Data-Driven Cruise-Speed Control Balancing Demand and Capacity for Transatlantic Arrivals at Schiphol under Uncertainties

by

Ahmed Kubba

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

Thesis committee:	
Supervisors:	Dr. Ir. Joost Ellerbroek
	Ferdinand Dijkstra
Chair:	Dr. Ir. O.K. Bergsma
External examiner:	Dr. Junzi Sun

Faculty:	Aerospace Engineering
Track:	Control & Operations
Profile:	Control & Simulation

TU Delft Delft University of Technology

Data-Driven Cruise-Speed Control Balancing Demand and Capacity for Transatlantic Arrivals at Schiphol under Uncertainties

Ahmed Kubba[1,2], under supervision of: Dr. ir. Joost Ellerbroek[1] and Ferdinand Dijkstra[2]
[1] *Control & Operations, Aerospace Engineering, Delft University of Technology, The Netherlands*
[2] *KDC Mainport Schiphol, The Netherlands*

European hub airports face persistent demand-capacity imbalances generating 22.4 million minutes of delays in 2024, costing €2.8 billion. Traditional reactive approaches prove inefficient for long-haul transatlantic flights. Moreover, long-haul flights have historically been granted exemptions from ground delay programs, meaning short-haul flights disproportionately absorb delays during capacity constraints. This research develops an uncertainty-aware cruise speed control strategy for managing transatlantic arrivals to Amsterdam Schiphol Airport. The methodology decomposes arrival time uncertainty using Johnson distributions conditioned on flight status and temporal horizon, continuously monitoring probability that demand at Initial Approach Fixes exceeds capacity. One-time strategic speed adjustments trigger only when exceedance probability surpasses a predetermined threshold of 70 %. Fast-time simulation of 30 days using BlueSky and EUROCONTROL trajectory data validates the approach. Across 79 interventions, the strategy achieved 35 % capacity exceedance certainty reductions on average. These required only 8.72 minutes average delay through conservative Mach reductions of 0.01-0.05. The approach offers a practical, environmentally beneficial pathway for long-range air traffic flow management at capacity-constrained hubs.

Keywords: Air traffic flow management; Uncertainty quantification; Arrival management; Speed control; Probabilistic forecasting; Extended AMAN; XMAN; Long-range ATFM; LR-ATFM; Demand-capacity balancing; Johnson distributions; Transatlantic flights

I. Introduction

Major hub airports frequently experience demand peaks **that exceed** their available capacity. In 2024, European air traffic suffered its worst delays in decades: en-route Air Traffic Flow Management (ATFM) delays reached 22.4 million minutes and the average delay per flight climbed to 2.13 minutes. Insufficient air-traffic-control capacity generated nearly 39% of all en-route ATFM delays, closely followed by adverse weather. The ongoing shortage of controllers, which is driven by slow hiring and delayed capacity projects, cost airspace users about €2.8 billion that year.[1]

These shortcomings have motivated demand-smoothing initiatives that shift delay from terminal stacks to cruise. Extending the horizon of an Arrival Manager (AMAN) beyond the terminal area, also known as Extended AMAN (XMAN), which allows delays to be absorbed en route, reducing holding, noise, and fuel burn [2]. At London Heathrow, cross-border XMAN trials showed that slowing aircraft up to 350 NM from destination can cut low-altitude holding and save thousands of tonnes of fuel and tens of thousands of tonnes of CO_2 annually [3]. Such results underline the environmental gains achievable by strategically managing arrival flows on a larger horizon.

Traditionally, air navigation service providers (ANSPs) manage demand-capacity imbalances through ground delay programs (GDPs) and near-terminal sequencing, but long-haul flights have unique challenges. Long-haul flights cross several Flight Information Regions, each with its own rules requiring coordination across different ANSPs and ATFM units. Ground delays are applied reactively, and standard Estimated Time-of-Arrival (ETA) models increase the problem by ignoring key variables such as wind patterns and airport congestion [4].

