighttime that needs to be predicted.

Another feature of interest is the CDM classification of the departure airport. In theory, the departure predictions of CDM enabled airports should be more accurate. This is due to the larger amount of messages sent with a theoretically more accurate prediction for the takeoff prediction. There might also be an effect on the departure route.

The departure airport itself might also have a major influence on the takeoff and departure route prediction. The prediction accuracy greatly depends on the procedures being followed at each airport, which might be modeled differently than the actual procedures or behavior. Furthermore, the geographic location and traffic density might also be of influence.

A summary of the different features selected is shown in Table 4.5

Table 4.5: Selected features per flight phase for uncertainty modeling

Flight Phase	Feature	Rationale
Takeoff	Prediction Lookahead Time	Large impact on uncertainty due to variable ground processes; 1-minute resolution used.
	CDM Airport Classification	CDM and AAT airports provide higher-quality ETOTs via offblock message.
	Departure Airport	Differences in procedures, infrastructure, and message behavior across airports.
Departure Route	CDM Airport Classification	Affects the availability and fidelity of trajectory prediction messages.
	Departure Airport	Influences typical routing, shortcut behavior, and SID complexity.
En-Route	Prediction Lookahead Time	Longer horizons lead to increased uncertainty in predicted arrival at IAF.

Table 4.6: Summary of processed EFD messages

Metric	Value
Dataset Overview	
Dataset period	2023-02-27 to 2024-04-26
Eligible EFD messages	2,122,265
Flights processed (total)	257,901
Takeoff Phase (ETOT Prediction)	
Total Flights	163,216
Flights from standard airports	71,349
Flights from CDM airports	70,957
Flights from AAT airports	20,910
Departure Route Phase	
Flights used	122,900
En-Route Phase (COP Analysis)	
Messages processed for COP analysis	13,135,881
Flights used in COP analysis	189,680

Table 4.6 provides an overview of the data used in the coming sections. The dataset uses over one year of EFD messages, in which over 250,000 flights inbound to Schiphol are represented. Approximately 2.1 million eligible messages were processed, of which subsets were used for various analyses including takeoff planning, departure route comparison, and en-route trajectory prediction.

4.5. Uncertainty in Takeoff Planning

Previous research has shown that the largest error source in the prediction using FUM and EFD messages is the uncertainty in the takeoff planning [19] [20]. This section aims to investigate this further, to be able to accurately represent it in future simulation. The EFD messages from more than one year have been processed, and takeoff errors calculated.

4.5.1. Methodology

The takeoff prediction error is defined as the difference between the ATOT and the ETOT:

$$\mathrm{E} = T_{\mathrm{ATOT}} - T_{\mathrm{ETOT}} \tag{4.2}$$

The lookahead time is defined as the difference between the ETOT and the message timestamp:

$$\text{Lookahead} = T_{\mathrm{ETOT}} - T_{\text{message}} \tag{4.3}$$

These values were calculated for each EFD message and grouped into 1-minute lookahead buckets. Results were then divided by airport CDM classification: CDM-enabled, AAT, and standard (non-CDM) airports. In the following figures, a positive error means a delayed aircraft with respect to the ETOT.

4.5.2. Impact of CDM Airport Type

Figure 4.6 shows the ETOT prediction error per message, divided by CDM type. CDM-enabled airports show a significantly lower spread in prediction errors, particularly between 8 and 14 minutes of lookahead time. This can be attributed to the off-block messages. AAT airports also perform better than standard airports, also due to off-block reporting mechanisms.

The increasing spread in prediction error closer to the ETOT is explained by a drop-off in message updates as most off-block messages have been sent, while some aircraft are already getting airborne.

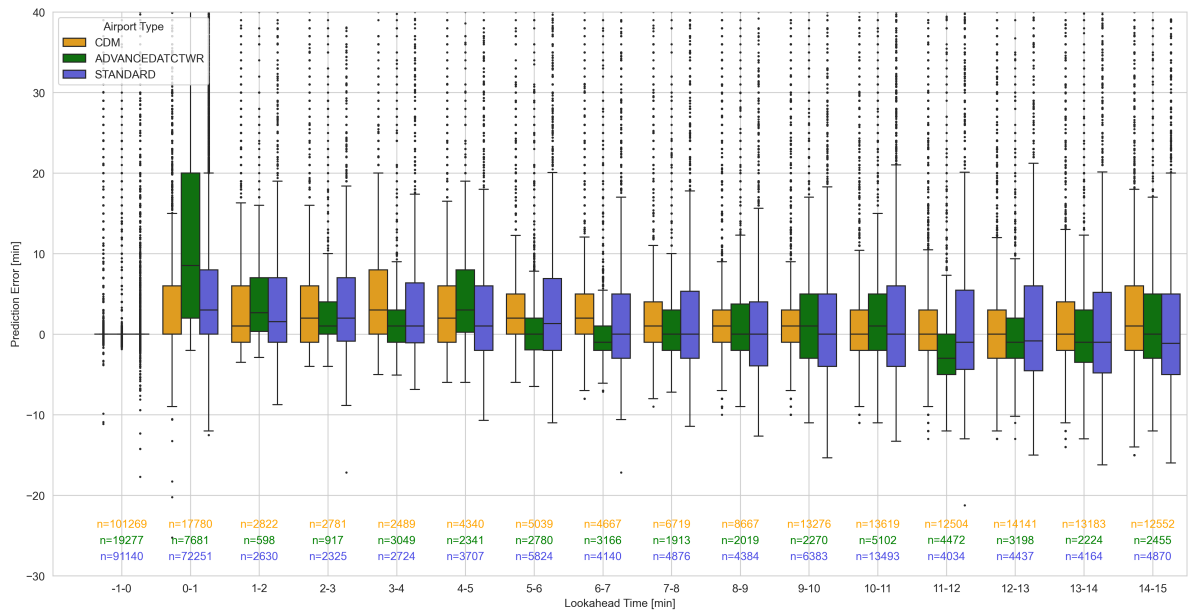

Figure 4.6: Estimated Take-Off Time (ETOT) error versus prediction lead time

4.5.3. Continuous Error Representation

To provide a continuous representation of takeoff prediction accuracy, prediction errors were extended from each message timestamp until the next available update. This is shown in Figure 4.7. Figure 4.6 essentially shows the error per message, thus discrete. To ensure that the error is shown in a continuous way, the errors per message are repeated until a new update is received. The result is shown in Figure 4.7. Equation 4.7 shows this in mathematical form. To ensure continuity in terms of actual time, not lookahead time, the message is repeated until the timestamp that the next message is received.

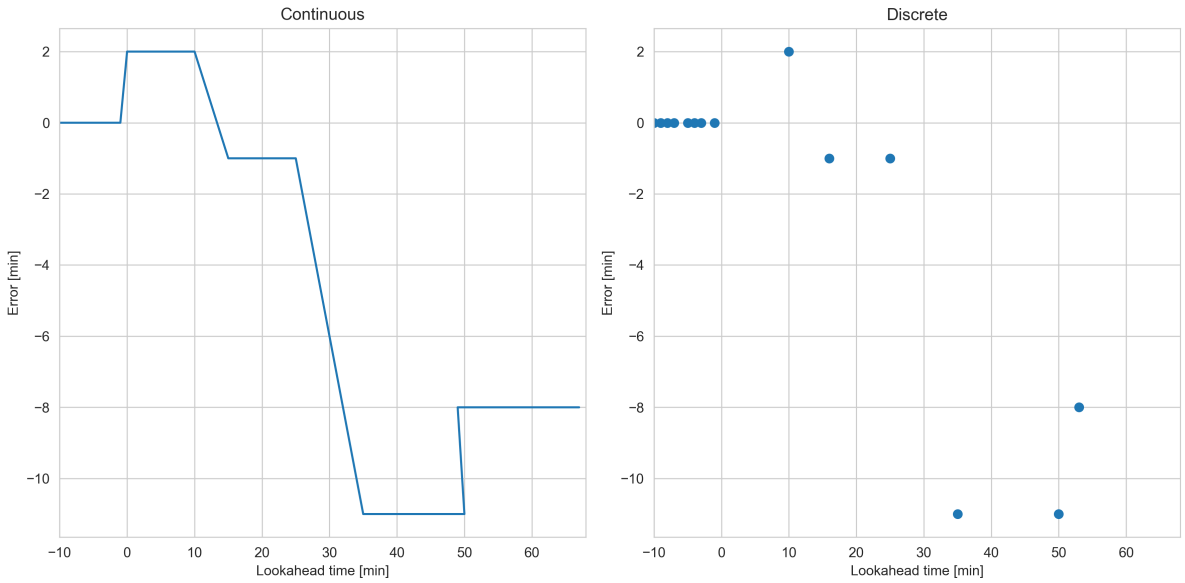

Figure 4.7: Example of a prediction error of an aircraft, continuous and discrete, per EFD message

$$Error = ATOT - ETOT \tag{4.4}$$

$$Repeatedmessages = Range(Lookahead, Lookahead + (T_{nextmessage} - T_{message}), 1) \tag{4.5}$$

$$\text{Error} = T_{\text{ATOT}} - T_{\text{ETOT}} \tag{4.6}$$

$$\text{Repeated messages} = \text{Range}\big(\text{Lookahead},\ \text{Lookahead} + \big(T_{\text{next message}} - T_{\text{message}}\big),\ 1\big) \tag{4.7}$$

The results of the continuous error function can be seen in Figure 4.8. In this figure it is clear that the accuracy of the takeoff predictions is much higher for the CDM airports in comparison to the standard airports, and somewhat higher when compared to the AAT airports. It is also clear that the spread of the error is reduced for each airport type when the lookahead time is reduced.

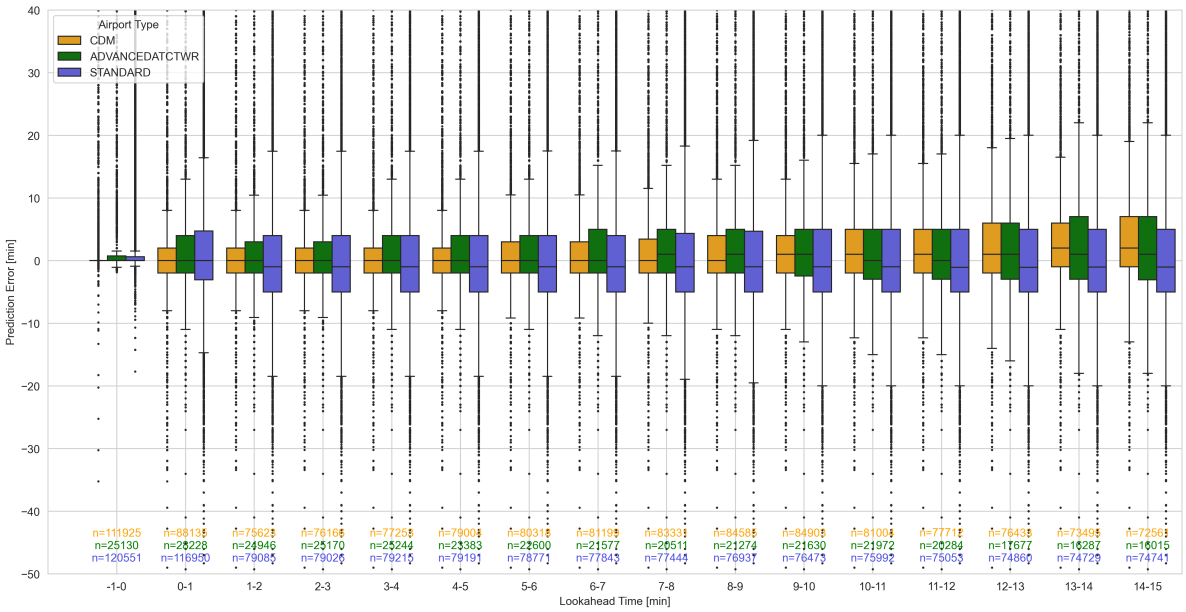

Figure 4.8: Continuous Estimated Take-Off Time (ETOT) Error versus prediction lead time, continuous per flight

4.5.4. Effect of CDM State

The takeoff time errors can also be represented according to CDM state, for CDM airports. Figure 4.9 shows the errors split according to the CDM status. The result clearly shows that the status ”actual off-block” highly increases the accuracy of the prediction. It can be concluded that this type of message causes the CDM and AAT enabled airports to provide the increase of accuracy in ETOT as seen in Figure 4.8 and previous figures.