Newer long-range ATFM (LR-ATFM) concepts intervene much earlier in flights, e.g. ICAO's Global Trajectory-Based Operations promotes cross-border data sharing and collaborative decisions to lengthen planning horizons [5]. While promising, these approaches can be complex to implement, requiring high predictability and extensive coordination. In practice, the uncertainties in long-horizon planning remain a significant obstacle. Over intercontinental distances, wind and weather forecast errors can significantly shift a flight's arrival time, undermining the accuracy of demand predictions. If not addressed, these uncertainties can reduce the effectiveness of any LR-ATFM tool, and if too conservative capacity might be wasted and missed interventions lead to residual congestion. Similarly, the departure timing of short-haul flights can vary enough that a hub airport's mix of long- and short-haul traffic becomes hard to schedule reliably multiple hours in advance. Moreover, transatlantic flights show significantly lower pop-up flight occurrence rates because of their operational characteristics [6]. This enhances predictability for ATFM systems. Most existing ATFM models still ignore stochastic demand–capacity effects, highlighting the need for uncertainty-aware strategies.

In light of these challenges, this paper proposes a unique approach to proactively redistribute ETAs for transatlantic flights bound for Amsterdam international airport (Schiphol), using uncertainty-aware speed control during the cruise phase. The core idea is to continuously predict the arrival demand in comparison to the available capacity at key waypoints and to issue a cruise speed adjustment only once for each flight if the forecasted demand is projected to exceed the capacity by a significant probabilistic margin. Unlike more elaborate multi-fix metering schemes [5], a single well-timed speed adjustment during cruise is far more practical than multiple changes that force the aircraft to alternately accelerate and decelerate. Furthermore, when applied early enough, it can absorb several minutes of delay. Additionally, by defining a demand-capacity threshold taking into account the uncertainties, it is ensured that speed advisories are triggered only when necessary i.e., when there is high confidence of capacity overload.

The novelty of this approach lies in combining probabilistic traffic forecasting with a threshold-based control law to initiate minimal interventions. As far as we are aware, prior LR-ATFM concepts have not employed an explicit demand-capacity probability threshold as the trigger for speed adjustments; this provides a new way to cope with forecast uncertainty to slow down flights. Furthermore, a proactive speed control strategy has clear environmental incentives alongside its operational goals [3]. To evaluate the proposed speed control strategy, we simulate transatlantic arrivals into Schiphol using the BlueSky open-source air traffic simulator [7].

This paper is outlined as follows: Section II reviews the theoretical background on which the proposed method is built with respect to recent literature; Section III details the concept; Section IV introduces the data set, scenario, and validation; Section V presents the numerical results; and Sections VII and VIII conclude with recommendations for further research.

II. Background

The ICAO Global Trajectory-Based Operations framework advocates cross-border information sharing and collaborative decision-making to extend ATFM planning horizons [5]. While Europe's centralized Network Manager issues Calculated Take-Off Times (CTOTs) to meter traffic, other regions rely on distributed or bilateral ATFM agreements, and globally there is no homogeneous ATFM system [8]. This can result in cross-border flow issues where upstream flights are not managed optimally for downstream congestion. Moreover, long-haul flights have historically been granted exemptions from ground delay programs, meaning short-haul flights disproportionately absorb delays during capacity constraints. [5] Such policies, intended to avoid penalizing flights already airborne, result in inequitable outcomes where arriving banks of long-hauls still overwhelm the runway capacity. Another limitation is the reactive nature of many ATFM measures for arrivals, interventions often occur only close to the destination (e.g. airborne holding or vectoring near initial approach fixes (IAFs)) rather than earlier in cruise. For instance, Schiphol's inbound flows are funneled via a few key IAFs, such as SUGOL and RIVER, where aircraft may be placed in holding patterns during peak periods. [9] Managing delay at these late stages is fuel-inefficient and contributes to congestion in the terminal airspace.

A. Long-Range Air Traffic Flow Management

LR-ATFM concepts under development aim to address these gaps by acting far upstream in the flight. Recent research has proposed, for example, setting target times over common waypoints hours before the arrival, effectively metering inbound flows well in advance [2]. Schultz et al. [8] introduced a concept of operations in the Asia-Pacific region in which speed adjustments as far as approximately 7 hours from landing were used to smooth arrival streams into Singapore. Enea and Bronsvoort [5] have explored using multiple sequential metering points along an aircraft's route, enabled by advanced Flight Management System capabilities (e.g. Required Times of Arrival at intermediate fixes) to gradually absorb delay. While promising, these approaches can be complex to implement, requiring high predictability and extensive coordination. In practice, the inherent uncertainties in long-horizon planning remain a significant hurdle.

B. Extended Arrival Manager

Traditional AMAN systems sequence inbound flights so that demand matches runway capacity, often through tactical vectoring or holding near the airport. Extending the AMAN horizon further upstream enables minor cruise-speed adjustments to absorb delay en route [10]. Trials at London, Rome, and Amsterdam reported up to 8%

fuel savings per flight and a 90% reduction in airborne holding when arrivals were pre-sequenced earlier [11]. Furthermore, Jones et al. [12] demonstrated that assigning CTAs ~150 NM from the airport (for flights up to 500 NM out) could shift about 20% of downstream delays into the en-route phase. However, their model assumed perfect information and did not account for uncertainties like weather or downstream congestion, highlighting the need to incorporate uncertainty in demand-capacity forecasts.

C. Uncertainties in Arrival Times and Demand and Capacity Balancing

Central to the success of any XMAN or LR-ATFM is the ability to anticipate traffic demand and available capacity with reasonable accuracy. This necessitates an explicit quantification of uncertainty. Recent advances in uncertainty quantification demonstrate potential for probabilistic trajectory management, yet existing methods focus primarily on continental airspace.[13]

Uncertainty quantification is critical when balancing demand and capacity in arrival management. Because predicted arrival times can deviate significantly from actual times especially when look-ahead horizons are extended, scheduling decisions made too early can be wrong leading to inefficient or even unsafe outcomes. Prior studies have documented that the earlier an arrival time is predicted, the larger the potential error, due to unknowns in winds, downstream ATC path stretching, and even departure airport delays propagating forward [10]. For instance, extending an AMAN horizon beyond ~100–180 NM means predictions might need to account for taxi-out delays or en route reroutes that are not yet realized [14]. If such uncertainties are ignored, an arrival sequence plan could be disrupted when flights do not arrive in the expected order, forcing last-minute interventions that negate the benefits of early planning [11].

Tielrooij et al. [15] established a method for predicting Estimated Time Over waypoint (ETO) uncertainty at an extended horizon by fitting Johnson curves to derive errors and confidence intervals. The literature consistently advocates for using uncertainty as information for decision-making rather than trying to eliminate it.

Oosterhof et al. [16] demonstrated a probabilistic debunching study, which introduced the idea of only intervening when there is a significant probability of bunching at an IAF. In Oosterhof's work, probability density functions were fitted to ETA prediction errors from historical data, allowing computation of a "bunching probability" that two or more aircraft will violate a separation threshold at the fix. A genetic algorithm then chose speed/delay interventions to reduce that probability. Notably, the results showed that acting too early, when uncertainty was still high, could impose unnecessary delay – before about 40 minutes from the IAF, the interventions didn't consistently reduce arrival holding because the predictions weren't reliable enough as the error was on the same order as the intended separation.