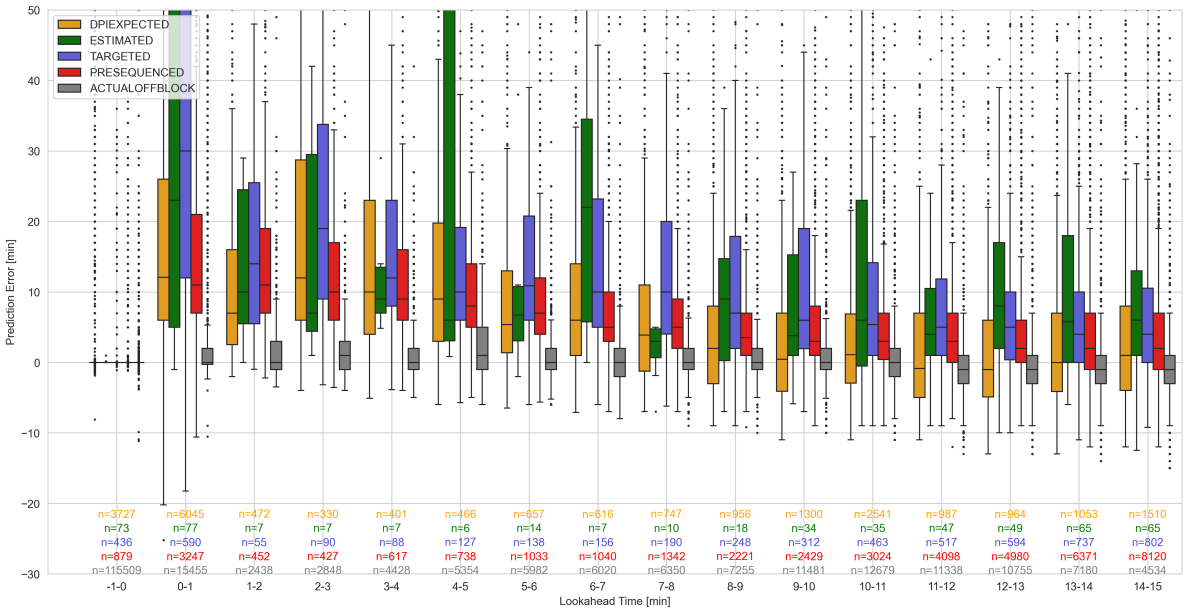

Figure 4.9: Discrete Estimated Take-Off Time (ETOT) Error versus prediction lead time, for different CDM statuses

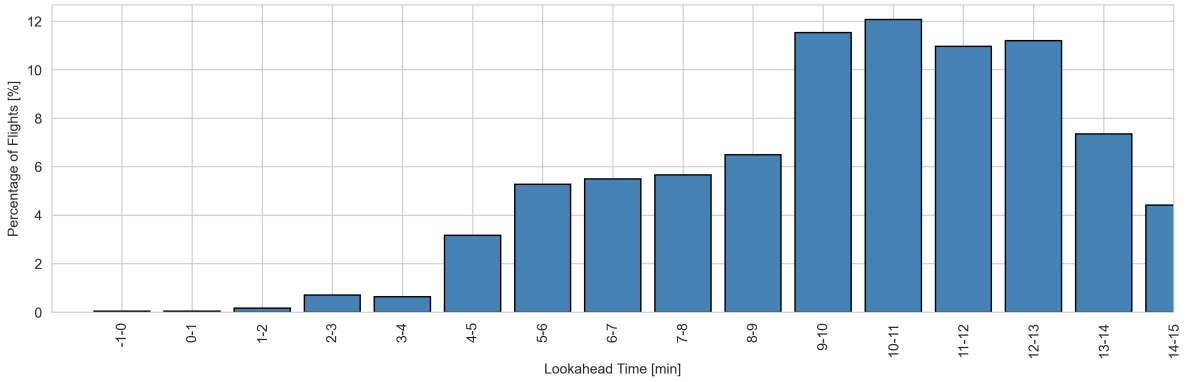

Figure 4.10: Off-block messages versus Estimated Take-Off Time (ETOT) lead time

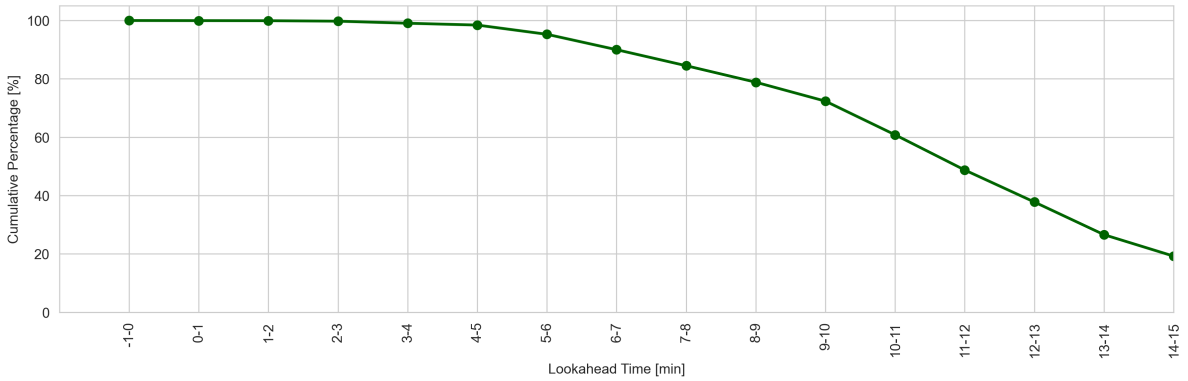

Figure 4.11: Cumulative percentage of off-block messages by Estimated Take-Off Time (ETOT) lead time

Figure 4.11 and Figure 4.10 show when an off-block message is sent for flights at an CDM or AAT airport, relative to their ETOT. It can clearly be seen that the reduction in spread in the takeoff prediction for the CDM airports and AAT airports coincide with the majority of aircraft going off-block. It is noted that at a lead time of 10 minutes, there are quite a lot of aircraft that are not yet off blocks, approximately 40%.

These results show that the lookahead time and the CDM state are of major influence on the accuracy of the ETOT. Further steps are to generate Probability Density Function (PDF)s to represent the uncertainty in these predictions in further research into the E-AMAN performance.

4.5.5. Analysis per Relevant Airport

In the figure below the same analysis as the previous subsections has been done for the airports within a 180NM radius of Schiphol. The figure shows that the difference between CDM, AAT and standard airports can also be explained as a result of the other differences between airports. The previous effects of the off-block messages for the CDM and AAT airports remain visible. Especially the error spread is reduced for the departures from those airports. Since these differences within CDM classification are quite substantial, and the spread in the previous CDM feature plots is likely influenced by the individual airports, the choice is made to continue the analysis with the takeoff prediction error statistics of individual airports. These distributions will be used to generate fitted PDFs for each airport in the simulation model.

ICAO	Airport Name	CDM Type
EDDL	Düsseldorf Airport	CDM
EBBR	Brussels Airport	CDM
EGSH	Norwich Airport	Standard
ELLX	Luxembourg Airport	Standard
EDDW	Bremen Airport	ADVANCEDATCTWR
EGSS	London Stansted Airport	ADVANCEDATCTWR

Table 4.7: CDM Classification of Selected Airports

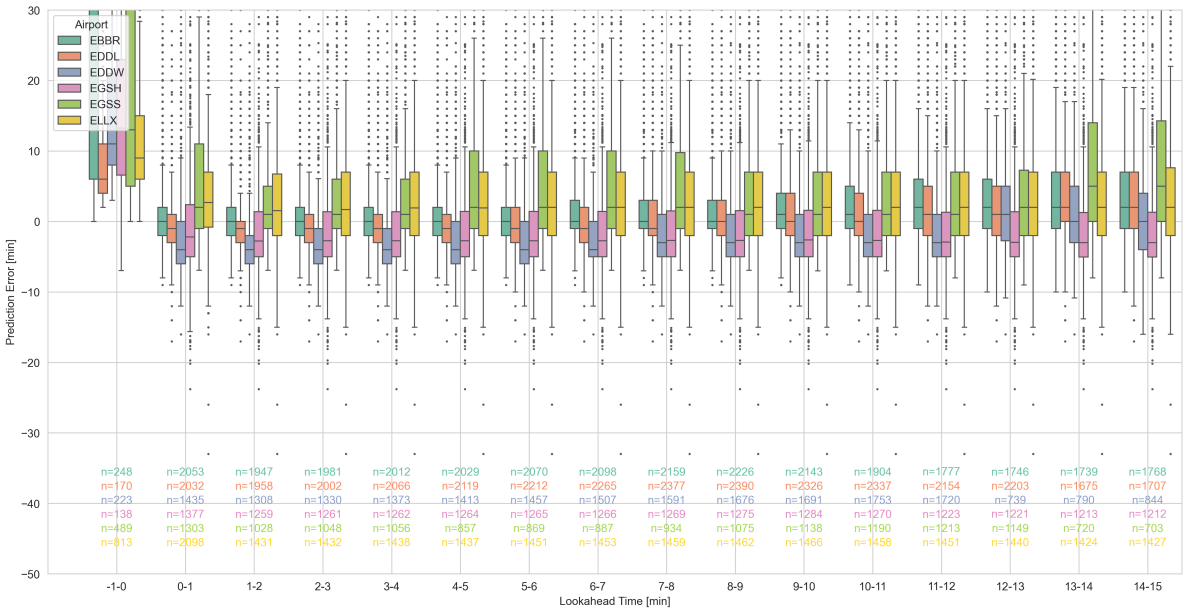

Figure 4.12: Continuous Estimated Take-Off Time (ETOT) error versus prediction lead time, per airport

4.6. Departure Route Prediction Error

As discussed in earlier sections, the departure route prediction might include errors that are significant to the (extended) arrival manager performance. This section investigates the errors in predicted versus actual departure routes, to be able to accurately represent it in future simulation. The investigation focuses on the estimated flight time to a fixed point above FL180 (or FL100 for Brussels). These errors are obtained by comparing the filed data (available at takeoff) and the actual flown trajectories derived from EFD messages The following data is shown in absolute time error, but also normalized for flighttime. This is done since the flighttime up to the point of calculation could be drastically different depending on multiple factors. The normalization ensures proper representation in future simulation. Once again, one of the features that could distinguish different error characteristics is the CDM type of the departure airport. The comparison in this section will first focus on the differences between the three types, as previous research had concluded that major differences exist between different CDM airport types in terms of prediction accuracy at the COP. The difference is at least partially due to the differences before takeoff, but a difference in the departure route prediction accuracy cannot be ruled out.

To calculate the errors (E) as shown in Figure 4.14 and Figure 4.15, the differences between predicted and actual flighttime up to a certain waypoint were calculated. A schematic representation is shown in Figure 4.13, while the following formulas were used.

$$E_{\text{departure route}} = \underbrace{\left(T_{\text{Point Actual}} - T_{\text{ATOT}}\right)}_{\text{Actual flight time}} - \underbrace{\left(T_{\text{Point Predicted}} - T_{\text{ETOT}}\right)}_{\text{Predicted flight time}} \tag{4.8}$$

$$\text{Normalized } E_{\text{departure route}} = 100 \cdot \frac{E_{\text{departure route}}}{\underbrace{\left(T_{\text{Point Predicted}} - T_{\text{ETOT}}\right)}_{\text{Predicted flight time}}} \tag{4.9}$$