D. Cruise Speed Control

Using cruise speed adjustments to manage arrival flows is an approach well-grounded in prior research and trials. In constrained airspace, traditional flow management often resorts to ground delays or airborne holding stacks to prevent overloads. Cruise speed control offers a fuel-efficient alternative: by reducing speed early, an aircraft can absorb delay while still en route, avoiding excessive vectoring or holding near the airport. Delgado and Prats [17] found that flying about 5–12% slower than nominal cruise speed stays in the regime of equal or lower fuel burn, essentially exploiting the flat part of the drag curve to save fuel while delaying arrival. Matsuno and Andreeva-Mori [18] similarly quantified that a modest 3% speed reduction could absorb roughly 2–3 minutes of flight time per half hour of cruise, with a 2–3% fuel savings as a side benefit. Schultz et al. [8] found that a shift of ±1 minute per remaining flight hour results in a reduction of 14% of the overshoot of continuous capacity due to long distance overhaul, shifting flights by ±2 and ±6 minutes reduces this by 20% and 80% in comparison to the original scenario, respectively.

In Europe's SESAR trials, controllers upstream of Schiphol and other hubs have been provided with speed advisories (via AMAN or network management tools) to delay certain flights by a minute or two, allowing them to seamlessly slot into the arrival sequence without downstream holding. This kind of cooperative speed management requires coordination across FIR boundaries and has been formalized in concepts like XMAN . [11]

A speed control strategy carries clear environmental incentives alongside its operational goals. First, absorbing delay by flying slightly slower at high altitude is far more fuel-efficient than absorbing it in low-altitude holding patterns or stop-and-go vectors near the airport. Jets cruising at their optimal or a reduced speed burn less fuel per minute than when in stacks; thus, delay absorption translates into fuel savings. Studies of XMAN have documented substantial fuel and emissions reductions when delay is transferred from holding stacks to the en-route phase. In the SESAR XMAN trials, for example, participating flights averaged a 48-second delay absorption in cruise with no increase in total flight time, yielding an estimated annual savings of ~4,700 tonnes of fuel and ~15,000 tonnes of CO_2 at a single airport. [3]

Finally, these findings underpin the idea that "strategic deceleration" is often more cost-effective than path stretching or holding. The extended metering concept tested by NASA and others showed many flights can absorb assigned delays entirely through cruise speed control if given early notice – in one study, over 70% of arrival flights were able to meet their delay absorption targets by slowing down, eliminating the need for vectoring in those cases. [19]

III. Concept

Building on sections I and II, this section presents the proposed uncertainty-aware cruise speed control concept for transatlantic arrivals into Schiphol. The concept combines long-range ATFM and Extended AMAN principles by shifting part of the delay-absorption task to the cruise phase, while taking into account the uncertainty in long-horizon arrival-time predictions. In line with the problem formulation in the Introduction, the core idea is to use probabilistic forecasts of arrival demand relative to available capacity and to issue at most one small speed reduction per flight, only when the forecast indicates a sufficiently high probability of demand exceeding capacity.

Based on this concept, the **hypothesis** of this study is that an en-route delay absorption under uncertainty-aware cruise speed control which issues at most one Mach reduction per transatlantic flight, applied several hours before arrival, reduces the probability of capacity exceedance at least below the speed control threshold while inducing only limited delay.

To implement this concept, the study uses a historical dataset of flights inbound to Schiphol from Eurocontrol Network Manager B2B messages. The data provides 4D trajectories consisting of waypoint sequences with associated estimated times over (ETOs) and flight metadata, which includes identifiers, aircraft types, and operational status. As done in previous research [15], the final B2B trajectory message for each flight, which is the last update of a flight, serves as the actual flight plan used for simulation.

Concept overview

The study proceeds through seven consecutive steps, each elaborated in the subsections that follow.

1) **Data Processing**
   Import and clean 4-D trajectory data and B2B messages, then remove any pre-assigned ATFM ground delays for flights within the ECAC area (Equations 1–2).
2) **Deterministic demand–capacity analysis**
   Compute raw inbound demand in 20-minute bins for each IAF (Equation 4) and compare it with fixed capacity.
3) **Uncertainty quantification**
   Fit Johnson distributions to prediction errors, conditional on time-to-arrival horizon, flight-status class, and time of day.
4) **Probabilistic demand computation**
   Convert individual Cumulative Distribution Functions to per-bin arrival probabilities and derive the exceedance probability $P(\text{demand} > \text{capacity})$.
5) **Speed-control logic**
   Apply a threshold strategy: if the exceedance probability exceeds T, strategic cruise-speed reductions are computed for selected flights.
6) **Trajectory recalculation and simulation**
   Recompute ETOs with the adjusted Mach numbers and run fast-time simulations to verify the impact on flow metrics and delays.
7) **Performance evaluation**
   The performance evaluation assessed how effectively the speed control interventions reduced capacity exceedance probability.

A. Data Processing and ATFM Delay Calculation

ATFM delay is a fundamental component of the methodology and is computed as the difference between the Calculated Take-Off Time (CTOT) and the Estimated Take-Off Time (ETOT) extracted from B2B messages, where the CTOT denotes the ATFM slot assigned to a flight to balance demand and capacity, and the ETOT represents the current best prediction of the flight's actual take-off time. EUROCONTROL's ATFM system employs this standardized approach to quantify ground delays imposed on flights due to capacity constraints or flow management measures. The methodology employs ATFM delay calculations based on Calculated Take-Off Times (CTOT) and Estimated Take-Off Times (ETOT) from B2B messages. ATFM delay is computed as shown in Equation 1 [20].

$$\text{ATFM Delay} = \text{CTOT} - \text{ETOT} \tag{1}$$

For trajectory analysis, the baseline ETO used for demand forecasting is adjusted to remove existing ATFM delays [20]:

$$\text{ETO}_{\text{baseline}} = \text{ETO}_{\text{original}} - \text{ATFM Delay} \tag{2}$$

This approach ensures that the speed control methodology operates on flight trajectories free from pre-existing ground delay assignments, providing a clean baseline for evaluating strategic cruise speed adjustments.

B. Demand and Capacity Analysis

Schiphol's transatlantic arrivals are channeled through specific IAFs in Dutch airspace: SUGOL and RIVER [9]. Each IAF operates under defined arrival capacity constraints, expressed as maximum aircraft within specified time windows. This study employs fixed capacity values per 20-minute interval at each fix, based on operational norms.

For each IAF i and time interval starting at t, we first define the contribution of a single flight f as

$$\delta_{f,i}(t) = \begin{cases} 1, & \text{if } t \leq \text{ETO}_{f,i} < t + \Delta t, \\ 0, & \text{otherwise}, \end{cases} \tag{3}$$

where $\text{ETO}_{f,i}$ is the estimated time over IAF i for flight f, and $\Delta t = 20$ minutes is the interval duration. The deterministic demand at IAF i in the interval $[t, t + \Delta t]$ is then

$$D_i(t) = \sum_f \delta_{f,i}(t). \tag{4}$$

This process creates discrete timelines divided into 20-minute slots, yielding time-series demand profiles for each fix that can be evaluated against capacity constraints.

C. Uncertainty Quantification

The uncertainty quantification methodology builds upon Tielrooij et al.'s foundational work [15], which established that empirical arrival-time prediction errors exhibit heavy-tailed and skewed distributions unsuitable for normal approximation. However, rather than implementing the full Hill-Hill-Holder parameter estimation routine, this study adopts a computationally efficient approach using simplified Johnson distribution fitting to balance accuracy with operational feasibility, which is similar to what has been done by Tielrooij [21].

1. Data Processing and Prediction Errors

Flight trajectory messages are processed to extract prediction errors and operational characteristics. Each flight generates multiple trajectory updates throughout its progression, containing flight status indicators mapped to four discrete categories: Filed (FI), Slot Issued (SI), Assumed Departed (TA), and Confirmed Airborne (AA) [22]. These categories represent increasing levels of operational certainty and information availability.