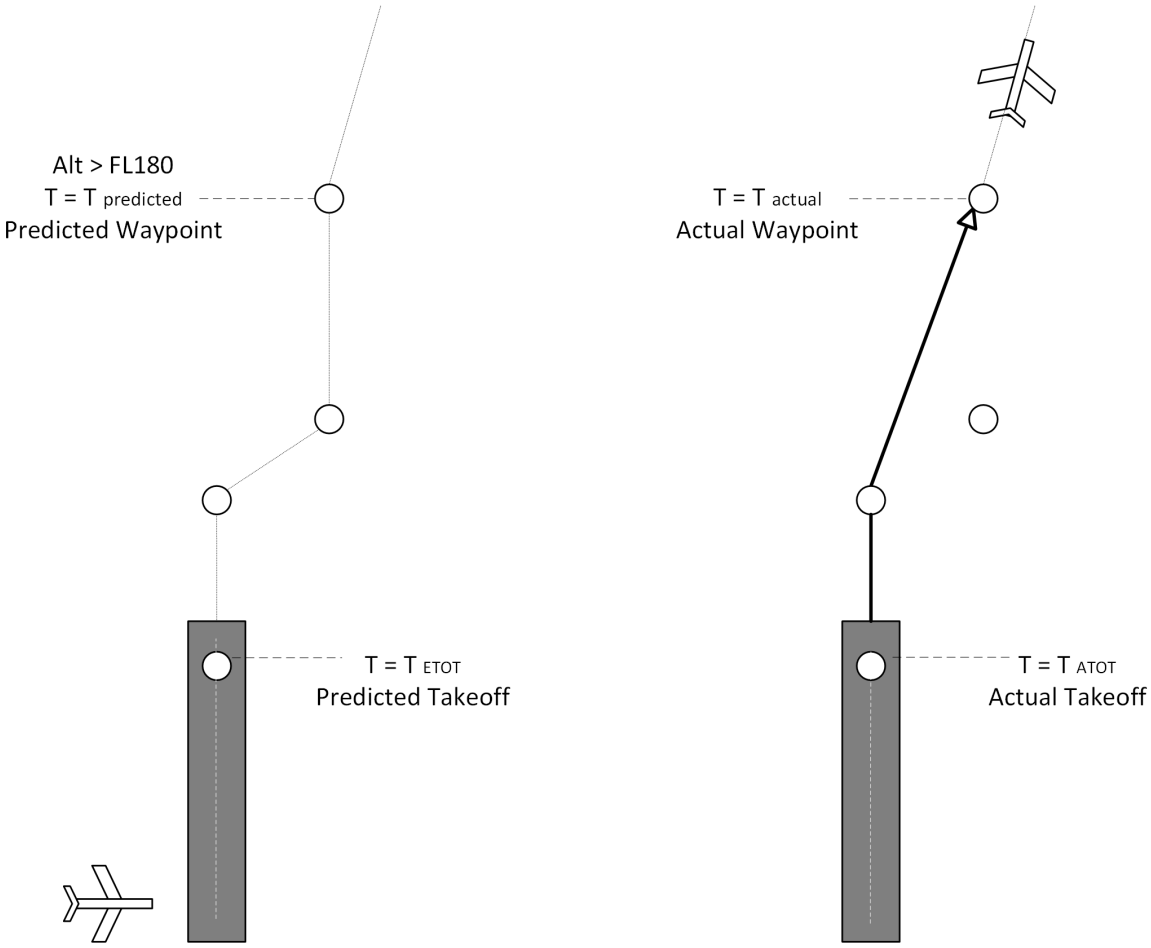

Figure 4.13: Schematic representation of departure route error calculation

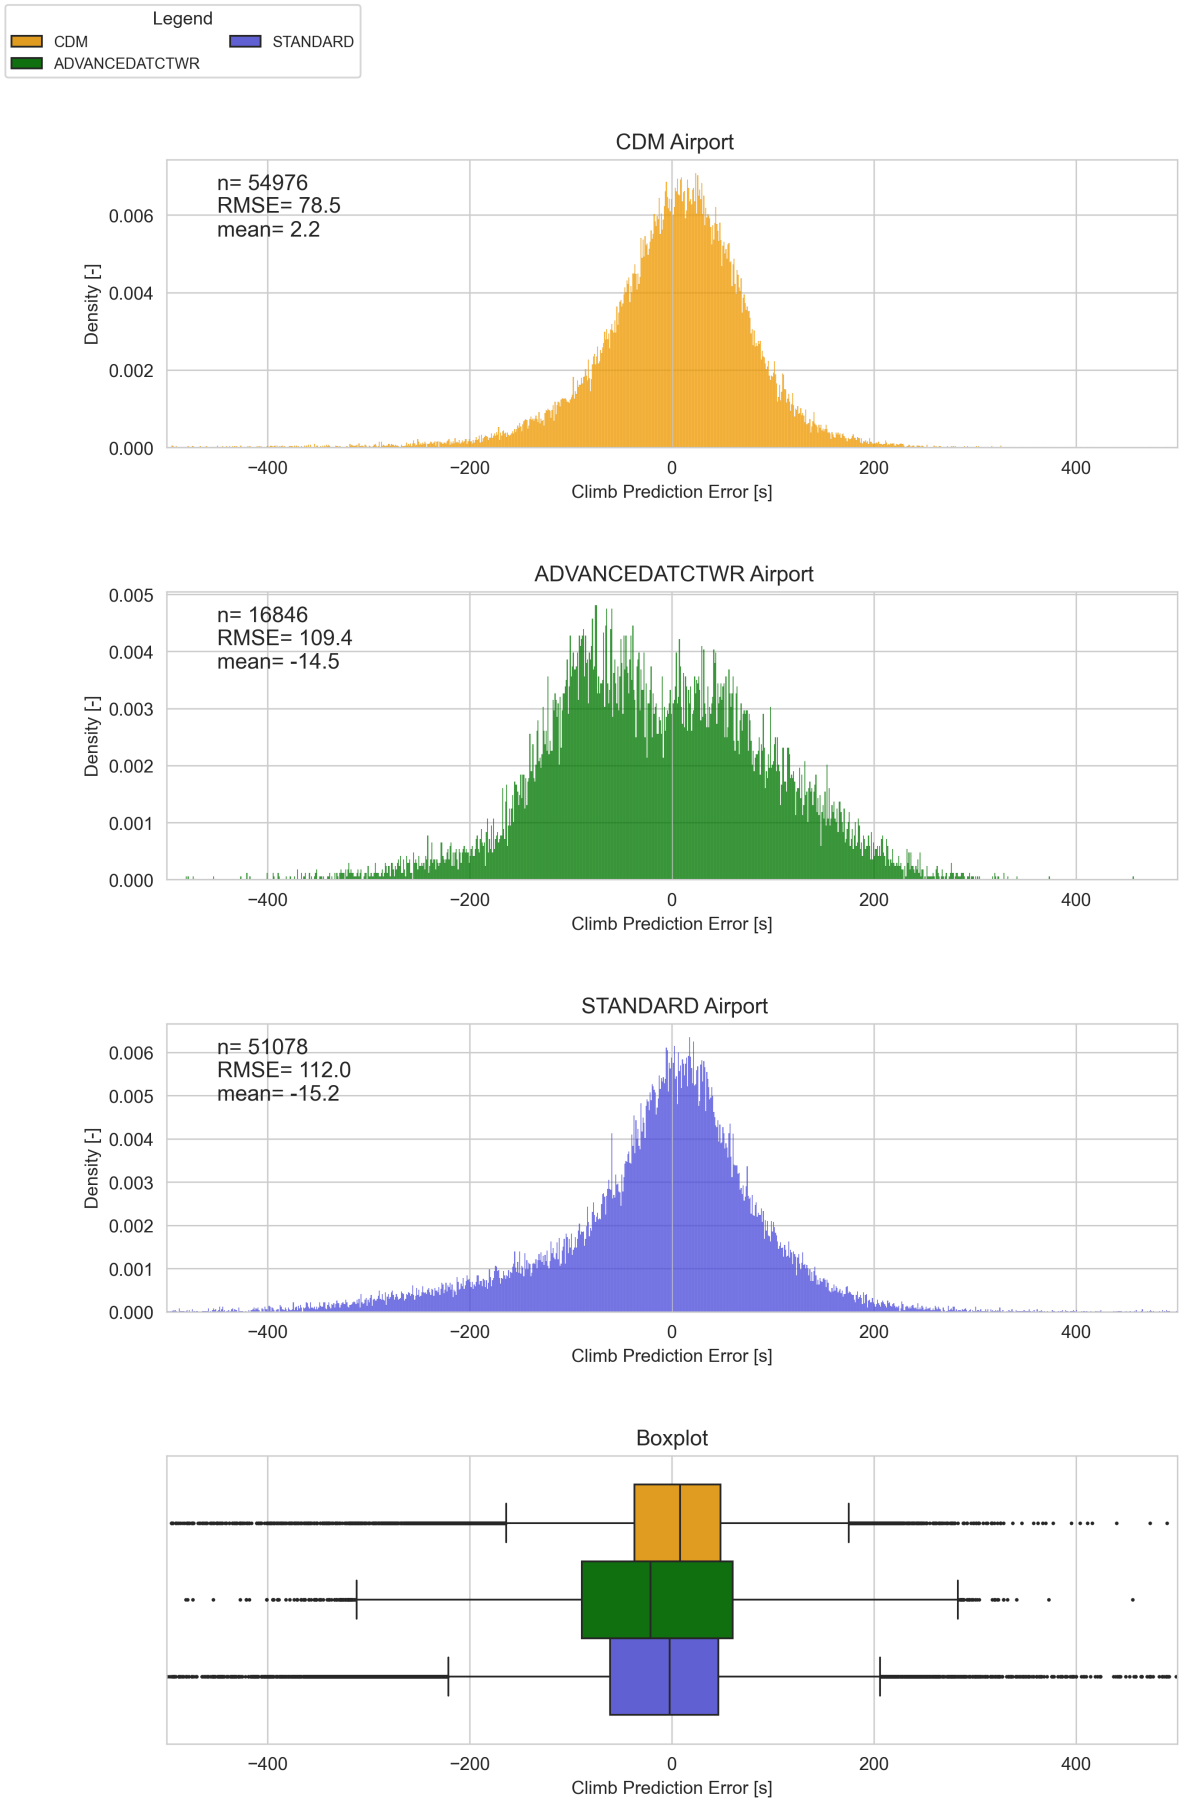

Figure 4.14: Departure flighttime error distribution

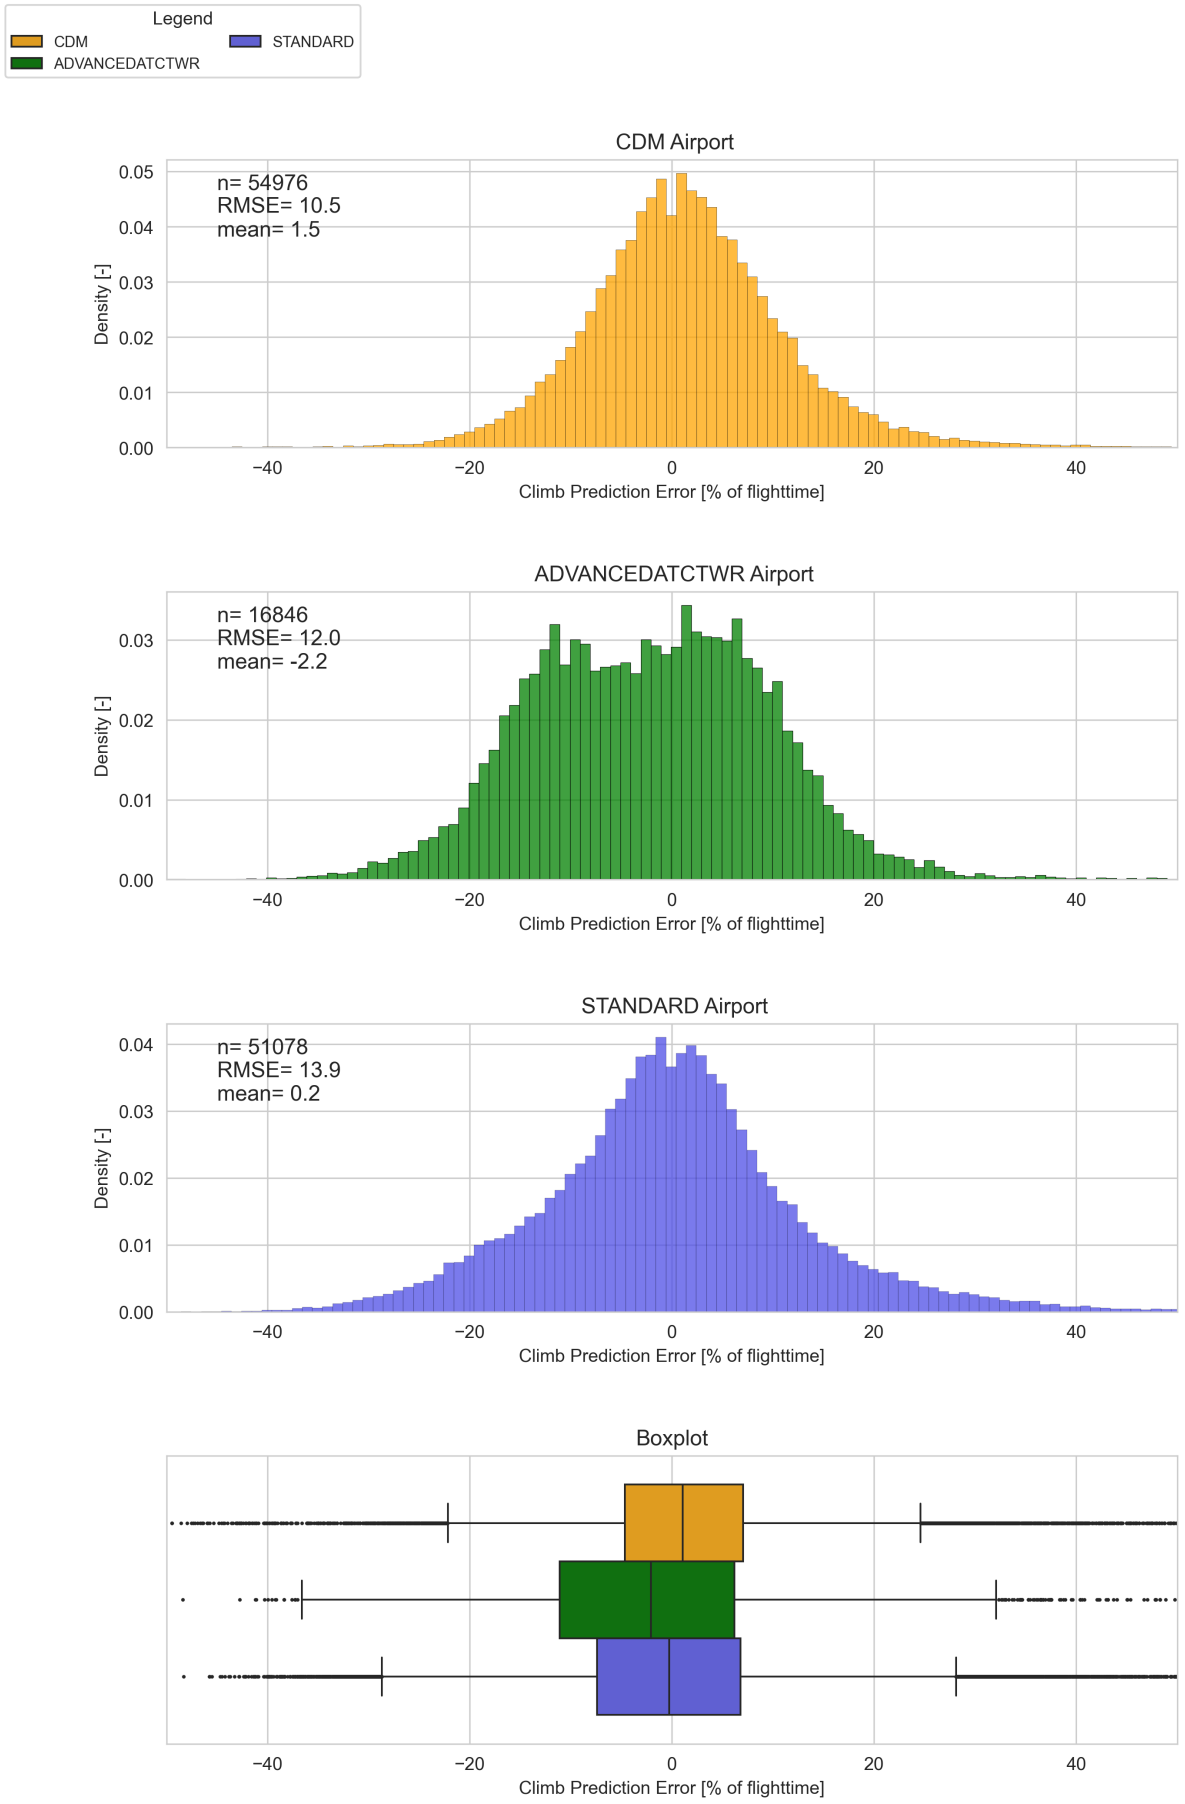

Figure 4.15: Normalized departure flighttime error distribution

In the figures above it is visible that the variation in departure route errors differs less between the different CDM airport types when compared to the departure planning errors. The magnitude of the errors is also much smaller. Though, these magnitudes could still have an effect on the performance of the E-AMAN.

Furthermore, it is clearly visible that the AAT airports have two peaks in the errors of the departure route plannings. This indicates that there might be major differences between airports and the accuracy of the departure route prediction. This suggests that departure route prediction errors might be quite airport specific.

These findings indicate that while overall error magnitudes are smaller than in the takeoff phase, there is still considerable variation between airports. To further investigate this effect, the following figures compare the normalized error distributions at two AAT airports.

As an example of the possible differences, the following figures show the normalized departure flight-time error distributions at London Gatwick and London City, both AAT airports. It is visible that the prediction error mean and spread is quite different. In general, it can be concluded that the feature of individual airports plays a larger role than the feature of CDM type of the departure airport. This also greatly argues for the representation in simulation of individual airports with respect to departure errors.