Prediction errors are calculated as temporal differences between message-contained ETAs and actual arrival times derived from final trajectory messages:

$$\epsilon_{f,m} = \text{ETA}_{f,m} - \text{ETA}_{f,\text{actual}} \tag{5}$$

where $\epsilon_{f,m}$ is the prediction error for flight f at message m, and $\text{ETA}_{f,\text{actual}}$ is the actual flight trajectory derived from the final message. The prediction horizon—defined as the time span between message generation and predicted arrival—provides the primary temporal context for uncertainty analysis.

2. Distribution Modeling and Probability Calculations

Error distributions are characterized using a three-dimensional parameter grid that captures the primary factors influencing prediction uncertainty:

Table 1 Parameter grid specification for uncertainty quantification

Feature	Bin Size	Range
Prediction horizon	10 minutes	0–1440 minutes
Flight status	categorical	FI, SI, TA, AA
Time of day	3 hours	24-hour coverage

The 10-minute horizon binning aligns with typical air traffic control decision intervals, while 3-hour time-of-day segments capture diurnal traffic pattern variations. Only grid cells containing at least 100 samples undergo distribution fitting to ensure statistical reliability. Sparsely populated cells employ nearest-neighbor parameter borrowing using weighted distance metrics that prioritize flight status matching over temporal proximity.

For each populated grid cell, Johnson distribution parameters $(\gamma, \delta, \xi, \lambda)$ are fitted to the empirical error distribution using simplified moment-based approximations rather than rigorous moment matching. Distribution type selection between unbounded (S_U) and bounded (S_B) forms follows the criterion $\kappa \geq \gamma^2 + 1$, where κ and γ represent excess kurtosis and skewness respectively.

The arrival time for flight f at IAF i is then modeled as:

$$T_{f,i} \sim \text{Johnson}(\gamma, \delta, \xi, \lambda) + \text{ETO}_{f,i} \tag{6}$$

The probability that flight f arrives within interval $[t, t + \Delta t]$ is calculated using the cumulative distribution function:

$$P_{f,i,t} = F_{\text{Johnson}}(t + \Delta t - \text{ETO}_{f,i}) - F_{\text{Johnson}}(t - \text{ETO}_{f,i}) \tag{7}$$

where F_{Johnson} is the cumulative distribution function of the fitted Johnson distribution.

3. Probabilistic Demand Calculation

For each flight f with estimated time of arrival (ETA) eta_f, the arrival uncertainty is modeled using flight status-dependent parameters. The predicted arrival time is:

$$\hat{t}_f = eta_f + \mu_f \tag{8}$$

where μ_f is the mean prediction error based on flight status and time to arrival.

The probability that flight f arrives during time interval $[t, t + \Delta t]$ is calculated using the normal distribution:

$$P_{f,t} = \Phi\left(\frac{(t + \Delta t) - \hat{t}_f}{\sigma_f}\right) - \Phi\left(\frac{t - \hat{t}_f}{\sigma_f}\right) \tag{9}$$

where σ_f is the standard deviation of arrival time uncertainty for flight f, and Φ is the standard normal cumulative distribution function.

The expected number of arrivals during interval t is:

$$\lambda_t = \sum_f P_{f,t} \tag{10}$$

For capacity exceedance analysis, the number of arrivals is modeled using a Poisson distribution. This is done because arrivals in a fixed time interval can be approximated as a count process of independent events with a stationary mean rate, for which the Poisson distribution is the standard model and is consistent with recent airport demand modeling approaches [23]:

$$N_t \sim \text{Poisson}(\lambda_t) \tag{11}$$

The probability of exceeding capacity C during interval t is:

$$P(\text{exceedance}) = P(N_t \geq C) = 1 - \sum_{k=0}^{C-1} \frac{\lambda_t^k e^{-\lambda_t}}{k!} \tag{12}$$