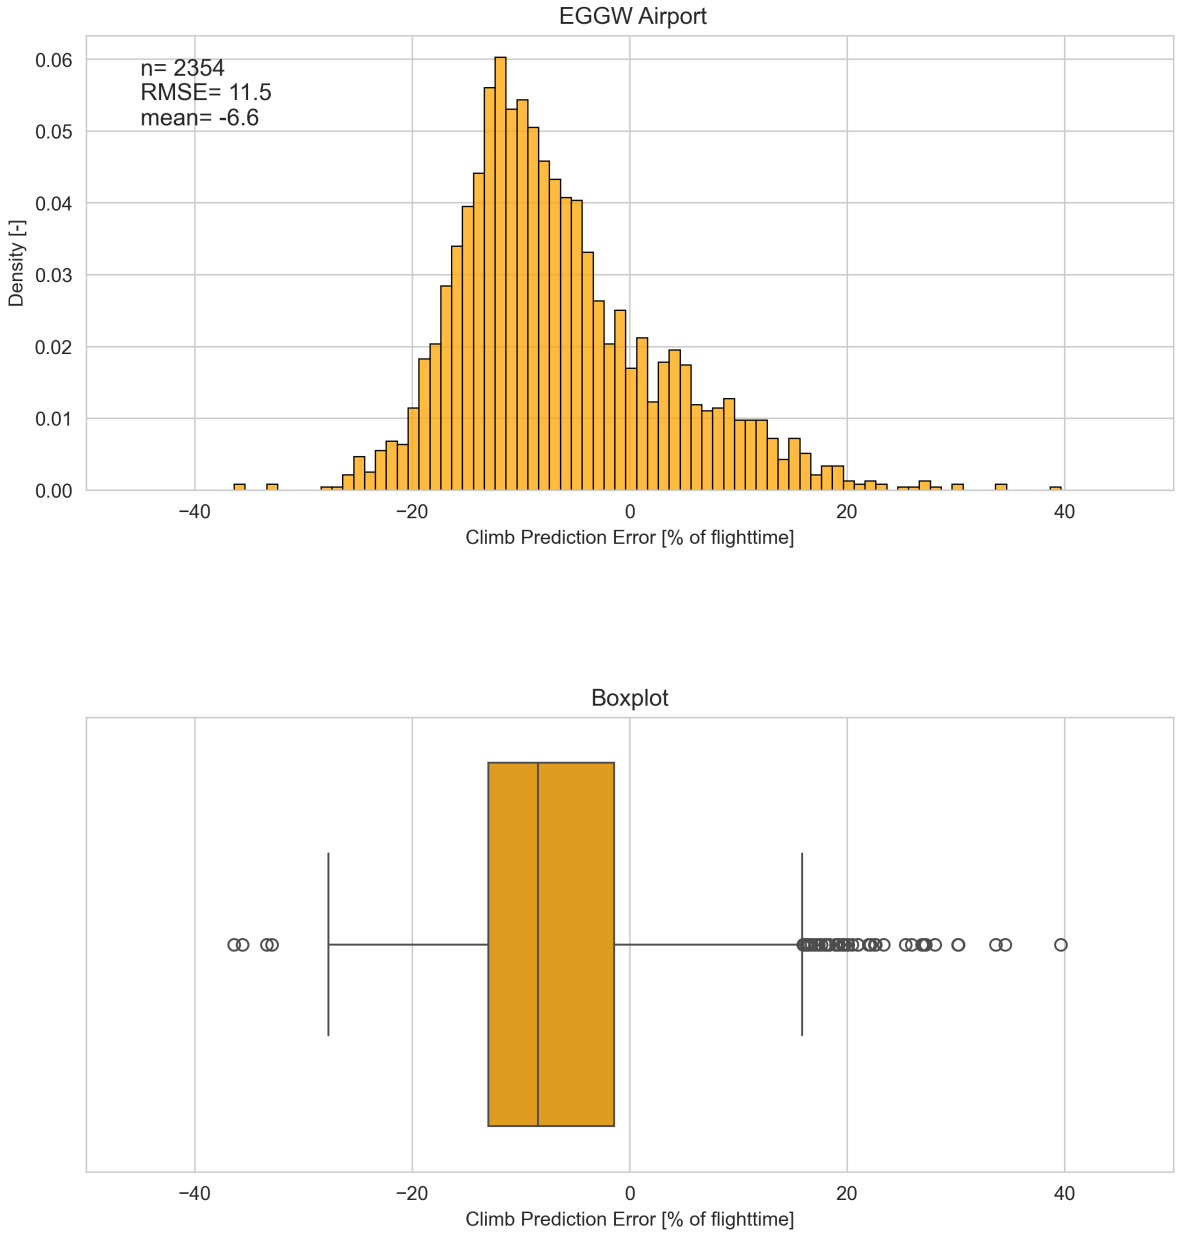

Figure 4.16: Normalized departure flighttime error distribution at London Gatwick

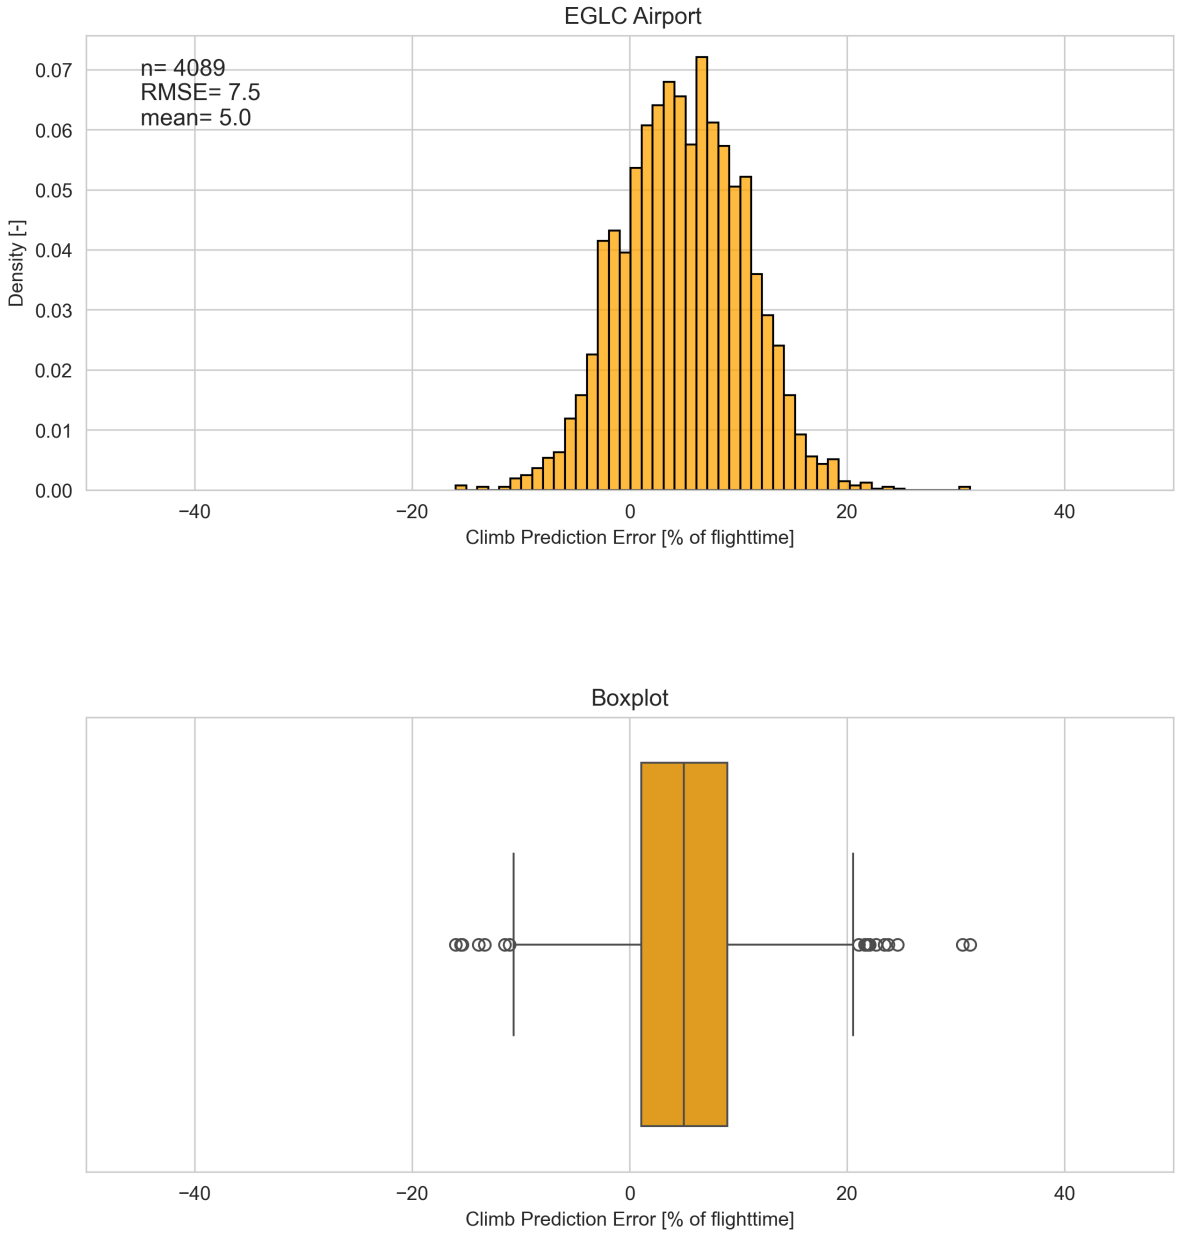

Figure 4.17: Normalized departure flighttime error distribution at London City

4.7. Uncertainty in En-Route Trajectory Prediction

The effect of the trajectory prediction uncertainty on the performance of the E-AMAN is still unknown, while the magnitude could be significant. The goal of this section is to investigate the en-route uncertainty, once again to be able to accurately represent it in future simulation.

The trajectory prediction uncertainty can be split into two categories: predictions made outside the local FIR, where controller intent is largely unknown, and those made within the FIR, where the aircraft is typically under the influence of the local ANSP and AMAN system. Each of these cases introduces different types and magnitudes of error.

4.7.1. Outside FIR

When the aircraft is still outside the Dutch FIR, there is less knowledge about the intent of an aircraft or controller. The unknown intent introduces errors with respect to the prediction of handover at the COP.

There is no clear quantification of these errors, so this subsection sets out to quantify these errors. Once again, they are normalized for flighttime.

The errors were calculated by comparing predicted and actual times at the COP, using Equation 4.10 and Equation 4.11, while Figure 4.18 shows a schematic representation of the calculation of these errors.

$$E_{\text{COP}} = \underbrace{\left(T_{\text{COP Actual}} - T_{\text{Prediction}}\right)}_{\text{Actual flighttime to COP}} - \underbrace{\left(T_{\text{COP Predicted}} - T_{\text{Prediction}}\right)}_{\text{Predicted flighttime to COP}} \tag{4.10}$$

$$\text{Normalized } E_{\text{COP}} = 100 \cdot \frac{E_{\text{COP}}}{\underbrace{\left(T_{\text{COP Predicted}} - T_{\text{Prediction}}\right)}_{\text{Predicted flighttime to COP}}} \tag{4.11}$$

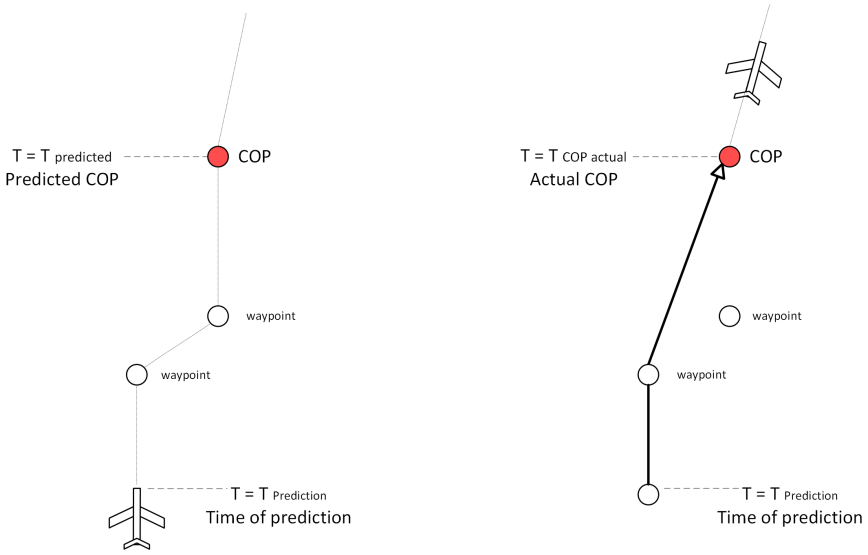

Figure 4.18: Schematic representation of departure route error calculation

Here, $T_{\text{Prediction}}$ represents the timestamp at which the trajectory prediction was made, i.e., the time of the EFD message. A lookahead window of 12 to 16 minutes was chosen as this is representative for the amount of planning time outside of the FIR with a FH of 20 minutes.

Figure 4.19 shows the distribution of the absolute time prediction error to the COP, while Figure 4.20 presents the same data normalized to the lookahead duration. These figures clearly indicate that the uncertainty in trajectory prediction outside the FIR is substantial. The effect of this uncertainty on the performance of the E-AMAN is unknown.