This can be efficiently computed using the Poisson cumulative distribution function:

$$P(\text{exceedance}) = 1 - F_{\text{Poisson}}(C-1; \lambda_t) \tag{13}$$

4. Uncertainty Parameters

The flight-specific uncertainty parameters are determined by flight status and time to arrival:

$$\mu_f = \mu_{\text{base}} \cdot \sqrt{\frac{T_f}{60}} \tag{14}$$

$$\sigma_f = \sigma_{\text{base}} \cdot \sqrt{\frac{T_f}{60}} \tag{15}$$

where T_f is the time to arrival in minutes, and $(\mu_{\text{base}}, \sigma_{\text{base}})$ are status-dependent base parameters.

D. Speed Control Logic

Before the speed control algorithm, the cruise speeds for each flight need to be derived from the historical B2B trajectory data. For every flight, the waypoint sequence in the final B2B trajectory is extracted, the great-circle distance between consecutive waypoints is computed using the Haversine formula. The total cruise distance is obtained by summing all cruise segment distances, while the corresponding cruise duration is taken as the difference between the earliest and latest times at the cruise waypoints. The ground speed is obtained as the ratio of total distance to flight time, ensuring that the speed metric reflects the actual route geometry. This ground speed is converted into the Mach number by correcting for the along-track wind component obtained from ERA5 reanalysis winds [24] at the corresponding location, time, and flight level to obtain a mean true airspeed, and dividing this by the local speed of sound under standard-atmosphere assumptions. The resulting Mach is the baseline for the incremental speed reductions applied by the control algorithm.

Building on recent research advocating probabilistic ATFM solutions [5, 16], the methodology implements a threshold-based control logic that triggers speed adjustments only when capacity exceedance probability exceeds a predetermined limit. This approach guards against overcorrection by avoiding unnecessary interventions when congestion forecasts remain highly uncertain or may resolve naturally.

The speed control algorithm triggers interventions when the probability of capacity exceedance exceeds a predetermined threshold T. For each flight f scheduled to arrive at IAF i during interval t, the system evaluates the probability that its scheduled arrival window will result in demand exceeding capacity. When this probability exceeds threshold T, proactive speed control algorithm is triggered for the contributing flight by iteratively checking a reduction of 0.01 Mach till 0.05 Mach. If this reduction would satisfy the constraints, the speed control is triggered.

The incremental cruise Mach number reductions is consistent with operational guidelines for strategic speed management [17]. Individual flight evaluation ensures that only aircraft having a high overload probability are subject to speed reductions, while others continue at planned speeds.

E. Performance Evaluation Framework

The effectiveness of the speed control strategy is evaluated through fast-time simulation. The primary effectiveness metric quantifies the reduction in probability that arrival demand exceeds available capacity. For each intervention, pre- and post-intervention probabilities is calculated by computing the individual aircraft arrival probabilities using normal approximations to Johnson-derived uncertainty distributions, and aggregating demand using Poisson distributions with rate parameter $\lambda_t = \sum_f P_{f,t}$, and (3) calculating exceedance probability as $P(N_t > C)$ where C represents acceptance capacity. The net reduction averaged across all temporal bins constitutes the intervention's effectiveness measure.

F. Simulation Framework

Evaluation of the speed control system employs fast-time simulation using an open-source air traffic control simulator. The simulation provides a realistic environment for modeling aircraft kinematics and air traffic scenarios while supporting custom implementations of operational concepts.

The simulation configuration incorporates all transatlantic inbound flights to Schiphol for the analyzed scenarios, using processed trajectories with baseline ETOs (Eq. (2)) as input data. A custom speed control implementation monitors projected demand at SUGOL and RIVER based on current aircraft trajectories. Detection of capacity exceedance probability above threshold T for upcoming intervals triggers speed reduction commands for affected aircraft during the cruise phase, well before reaching Dutch airspace.