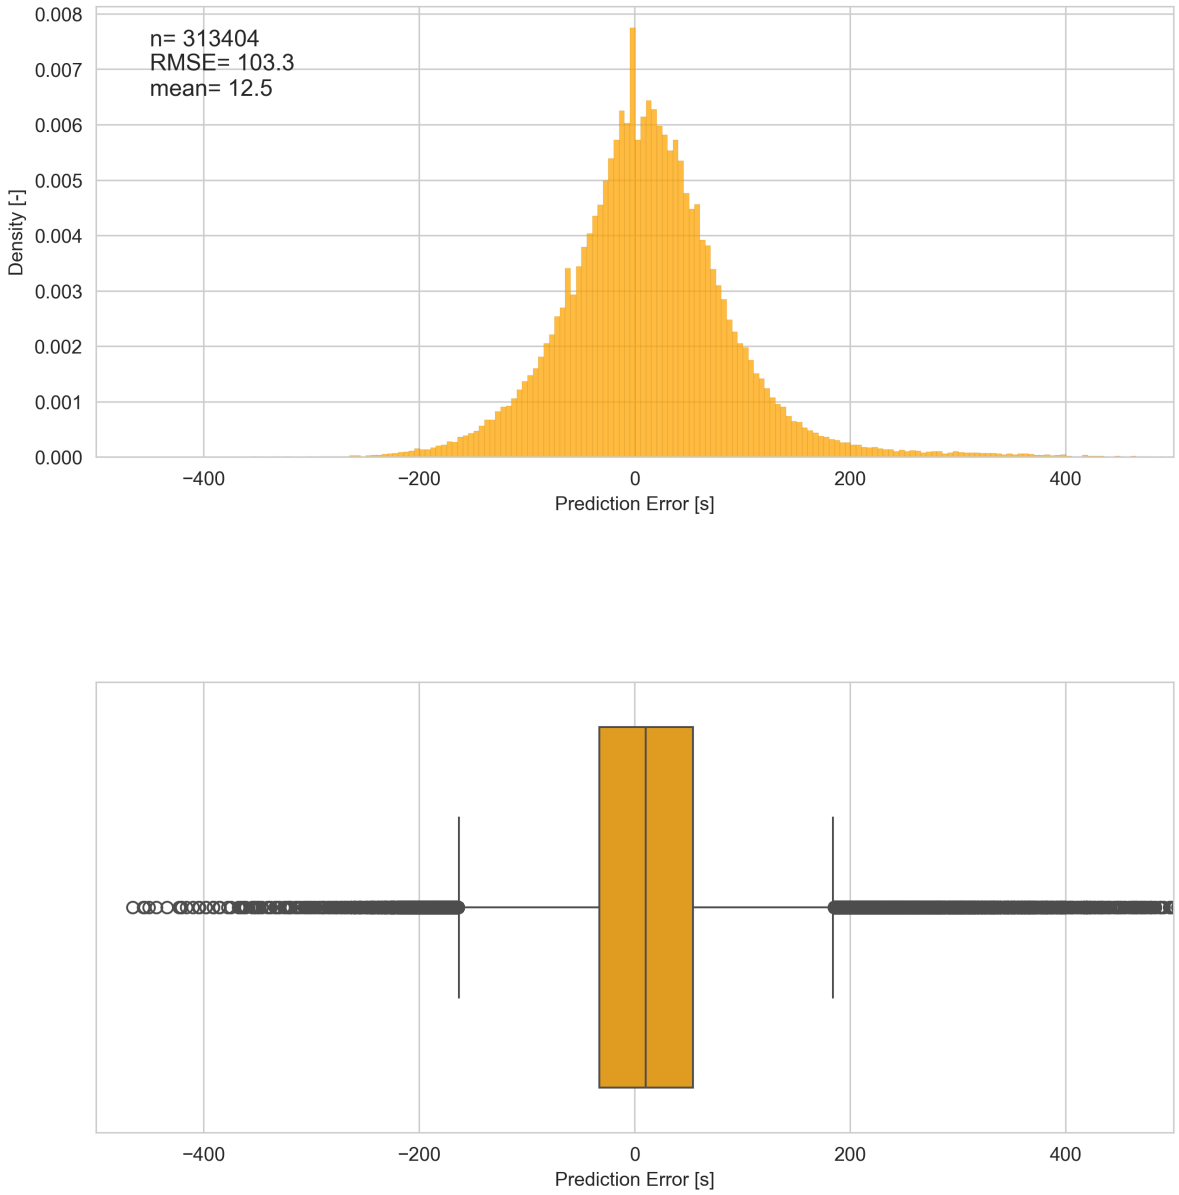

Figure 4.19: En-route flighttime error distribution to predicted Estimated Time Over (ETO) COP

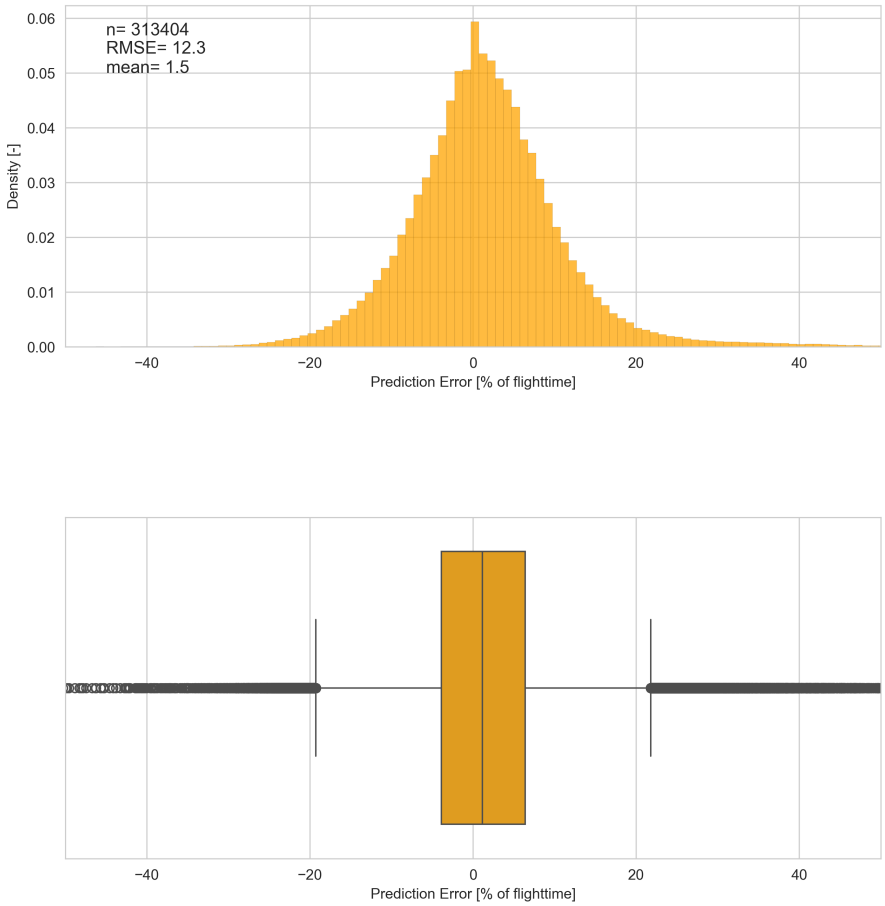

Figure 4.20: Normalized en-route flighttime error distribution to predicted Estimated Time Over (ETO) COP

4.7.2. Within FIR

Once an aircraft enters the Dutch FIR, the situation with respect to intent changes. The controller intent is now controlled by the AMAN, which causes the largest factor of uncertainty to be removed. Though, the error cannot be measured with EFD messages, as this would lead to an incorrect representation of the TP error within the FIR, as the intent based error would be mainly driven by the AMAN. The representation of the uncertainty within the FIR will be done using statistics from literature. As presented in subsection 3.5.1, the internal error used by controllers is approximately 1.5%, while [15] has shown that the worst case within the FIR is an absolute error is 15.9 seconds.

4.8. Solution Space

A consultation with controllers [24] has been done to have a better indication where the limits of the solution space lie. During the consultation, the controllers considered a delay of 6 minutes, or a expedition of 2 minutes feasible, for a flight from Brussels through the RIVER IAF, in Dutch airspace. This case represents one of the more constrained situations due to the short distance between the FIR boundary and the IAF. For other IAFss, the available solution space is larger, due to the increased distance from the FIR boundary. Holding patterns and communication with adjacent centers would allow for more delay absorption, though it is not desirable.

4.8.1. Delay

When encountering a positive Time To Lose or Gain (TTLG), a delay must be achieved for that aircraft to achieve crossing the IAF within the approach window. There are two main possibilities to achieve a delay: a speed reduction or a longer route, or a combination of these.

A speed reduction can be done during the cruise or initial descent phase, when the aircraft is controlled

using a mach number, or using airspeed. After crossover, when the speed control is switched from mach to airspeed, the airspeed generally remains constant, until crossing FL 100. Below FL 100, the airspeed must be at maximum 250 knots [4]. Crossing the IAF is normally done at FL 100 or FL 090.

Depending on aircraft type and weight, controllers can instruct pilots to reduce to minimum clean speed, which is the lowest safe speed without deploying flaps. Typical values include: • 230 knots for heavy aircraft (B777, B787, A330) • 210 knots for medium aircraft (A320, B737) • 190 knots for regional jets (E190, E175)

During cruise and initial descent, a mach reduction of 0.04 should be possible for nearly all aircraft [24].

If speed instructions within the FIR do not provide enough delay to cross the IAF within the approach window, there are two main options to increase the track length of the aircraft: a dogleg, or a holding. A dogleg is a detour, often instructed by giving the aircraft a heading slightly away from the next waypoint. After a while, the aircraft is instructed to turn to the next waypoint once again. A holding pattern entails circling near a certain waypoint, often near the IAF.

4.8.2. Speed-Up

In cases of negative TTLG, when an aircraft is expected to arrive too late, flight time must be reduced. There are two ways to achieve this: shortening the route, and speed increases.

The most straightforward method is to instruct the aircraft to skip waypoints, known as issuing “directs.” This can save several nautical miles and reduce flight time.

During cruise, most aircraft already operate close to their optimal or maximum Mach number, limiting opportunities for further speed-up. During descent, however, many aircraft can perform high-speed descents, maintaining higher airspeed until FL100.

However, the maximum achievable time gain is relatively limited. This shows that absorbing delay to adhere to the approach window is achieved much more easily.

4.9. Proposed Solutions to Uncertainty

4.9.1. Required Time At

The RTA solution, as presented in [35], offers a promising approach to managing pop-up aircraft. By calculating a Required Time of Arrival at the IAF before departure, delay can be absorbed on the ground. This allows aircraft to arrive at the IAF within the approach window without further instructions. This reduces taskload and fuel burn.

However, there are some downsides to this mitigation strategy when implemented at Schiphol. First, there are no suitable communication standards present at Schiphol, to communicate to nearby airports. In previous trials, such coordination was done by telephone, which is not scalable for daily operations. SWIM might be able to overcome this hurdle, though the implementation of a standard has not progressed far enough to start using it in systems in the coming years.

Second, the takeoff environment at departure airports near Schiphol is highly constrained. RTA trials have been done at airports with relatively few departures, and thus few departure sequence constraints. To implement a RTA system at a major airport successfully, a lot of integration needs to take place at each departure airport to be able to accommodate aircraft and their RTA calculated takeoff time. It could be that it is not feasible at all.

A more realistic option is to regulate aircraft using the existing CTOT infrastructure. This would provide a takeoff window of approximately $[-5, +10]$ minutes around the target time [36]. While this allows for strategic regulation, it does not reduce uncertainty at the FIR entry to below the $2\,\mathrm{min}$ threshold that [19] shows is necessary for E-AMAN to outperform current AMAN systems.

Therefore, while RTA is promising mitigation strategy, the practical implementation at Schiphol within an extended AMAN is currently infeasible without proper information sharing systems and procedures. However, CTOT assignment could still be used together with other solutions discussed below.

4.9.2. Delayed Slot Strategy

Based on the conclusions in [19] and the statistical analysis of takeoff and departure route uncertainty presented above, many aircraft are expected to arrive at the FIR within a relatively narrow time window. However, when an aircraft arrives late for its assigned slot, the impact on the sequence is substantial [19]. Speeding up en route is difficult, while absorbing some additional delay is relatively easy. A strategy in which pop-up aircraft are scheduled later than their expected arrival could take advantage of this. This transfers delay from the entire fleet to the smaller amount of pop-up flights, possibly leading to an overall improvement in E-AMAN performance.

The scheduling would be done according to the statistics found above. For each airport, the expected maximum delay is used, at the time before ETOT. The expected maximum delay which would be suggested to use in this solution would be the 75th percentile. This would ensure that most aircraft would be able to meet their EAT and slot. The formula for the aircraft slot would be the following:

$$ETA_{slot} = ETA_{EFD} + \text{expected maximum delay} \tag{4.12}$$

With the slot calculation based on the ETA, the slot would be pushed back to a time at which the aircraft is nearly certain to be able to arrive. In this case, the majority of pop-up aircraft are required to absorb additional delay. In case the pop-up aircraft is unable to arrive at the IAF within the approach window, a new slot must be assigned. If the pop-up aircraft is estimated to arrive at the IAF before the approach window, the aircraft can either absorb the additional delay, or receive an earlier slot if available.

A short approximation can be made to test whether this strategy is worth investigating further.

In [19] it is shown that increasing the horizon reduces the average Low-Level Delay Absorption (LLDA) from 141 to 116 seconds. If 3.7% of flights are popup flights, and it is assumed that an expected delay of 5 minutes is used, with the average aircraft taking off at the ETOT, the worst case LLDA can be estimated in the following manner:

$$\begin{aligned} LLDA_{\text{delayed slot}} &= \text{Average } LLDA + \text{pop-ups} \cdot \text{average additional delay} \\ &\approx 116\,\text{s} + 3.7\% \cdot (5 \cdot 60)\,\text{s} \\ &\approx 116\,\text{s} + 11.1\,\text{s} \\ &\approx 127.1\,\text{s} \end{aligned} \tag{4.13}$$

This approximation shows that the gains in the reduction in LLDA are still quite substantial, even though some aircraft do have to absorb more LLDA. The estimate for the average additional delay might be off. However, the estimation assumes that the pop-ups also have the normal average LLDA which would make the estimated result less optimistic.

In the longer term, this delayed slot strategy could be combined with the RTA concept described above. The RTA could be based on the delay statistics, and communicated with airports that are equipped for RTA operations.

4.9.3. Optimized Rescheduling

Another method of reducing workload and possibly increasing efficiency is automatically calculating a solution based on the current planning. A support tool would remove the ambiguity of the human side of the approach planner, and could provide the planner with a better solution. An approach would be to employ a cost function based on the KPI's suggested later in chapter 5, which could automatically determine which slot would be the best for a pop-up aircraft. Certain boundaries would have to be set to make sure the solution presented to the approach planner can stay within the solution space as presented in section 4.8.

4.9.4. Planning Off-block Flights

Once an aircraft is off-block, the trajectory prediction becomes significantly more accurate, increasing the likelihood of arriving within the approach margin at the IAF. Zürich airport is currently using this

method as a pop-up mitigation strategy [17]. The downside of this method is that at Schiphol at least 40% of pop-ups cannot be planned in an E-AMAN environment due to their departure from non-CDM airports. Furthermore, some aircraft have a planned takeoff time shortly after the off-block message is received, when the FH has already passed. Additionally, some uncertainty remains in the off-block flights.

4.9.5. Reducing Uncertainty

Reducing uncertainty would only be a standalone option when the error of most pop-up aircraft at the FIR entry would be reduced to approximately two minutes, as shown by [19], or in combination with one of the options above. To be able to reduce the uncertainty, solutions such as EPP could be employed, or the CTOT presented above. The issue with these options is that the infrastructure is not present, and the gains from enabling CDM at airports is also limited. It is noted that the options presented above could possibly improve the performance of the E-AMAN further when combined with a reduction of the uncertainty in takeoff.

4.10. Conclusion

This chapter has investigated the sources and magnitude of time prediction uncertainty affecting the performance of an Extended Arrival Manager (E-AMAN) at Schiphol. Using a full year of EFD data, the uncertainty has been analyzed in three main flight phases: takeoff planning, the departure route and en-route.

The results show that:

- A significant proportion of inbound traffic originates from airports near the freeze horizon. If the horizon is extended from 120NM to 180NM, the share of pop-up aircraft more than doubles. A lot of traffic departs from near-horizon airports.
- The largest source of uncertainty is found in the takeoff phase, especially at standard (non-CDM) airports. CDM and AAT airports show substantially improved accuracy due to off-block messages. The results also vary between airports.
- Departure route prediction errors are smaller in magnitude but vary significantly between individual airports.
- En-route trajectory prediction uncertainty is substantial when aircraft are outside the FIR, mainly due to unknown controller intent. Inside the FIR, uncertainty is lower, based on literature.

Based on these findings, several mitigation strategies were discussed. Next, chapter 5 will dive deeper into how the results from this chapter will be used in simulating an AMAN environment.

5

Research Design

This chapter introduces the research design used to evaluate the impact of takeoff and trajectory prediction uncertainty on the performance of the Extended Arrival Manager (E-AMAN). Building on the findings from chapter 4, it outlines the overall methodology in section 5.2, the independent and dependent variables used in the experiments in subsection 5.2.2 and 5.2.3 and the simulation setup in section 5.3. Finally, section 5.5 discusses the future steps to be taken in the development.

5.1. Relevance

The previous chapters have shown that the Arrival Manager (AMAN) is a vital component of the air traffic management system. Extending the AMAN planning horizon could improve performance. However, current limitations related to stability, controller and pilot taskload hinder such extensions. A key challenge lies in the handling of aircraft that take off within the planning horizon, often referred to as ”pop-up” flights.

In chapter 4, the uncertainty in trajectory prediction was analyzed for each flight phase. One year of ETFMS Flight Data message (EFD) data was processed to quantify the statistical error characteristics of (1) the pre-departure phase, (2) the departure route, and (3) the en-route trajectory. These errors were represented using probability density functions to be used in simulation. From this analysis, several conclusions can be drawn. First, the pre-departure phase remains the main source of uncertainty. This uncertainty decreases closer to the estimated takeoff time (Estimated Takeoff Time (ETOT)), particularly for flights departing from Collaborative Decision Making (CDM) and Advanced ATC Tower (AAT) airports. This improvement is partially due to off-block messages being sent at those airports, while they also vary per airport. Second, the uncertainty in the departure route varies by airport but is still significant. Third, the en-route trajectory prediction, especially outside the Flight Information Region (FIR), also introduces non-negligible errors.

All three phases of uncertainty should therefore be considered when evaluating E-AMAN performance. Additionally, an analysis of freeze horizon extension revealed that the percentage of pop-up flights increases from 1.8% to 3.7% when extending the freeze horizon, with many aircraft departing just outside of the freeze horizon. The share of such departures from CDM and AAT airports is approximately 60 %.

To date, no study has evaluated the impact of trajectory prediction uncertainty on E-AMAN performance. Earlier research often assumed a constant or limited influence. However, the substantial errors found in en-route phase suggest that trajectory uncertainty may have greater effect on E-AMAN performance than previously thought.

While various solutions for handling pop-up traffic have been proposed, no effective mitigation strategy has been found. A few mitigation strategies could be explored further. One such strategy is the delayed slot strategy, which schedules aircraft using expected delay statistics specific to the departure airport. This could reduce the negative effects of pop-ups. Another approach is to insert pop-up aircraft at

the optimal point in the sequence once they become airborne. Finally, limiting planning to aircraft that have issued an off-block message may also improve robustness if other mitigation strategies are not effective enough.

Previous simulations of the E-AMAN relied on pragmatic assumptions to simplify scenario modeling. However, these simplifications reduce fidelity, and it could lead to results that are less accurate. For further development, a higher-fidelity simulation is required to be able to investigate the effects of the above mentioned uncertainties.

This chapter presents a research plan to investigate the impact of takeoff and trajectory prediction uncertainty on the performance of the (extended) arrival manager, and to evaluate the effectiveness of various mitigation strategies, according to the research goal and questions posed in chapter 1.

5.2. Methodology

A high-fidelity simulation will be developed to investigate the posed research questions outlined in chapter 1. Simulated air traffic will be based on historical traffic scenarios, which will be controlled using a simulated E-AMAN. The resulting behavior will be measured to evaluate and compare different configurations. The scenarios used are introduced in subsection 5.2.1, with the independent variables and dependent variables respectively introduced in subsection 5.2.2 and subsection 5.2.3.

Uncertainty in Trajectory Predictor (TP) and takeoff and resulting prediction errors will be simulated using a Monte-Carlo approach. This allows for measuring the effects of this uncertainty and resulting errors, and allow for statistically meaningful results. The uncertainty statistics are based on the results of chapter 4.

The delayed-slot strategy, as well as the optimized rescheduling are both to be tested using this Monte-Carlo based uncertainty simulation, together with the effect between the AMAN and E-AMAN.

The simulation will be developed using BlueSky as a basis, with different new modules to be developed to simulate the AMAN.

5.2.1. Traffic Scenarios

The evaluation of the different E-AMAN configurations is to be done using traffic scenarios simulated in BlueSky. Since the AMAN is designed to work in high-capacity operations, the basic traffic scenarios will consist of a three-hour period with enough inbound demand to saturate two runways, also known as an inbound peak. The results from this scenario are to be used quantitatively to compare performance between different configurations. At least two of these scenarios are to be used. The scenarios are to be constructed from a real-life traffic scenario using the Demand Data Repository (DDR) [37] data format. Further details of the scenario generation can be found below in subsection 5.3.6. Aircraft are all planned via an Initial Approach Fix (IAF), with fixed routes within the Terminal Control Area (TMA) for simplicity.

The most common, and thus one of the relevant scenarios to evaluate is when there is wind from the south or south-west [38], and aircraft land in southernly direction onto 18R and 18C (the Polderbaan and Zwanenburgbaan). Aircraft inbound via IAFs RIVER and SUGOL plan to land at 18R, while aircraft inbound via ARTIP plan to land at 18C [23].

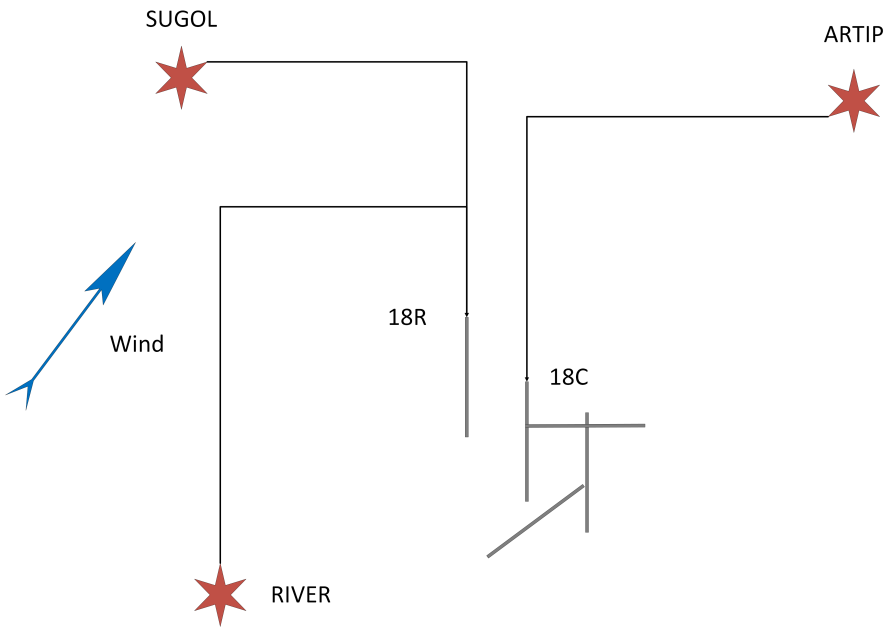

Figure 5.1: Main Scenario

Together with this quantitative analysis, a qualitative analysis will also be performed on scenarios that occur less frequent, but might be limiting towards E-AMAN implementation. The most limiting scenario is when aircraft are landing at runways 36R and 36C, with lower level winds coming from the north-east, with higher level winds coming from the south-west [23]. This scenario greatly reduces the solution space for absorbing delay for controllers, due to the limited time of aircraft arriving from the south in the Dutch FIR. Furthermore, the off-peak performance of the E-AMAN has not been investigated. This is also one scenario that is to be qualitatively investigated.

5.2.2. Independent Variables

Since this research aims to develop and evaluate solutions to uncertainty in arriving traffic, the independent variables are chosen such that the solutions can be evaluated against a baseline. The main independent variable is the prediction error, which is split up in takeoff, departure and trajectory prediction error. The influence of the error is to be evaluated, thus introduced to the simulation through an error module. The error is generated by sampling from a distribution based on statistical uncertainty analysis. This is done separately for take-off, departure route, en-route and within FIR.

The effectiveness of mitigation strategies will also be evaluated. The effectiveness will be tested for both the delayed slot strategy and the optimized rescheduling separately, and combined in configurations 7 as visible in Table 5.3. Finally, the freeze horizon will also be varied, to investigate the effect of the uncertainty and their solutions on the different possible E-AMAN configurations: current, required and extended.

Table 5.1: Independent variables with definitions

Variable	Definition	Values / Type
Freeze Horizon	Time (in minutes) after which the sequence is frozen and no further adjustments are made.	Integer (14, 20)
Trajectory Prediction Uncertainty	Whether trajectory prediction errors are included in the simulation.	Boolean (On / Off)
Takeoff and Departure Route Uncertainty	Whether uncertainty in takeoff time and departure routing is included.	Boolean (On / Off)
Delayed Slot Strategy	Whether pop-up flights are scheduled based on airport-specific expected delays.	Boolean (On / Off)
Optimized Rescheduling	Whether slot rescheduling is applied when delays or pop-ups occur.	Boolean (On / Off)

The freeze horizon is different than the values of 120Nautical Mile (NM) and 180NM as presented in chapter 4. The AMAN at Luchtverkeersleiding Nederland (LVNL) uses a time-based Freeze Horizon (FH), while the Common Project 1 (CP1) regulations presented in [14] use a distance based FH. The previous analysis of the increase in pop-up traffic is valid for the distances discussed, but the values could differ when considering a time-based FH. The choice to use a time-based FH is made to be able to directly draw conclusions that would apply to Advanced Schiphol Arrival Planner (ASAP).

The current FH is 14 minutes from the estimated time over IAF. An estimation of the time that would represent a FH of 180NM from Schiphol airport resulted in an E-AMAN FH of 20 minutes, from the IAF.

To measure the effect of TP uncertainty, and the effect of takeoff and departure route uncertainty, the errors caused will be introduced in different measurement configurations.

5.2.3. Dependent Variables

The performance of the E-AMAN will in essence be measured using the aircraft state of the simulated aircraft. The aircraft state and thus the flown 4D trajectory can be converted into several dependent variables, together with the commands of the ATC module. These variables are based on section 3.2. The goal is to be able to answer the research questions as presented in chapter 1 by accurately representing E-AMAN performance in terms of stability, accuracy, capacity and taskload.

KPA	Variable	Definition
1 Stability		
1.1	EAT revisions after freeze [-]	Number of times the expected arrival time (EAT) is revised after the freeze horizon per aircraft.
1.2	Sequence position changes after freeze [-]	Number of changes in sequence position of aircraft after the freeze horizon per aircraft.
2 Accuracy		
2.1	Mean Low Level Delay Absorption [s]	Amount of delay is absorbed via speed or route modifications near the TMA, per aircraft.
2.2	Delay energy cost [MJ]	Estimated extra energy consumption due to delay maneuvers per aircraft.
2.3	Count of no LLDA [-]	Number of flights for which no low-level delay absorption was required.
2.4	Delay type overview	Breakdown of delay absorption maneuvers by type: speed adjustments, dogleg route changes, and holding-stack entries.
3 Capacity		
3.1	Mean inter-arrival time [s]	Mean of time intervals between consecutive runway arrivals.
4 Taskload		
4.1	Number of ATC instructions [-]	Total number of ATC instructions generated during the simulation, per aircraft.
4.2	Number of sequence position changes after freeze [-]	Number of sequence changes after freeze horizon, per aircraft.

Table 5.2: Dependent variables split according to Key Performance Areas (KPA)

5.2.4. Controlled Variables

Some variables in the simulation are to be kept constant between configurations. These controlled variables could influence results and are chosen to be representative of the actual situation. They might be of interest in future research. The most significant controlled variable is the traffic scenario, which was previously introduced. Further controlled variables include aircraft performance and speed scheduling, the freeze horizon, the eligibility horizon, update frequency, the approach margin, and all

other simulation settings.

5.2.5. Experiment Matrix

Table 5.3 gives an overview of the different configurations resulting from the independent variables.

Table 5.3: Experiment Matrix of Simulation Configurations

Configuration	Measurement	FH (min)	En-route TP Error	Takeoff & Departure Route Error	Delayed Slot Strategy	Optimized Rescheduling
1	Idealized Reference	14 20	0	0	–	–
2	Realistic Baseline	14 20	✓	✓	–	–
3	Isolated TP Uncertainty	14 20	✓	0	–	–
4	Isolated Takeoff Uncertainty	14 20	0	✓	–	–
5	Delayed Slot Strategy	14 20	✓	✓	✓	–
6	Rescheduling Strategy	14 20	✓	✓	–	✓
7	Combined Delayed Slot + Rescheduling	14 20	✓	✓	✓	✓

5.2.6. Monte-Carlo

To evaluate the impact of uncertainty on E-AMAN performance under different configurations, a Monte-Carlo based approach will be used. This approach introduces stochasticity into the simulation by sampling from probability distributions that reflect real-world prediction errors, as derived in chapter 4. Each Monte-Carlo iteration simulates a complete inbound peak scenario using a unique random seed to generate consistent sets of errors.

Each configuration from the experiment matrix in Table 5.3 will be simulated using the same random seeds and traffic scenarios, to be able to compare configurations. The variables varied across Monte-Carlo iterations include:

- Takeoff time uncertainty, dependent on airport and lookahead time,
- Departure route time uncertainty, dependent on airport,
- Trajectory prediction uncertainty outside of the FIR,
- Trajectory prediction uncertainty within the FIR

The amount of iterations is based on whether the results converge, and statistical tests are able to show any statistical significance between Key Performance Indicator (KPI)'s of different configurations. There is also an upper limit to the amount of iterations, based on the available time each iteration takes.

The results for each configuration, together with the settings, seed and samplings will be stored for reproducibility.

5.2.7. Representation of Uncertainty in Simulation

To simulate realistic aircraft behavior under uncertainty, prediction errors must be modeled and included in the Monte-Carlo simulation. These errors are derived from probability density functions (PDFs), based on the numerical results presented in chapter 4. Each type of uncertainty is represented by a distribution, with parameters derived from historical EFD data, except en-route uncertainty within FIR.

Following the approach of [15], the Johnson SU distribution is selected to represent most of the uncertainty sources due to its ability to model skews and tails. The distributions are fitted using SciPy's Johnson SU method. For the trajectory prediction within the FIR, a standard normal distribution is used instead.

Separate PDFs are constructed for each uncertainty source:

- **Takeoff Error:** A unique Johnson SU distribution is generated per airport and per lookahead time.

- **Departure Route Error:** Each departure airport has its own Johnson SU distribution representing the deviation in time to reach FL180 (or FL100 for Brussels), normalized as a percentage of estimated flight time.
- **Trajectory Prediction (TP) Error – En-route:** A single Johnson SU distribution is used for the prediction error to the Coordination Point (COP), normalized as a percentage of estimated flight time.
- **Trajectory Prediction Error – Within FIR:** This phase uses a normal distribution with a standard deviation.

The distributions are sampled using a scenario-specific seed. This allows for comparison between configurations, and ensures reproducibility. The sampling logic varies by flight phase:

- **Takeoff Error:** For in-horizon departures, a takeoff delay is sampled once at the freeze horizon using the lead time:

  $$T_{\text{lead}} = \text{ETOT} - T_{\text{Freeze}}$$

  The resulting error value (in minutes) is added to the aircraft's simulated off-block time, while its predicted ETOT remains unchanged. Once the aircraft has taken off, this error is reset to zero.
- **Departure Route Error:** A normalized time error is sampled for each aircraft. This error is introduced into the trajectory predictor. Once the aircraft reaches FL180, the error is reset to zero.
- **TP Error – En-route (outside FIR):** A single time error is sampled per flight and applied to the predicted time from COP to IAF. This error linearly decreases in discrete steps as the aircraft approaches the FIR boundary, reaching zero at FIR entry.
- **TP Error – Within FIR:** A time error sampled from a normal distribution is applied to the final phase of the trajectory prediction. Similar to the en-route error, it linearly decreases to zero at the IAF.

All errors introduced in the TP plugin. The errors modify the predicted arrival times used by the AMAN module.

5.3. Simulation Setup

To investigate the questions presented previously, the AMAN and E-AMAN are going to be simulated, with BlueSky as a basis. This air traffic simulator has been chosen since it provides an easily adaptable platform, and of which the use is well known by the researcher. Furthermore, most previous simulation research into the E-AMAN has been conducted using the BlueSky simulator, which would possibly allow for the re-use of certain modules or ideas. This in turn would simplify development. The AMAN simulator will consist of different modules, presented below. The modules are quite similar to the modules of the actual AMAN.

In general, the simulator development shall be conducted such that the developed product can be easily adapted to future research. Each module is to be implemented as a plugin, while simple commands in the BlueSky scenario can easily adapt the AMAN settings, such as freeze horizon or approach margin.

5.3.1. Trajectory Predictor

The TP is the basis for the AMAN, as described in chapter 2. There would be two suitable options for the trajectory predictor implementation in BlueSky: one previously used in [19] and [13] specifically for a Boeing 747, or a newly developed deterministic TP using BlueSky as its' own TP [39].

Due to the nature of the deterministic option being available without much further development, the compatibility with any aircraft, and the possibility to introduce a controllable error for each prediction, the option presented in [39] is chosen. Another reason for choosing this option is that a beginning has been made in the development of the AMAN using this TP. This again simplifies development and allows for increased fidelity.

The TP consists of a second BlueSky instance, or "node". The main BlueSky instance runs a scenario file, and sends the scenario commands to the node. The node then returns flighttimes to relevant

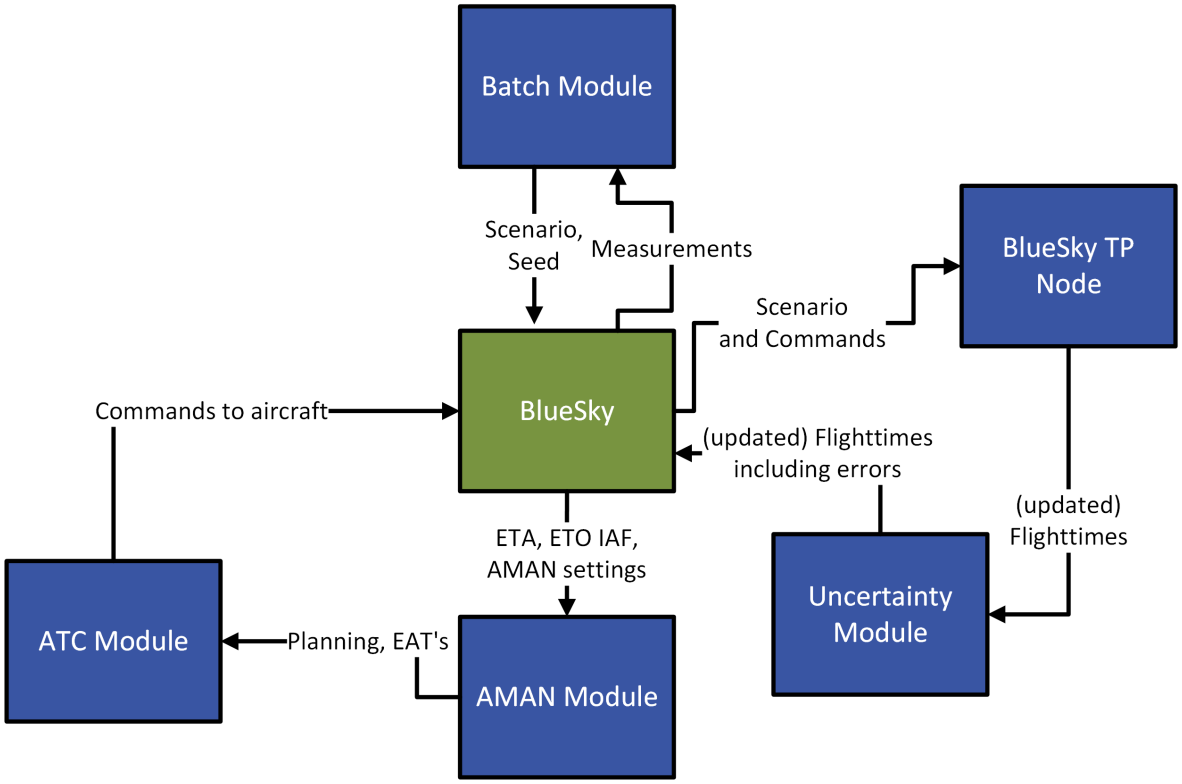

Figure 5.2: AMAN simulation visualization

waypoints per aircraft, which get processed in the main node. If a command is given to an aircraft in the main BlueSky instance, the aircraft is created again in the node, and the flighttime updated.

5.3.2. AMAN Module

The AMAN module will be based on the work presented in [39], and is modeled after ASAP as presented in chapter 2. The AMAN must create a sequence based on the predicted arrival times at the runway, in combination with ensuring separation at the runway by generating slottimes in combination with Time To Gain (TTG) or Time To Lose (TTL). The module must communicate this to the ATC module or the operator. The AMAN module must also be able to assign an expected delay to a pre-planned pop-up aircraft, and use this value in its' planning. Furthermore, the AMAN module must have a freeze horizon and an eligibility horizon, which can be changed through a setting or scenario. Finally, the AMAN module must be able to accommodate the actions a controller can take to optimize the planning.

The AMAN planning will be based on a dataframe, since it is easily modified, can be shown graphically easily, and filtered easily. Another option would be to use a singleton, similar to the traffic object in BlueSky. This option would make the simulation run faster, but seems more complicated to develop.

5.3.3. ATC Module

The ATC module should accurately simulate the actions of an air traffic controller working with AMAN. As can be read in subsection 3.5.5, previous research into the E-AMAN recommend that a closer look is taken at the actions available to the controller using the E-AMAN. The previously unavailable actions to the simulated controller are front-loading and ad-hoc runway changes. These actions are not solutions to pop-ups, but tools available to increase efficiency. There might be an effect on the impact of pop-up traffic since they increase the flexibility in the planning.

The ATC module must optimize the sequence in a similar way as the approach planner. Primarily by swapping aircraft, front loading or swapping runways. It must also be able to appropriately deal with revising the planning in case of pop-up aircraft.

Furthermore, the ATC module must be able to instruct incoming aircraft to absorb delay. This will either be done by increasing or decreasing aircraft speed, adding or removing track miles, or assigning the aircraft to a holding position. The ATC module must be able to convert TTG and TTL into these instructions per aircraft. In previous research, a similar module was developed, called Speed and Route Advisor (SARA).

Required ATC actions for AMAN operation:

- Reduce speed instruction
- Reduce mach instruction
- Dogleg instruction
- Holding instruction
- Direct instruction
- Speed up instruction
- Swap slots
- Front-load
- Swap runway

5.3.4. Batch

Since a BlueSky instance takes up a single thread, the time efficiency of running the simulation can be increased if more threads can be used. To ensure a relatively short run time in combination with the Monte-Carlo simulations, one or multiple computers with 16 CPU cores are to be used. This allows for at least 8 instances of BlueSky with 8 TP nodes. During initial research, the batch functionality of BlueSky was inoperative, causing development to be required. The choice has been made to redevelop the functionality, with the TP plugin as a basis. This new plugin will also be used to control the entire simulation, such as setting the scenario, the seed, and controlling the data storage.

To speed up start-up of each scenario, a cache can be made to store predictions for each scenario. If the scenario is equal, the initial TP prediction is exactly the same. If the TP is used, the communication of the predicted times can take up to four minutes for a large scenario.

5.3.5. Prediction Errors

A plugin is to be developed to calculate errors for the prediction in each flight phase. The inputs will be the Probability Density Function (PDF)'s generated using the data examined in chapter 4. The plugin will generate errors for all phases of relevant simulated flights, in a reproducible manner. The plugin will also record all generated values to enhance reproducibility. The error will not change between different E-AMAN configurations to be able to directly compare the different configurations.

5.3.6. Scenario Generation

The scenarios will be generated using data from the Eurocontrol DDR [37], specifically the filed trajectories. In previous research this data source has been used, and makes realistic scenarios possible. Per aircraft, the flown trajectories will be processed into routepoints in a BlueSky scenario. From the COP towards the IAF, the routepoints will either be sourced from the same DDR source, or from the standard LVNL routes depending on suitability of the data. From the IAF, standard routes from LVNL will be used. These fixed routes in the TMA will have no effects on the performance of the AMAN. Aircraft speed scheduling will be based on the ASAP performance numbers, while aircraft performance will be based on OpenAP or BADA. The scenario generation is done using code based on the code used by [19] and [13].

5.3.7. General Assumptions and Limitations

To simplify development, several assumptions have been made. These assumptions either have negligible impact on AMAN performance or would require major development effort to model accurately.

1. No departures are simulated from Schiphol, as arrival and departure flows are largely separated outside the TMA.
2. The arrival sequence is not de-conflicted.
3. Within the TMA, fixed routes are used instead of vectoring. Vectoring is assumed to have negligible influence on E-AMAN performance.

4. Only wind is modeled; other weather effects are excluded due to their relatively minor impact.
5. Aircraft departing from Dutch airfields are not considered as pop-up aircraft due to pre-takeoff coordination.
6. The sequence remains First-Come-First-Served (FCFS), except for pop-up aircraft.
7. The approach margin is fixed at 2 minutes, but can be modified for future research.
8. Holding is assumed to occur near published holding patterns [4].
9. DDR updates occur only when flights deviate by more than 5 minutes. This is assumed to have limited influence on scenario accuracy.
10. Radar range is assumed to cover areas beyond the freeze horizon, in line with future System Wide Information Management (SWIM)-enabled surveillance.
11. LVNL receives a message at actual takeoff time.
12. All aircraft are assumed to file a flight plan at least 60 minutes prior to takeoff.

Two major limitations are present in the simulation: the only goal of Air Traffic Control (ATC) instructions is to adhere to planning, with no tactical or conflict resolution instructions being made. The effects of the different configurations on these controller tasks are unknown, and have an effect on the E-AMAN performance. These secondary effects are very difficult to model without a model of a controller, or a human-in-the-loop experiment. Though, the research goals are such that these options are out of scope, and would not answer the research questions properly. Depending on the outcome of the current research, a human-in-the-loop experiment could be recommended for further research. The other limitation is mainly set in the accuracy and variation in the scenarios. The measurements are only valid for the simulated scenarios. The scenarios inevitably take shortcuts, simplifications, and do not represent all traffic scenarios accurately. Though, the goal of this research is to investigate the general impact of uncertainty and mitigation strategies on the performance of the E-AMAN. It is not the goal to provide the impact on the performance in all scenarios.

5.4. Verification and Validation

Verification and validation will be done throughout the development of the simulation, to ensure that the results are accurate and representative.

The verification focuses on the accuracy of each functionality. The verification begins with testing if each added function to each module reproduces the expected result in each type of possible input. Simple test scenarios will be used to check the behavior of each function. BlueSky has the benefit of visually representing the aircraft behavior, which greatly simplifies this process. Furthermore, the functionality that has influence on the planning can also easily be verified using the AMAN planning.

- **Module** Each newly developed function is tested against expected outputs using simplified input cases. These unit tests verify correctness for the expected range of inputs.
- **Visual behavior checks:** BlueSky allows for verification of aircraft behavior due to instructions being given by any of the modules.
- **AMAN planning verification:** The planning accuracy and freeze horizon handling is directly observable via the AMAN planning output.

The validation focuses on two areas. First, the correct behavior of the complete simulation and individual modules.

- **Trajectory Predictor:** Prediction accuracy is checked by comparing predicted versus actual arrival times.
- **ATC Module:** The correctness of ATC instructions is verified by checking adherence to Estimated Approach Time (EAT).
- **AMAN Module:** Validated by comparing the output sequences with expected First-Come-First-Serve (FCFS) baselines and by testing the effects of freeze horizon logic and pop-up handling.
- **Batch Module:** Ensures consistent outputs when using identical seeds across repeated runs.

- **Uncertainty Module:** Validation by plotting sampled errors and comparing their distribution to the PDFs derived from EFD data.
- **Complete Simulation:** Validation by comparing simulation behavior with the expected behavior.

The second focus is the validation of the results. This will be done using a statistical analysis. The goal is to test whether the differences in mean values of dependent variables are statistically significant between configurations, and thus whether observed performance differences are due to real effects rather than random variation. A further look into which statistical test(s) is/are relevant to the Monte-Carlo simulation is taken later.

5.5. Current Development Status

The current status of the development is that many functionalities have been implemented. The TP module is complete, with the AMAN module providing the planning and update functionalities. The essential functionalities of the ATC module are also complete, while the batch module must only be validated. The main areas of development lie in the uncertainty module, the creation of the final scenarios, and the data processing pipeline. Furthermore, the functionality of changing the planning produced by the AMAN module must still be implemented, to be able to represent the controller actions, but also the mitigation strategies.

It is expected that at the beginning of September 2025, the development of the simulation is complete, and the first configurations can be tested.

References

[1] Jacco M. Hoekstra and Joost Ellerbroek. “BlueSky ATC Simulator Project: An Open Data and Open Source Approach”. en. In: *Proceedings of the 7th International Conference on Research in Air Transportation (ICRAT)*. Paper accepted at ICRAT 2016; TU Delft Research Portal. Philadelphia, USA, 2016, pp. 131–138. URL: https://resolver.tudelft.nl/uuid:d1131a90-f0ea-4489-a217-ad29987689a1 (visited on 07/24/2025).

[2] Jacco Hoekstra and Joost Ellerbroek. “AE4321 Air Traffic Management Human Machine Interface”. en. Lecture slides for AE4321, TU Delft. Nov. 2023.

[3] EUROCONTROL Maastricht Upper Area Control Centre (MUAC). *Cross-border air navigation services*. en. Technical Report / Web Publication. Overview of cross-border civil□military air navigation services by MUAC, Europe’s sole cross□border ANSP. EUROCONTROL, Sept. 2024. URL: https://www.eurocontrol.int/muac (visited on 09/09/2024).

[4] Air Traffic Control the Netherlands (LVNL). *Aeronautical Information Publication (AIP) – The Netherlands*. en. Official eAIP issued by AIS Netherlands (LVNL), updated per AIRAC cycle. Jan. 2025. URL: https://www.lvnl.nl/diensten/aip.

[5] Hans Koolen and Ioana Coliban. *Flight Progress Messages Document*. en. Technical Report. Version 2.600. Edition 2.600 dated 1 July 2020; detailed description of ETFMS flight progress messages. EUROCONTROL Network Manager, July 2020. URL: https://www.eurocontrol.int/sites/default/files/2020-06/flight-progress-msg-update-230620.pdf (visited on 02/01/2024).

[6] EUROCONTROL. *EUROCONTROL Specification for Airport Collaborative Decision Making (A□CDM)*. en. Specification. Draft Edition 1.0 (05 Dec 2024); supports harmonised A□CDM implementation at European airports. EUROCONTROL, July 2024. URL: https://www.eurocontrol.int/sites/default/files/2024-07/eurocontrol-draft-specification-acdm-ed1-0.pdf (visited on 09/09/2024).

[7] EUROCONTROL. *Airport Collaborative Decision□Making (A□CDM)*. en. Published 31 Mar 2017; guides harmonised implementation of A□CDM at European airports. Mar. 2017. URL: https://www.eurocontrol.int/concept/airport-collaborative-decision-making (visited on 09/09/2024).

[8] Hans Koolen and Ioana Suciu. *DPI Implementation Guide*. en. Technical Report. Version 2.700. Edition 2.600 issued 1 Oct 2023; authors Hans Koolen and Ioana Suciu. EUROCONTROL, Oct. 2023. URL: https://www.eurocontrol.int/publication/departure-planning-information-dpi-implementation-guide (visited on 02/01/2024).

[9] EUROCONTROL. *Arrival Manager Implementation Guidelines and Lessons Learned*. en. URL: https://skybrary.aero/sites/default/files/bookshelf/2416.pdf (visited on 03/29/2024).

[10] Luchtverkeersleiding Nederland (LVNL). *Advanced Schiphol Arrival Planner (ASAP) – Ontwerpdocument: Procedures*. nl. Technical Report. Version 2.0; internal design document for AMAN (Arrival Manager) implementation. LVNL – Luchtverkeersleiding Nederland, Mar. 2019.

[11] Anouk Hollebeek. “Arrival Management Evaluation: An insight into Advanced Schiphol Arrival Planner”. en. Graduated spring 2019; collaboration between Aviation Academy, Amsterdam University of Applied Sciences and TU Delft. BSc thesis. KDC Mainport Schiphol – Centre of Excellence (AUAS & TU Delft), 2019.

[12] *RECAT-EU: European Wake Turbulence Categorisation and Separation Minima on Approach and Departure*. Tech. rep. Released Issue, Edition 2.0, 08 November 2024. EUROCONTROL, Nov. 2024. URL: https://www.eurocontrol.int/publication/recat-eu-wake-turbulence-categorisation-and-separation-minima.

[13] M. Van Horssen. “Cross-border Arrival Management to Reduce Traffic Bunching at Schiphol Airport”. en. TU Delft Master Thesis in Aerospace Engineering. MSc thesis. Delft University of Technology, July 2017. URL: https://repository.tudelft.nl/islandora/object/uuid:bf30d6f8-4ac3-4cbe-9869-4fe923118e6f (visited on 01/11/2024).

[14] European Commission. *CP1 Regulation: Commission Implementing Regulation (EU) 2021/116 on the establishment of the Common Project One*. en. Commission Implementing Regulation. Entered into force 22 Feb 2021; establishes Common Project One (CP1) supporting the ATM Master Plan. Official Journal of the European Union, Feb. 2021, pp. 10–38. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32021R0116 (visited on 09/09/2024).

[15] Edzer Oosterhof. “Effect of Trajectory Prediction Uncertainty on a Probabilistic Debunching Concept for Inbound Air Traffic”. en. MSc thesis. Delft University of Technology, Dec. 2022. URL: https://repository.tudelft.nl/record/uuid:6556ed83-f28c-40a1-baa8-a46ddda804db (visited on 07/24/2025).

[16] Onbekend. *QRC8 Baancombinaties en LIV – Capaciteit in ASAP (APLN)*. nl. Technical Report. QRC8-rapport over baancombinaties en LIV-capaciteit in het ASAP□systeem (APLN). Interne rapportage, 2024.

[17] EUROCONTROL. *Supporting Material to the SDP Implementation Update 2024.1*. en. Technical Report. Supporting material on CP1 implementation, including AMAN horizon extension. EUROCONTROL – SESAR Deployment Manager, Jan. 2024. URL: https://www.eurocontrol.int/sites/default/files/2024-01/sdp-implementation-update-2024-1-supporting-material.pdf.

[18] EUROCAE / EUROCONTROL. *ED□254: Arrival Sequence Service Performance Standard*. en. Performance Specification. Available as SWIM service specification; current version released June 2018. EUROCAE (EUROCONTROL), June 2018. URL: https://eshop.eurocae.net/eurocae-documents-and-reports/ed-254/.

[19] A. Vanwelsenaere. “Effect of Pop-Up Flights on Extended Arrival Manager”. en. TU Delft Master Thesis in Aerospace Engineering. MSc thesis. Delft University of Technology, July 2016. URL: https://repository.tudelft.nl/islandora/object/uuid:fe4cb9c1-59a9-45a5-853d-1da2dccc74a1 (visited on 01/11/2024).

[20] Maarten Tielrooij. “Arrival Management Support in the Presence of Prediction Uncertainty”. en. Defended on 21 October 2022. PhD thesis. Delft University of Technology, Oct. 2022. URL: https://repository.tudelft.nl/record/uuid:a403112b-48a2-40d7-9db3-a5b754f31eee (visited on 07/24/2025).

[21] Nelson M. Guerreiro and Matthew C. Underwood. “Understanding Extended Projected Profile (EPP) Trajectory Error Using a Medium-Fidelity Aircraft Simulation”. en. In: *2018 Aviation Technology, Integration, and Operations Conference*. 2018 AIAA Aviation Forum. Atlanta, Georgia: American Institute of Aeronautics and Astronautics, June 25, 2018. ISBN: 978-1-62410-556-2. DOI: 10.2514/6.2018-3044. URL: https://arc.aiaa.org/doi/10.2514/6.2018-3044 (visited on 01/30/2024).

[22] R. A. Vos. “Machine Learning Based Trajectory Prediction to Support Demand Forecasting”. en. TU Delft Master Thesis in Aerospace Engineering. MSc thesis. Delft University of Technology, July 2023. URL: https://repository.tudelft.nl/record/uuid%3A0c7dc66f-84c3-442b-a77e-41eb17da68ee? (visited on 07/24/2025).

[23] Ferdinand Dijkstra. “Expert Interview on Traffic Scenarios”. Personal communication during expert interview, conducted by the author. June 2025.

[24] Ron Keizer. *Expert Interview with Approach Supervisor*. Personal communication. Nov. 2024.

[25] H. J. Hoogendoorn. “Dynamic Airline Centric Inbound Priority Sequencing”. en. Graduation date 13 Dec 2022; dynamic IPS model for KLM at Schiphol. MSc thesis, Air Transport and Operations. Delft University of Technology, Dec. 2022. URL: https://resolver.tudelft.nl/uuid:4c08bbcf-21c9-4074-81bf-b66099b67734 (visited on 07/24/2025).

[26] Songwei Liu et al. “Identifying Operational Benefits of the Arrival Management System – A KPI-Based Experimental Method by Evaluating Radar Trajectories”. en. In: *Promet – Traffic & Transportation* 33.5 (Oct. 2021), pp. 633–645. ISSN: 0353-5320. DOI: 10.7307/ptt.v33i5.3786. URL: https://traffic.fpz.hr/index.php/PROMTT/article/view/3786 (visited on 03/21/2024).

[27] Zihan Peng et al. “Benefits Derived from Arrival Management and Wake Turbulence Re-Categorization in China”. en. In: *Transportation Research Record: Journal of the Transportation Research Board* 2675.11 (Nov. 2021), pp. 373–383. ISSN: 0361-1981. DOI: 10.1177/03611981211017898. URL: https://journals.sagepub.com/doi/10.1177/03611981211017898 (visited on 03/22/2024).

[28] Jane Thipphavong and Steven Landry. “Effects of the Uncertainty of Departures on Multi-Center Traffic Management Advisor Scheduling”. en. In: *AIAA 5th ATIO and 16th Lighter-Than-Air Systems and Balloon Systems Conferences*. AIAA 5th ATIO and 16th Lighter-Than-Air Systems and Balloon Systems Conferences. Arlington, Virginia: American Institute of Aeronautics and Astronautics, Sept. 2005. ISBN: 978-1-62410-067-3. DOI: 10.2514/6.2005-7301. URL: https://doi.org/10.2514/6.2005-7301 (visited on 03/21/2024).

[29] Katsuhiro Sekine et al. “Validating Flow-Based Arrival Management for En Route Airspace: Human-In-The-Loop Simulation Experiment with ESCAPE Light Simulator”. In: *Aerospace* 11.11 (2024). Open Access under CC BY 4.0 license, p. 866. ISSN: 2226-4310. DOI: 10.3390/aerospace11110866. URL: https://www.mdpi.com/2226-4310/11/11/866.

[30] NATS. *XMAN: Cross-Border Arrival Management*. Tech. rep. Accessed: 2025-08-03. NATS - UK Air Traffic Control, 2017. URL: https://www.nats.aero/news/arrival-manager-improvements-benefit-arrivals-london-heathrow-airport/.

[31] NATS and Harris Orthogon GmbH. *NATS Arrival Manager powered by Harris Orthogon*. https://www.nats.aero/news/new-streamlined-heathrow-arrivals-start-europe/. Accessed: 2025-08-03. 2017.

[32] M. Vella and X. Olive. “Data□Driven Prediction of Aircraft Holding Times Using ...” In: *Journal of Air Transportation Systems* (2024). States that 36 % of arrivals spend time in stacks at Heathrow. URL: https://journals.open.tudelft.nl/joas/article/view/7890.

[33] NATS. *Heathrow holding times on the decline thanks to new technology*. Tech. rep. Reports reductions from 8.5 min to 6–7.5 min average hold per flight. 2016. URL: https://www.nats.aero/news/heathrow-holding-times-on-the-decline-thanks-to-new-technology/.

[34] L.M.C. Adriaens. “MUAC En Route Delay Absorption Capabilities for Schiphol Inbounds”. en. Control & Operations – ATM, Airports & Safety (Aerospace Engineering). MSc thesis. Delft University of Technology, Aug. 2015. URL: http://resolver.tudelft.nl/uuid:6c752803-6325-4ed9-8b8c-ec52e9fae5db.

[35] Sofia Rydell. “Arrival and Departure Manager Cooperation for Reducing Airborne Holding Times at Destination Airports”. en. Available via Cranfield University repository (DSpace). MSc thesis. Cranfield University, Aug. 2011. URL: http://dspace.lib.cranfield.ac.uk/handle/1826/8022 (visited on 07/24/2025).

[36] SKYbrary. *Calculated Take-Off Time (CTOT)*. SKYbrary Aviation Safety. 2023. URL: https://skybrary.aero/articles/calculated-take-time-ctot (visited on 07/24/2025).

[37] EUROCONTROL. *DDR2 User Guide*. en. User Guide. Generic user guide for the DDR2 Demand Data Repository service. EUROCONTROL Network Operations Library, 2022. (Visited on 08/02/2024).

[38] Bewoners Aanspreekpunt Schiphol (BAS). *BAS 2024 Jaarreportage*. nl. Tech. rep. Bewoners Aanspreekpunt Schiphol (BAS), Royal Schiphol Group, Mar. 2025. URL: https://bezoekbas.nl/nieuws/jaarrapportage-2024-beschikbaar/ (visited on 03/26/2025).

[39] Kubba Ahmed. *README: Trajectory Predictor & AMAN Plugins*. en. Internal report. README file describing Trajectory Predictor and AMAN plugin functionality. Internal Documentation, LVNL / KDC Mainport Schiphol, 2024.