Simulation output recording captures complete trajectories and arrival times at IAFs, enabling a thorough evaluation of the speed control's effectiveness in managing arrival flow while maintaining demand within capacity limits.

IV. Experiment Setup

This section summarises the simulation-based setup used to evaluate an uncertainty-aware, *one-time* cruise speed control strategy for smoothing inbound arrival peaks to Amsterdam Schiphol (AMS). To avoid duplication, we reference the modelling details and equations introduced in section III.

A. Environment & Scope

- **Simulator:** *BlueSky* fast-time simulator [7] with a custom controller that (i) monitors bin-level demand and exceedance risk, (ii) issues at most one cruise Mach reduction $\Delta M <= 0.05$ per flight when the threshold rule is met, and (iii) recomputes times over fixes with wind-adjusted groundspeed; lateral paths remain unchanged.
- **Study window:** 30 consecutive days, **20 July–18 August 2024**.
- **Airspace focus:** Inbound flows to AMS via IAFs **SUGOL** (primary) and **RIVER**.
- **Time grid & capacity:** Fixed **20-minute** bins with nominal per-IAF capacity constants (as operational norms).

B. Input Data

- **Trajectories:** EUROCONTROL *Network Manager B2B* messages (4D waypoints with ETO/ETA updates, identifiers, aircraft type, operational status).
- **Wind:** Meteorological *reanalysis* ERA5 winds sampled along route (lat/lon/alt/time) and projected to the along-track component to correct groundspeeds after any ΔM [24].
- **IAF definitions:** AMS IAFs and routing constraints as per AIP [9].

C. Data Filtering

1) **Scope filters:** Keep inbound transatlantic flights to AMS that route via SUGOL or RIVER and have valid IAF ETOs inside the study window.
2) **Message sanity:** Use the last trajectory message as the flown flight path for simulation.
3) **Neutral baseline:** Remove pre-assigned ground delay per Equations 1–2.

D. Demand Probabilities

- **Deterministic demand:** Build IAF timelines in 20-minute bins using Equation 4.
- **Uncertainty model (after Tielrooij et al.):** Condition ETA-error fits by look-ahead horizon, flight status (FI/SI/TA/AA), and time-of-day; use Johnson $\mathbf{S_U}$ for heavy-tailed, unbounded long-horizon errors and $\mathbf{S_B}$ close-in where support is effectively bounded, then convert ETAs to per-bin arrival probabilities as in Equation 7 [15, 21].
- **Exceedance risk:** Aggregate per-flight probabilities to expected arrivals λ_t and compute $P(X_t \geq C_t)$ using the Poisson-binomial (exact) or the Poisson tail approximation (Equation 12).
- **Threshold rule:** If an upcoming bin's $P(\text{demand} > \text{capacity})$ exceeds a threshold of 70 %, select a minimal set of contributing flights and issue *one* small cruise $\Delta M < 0$ each to shift mass into a later bin, then recompute ETOs with winds.

E. Evaluation Metrics

- **Intervention footprint:** Count of advisories; per-flight absorbed delay (avg/median/percentiles, min/max); number of unique aircraft.
- **Uncertainty reduction:** For each flight, define $U_t = P(X_t \geq C_t)$ before control and U'_t after recomputation; report the mean percentage drop $U_t - U'_t$.
- **Where/when:** Distribution by IAF (SUGOL/RIVER) and by day.

V. Results

This section presents the outcomes of applying the uncertainty-aware cruise speed control methodology to 30 consecutive days (20 July–18 August 2024) of transatlantic inbound traffic to Amsterdam Schiphol. The analysis focuses on intervention characteristics, time patterns, capacity exceedance probability reduction effectiveness, and operational efficiency metrics. All interventions were applied at least one hour before predicted arrival to ensure strategic rather than tactical implementation.

A. Overview and Intervention Frequency

Over the 30-day study period, the speed control system identified and executed **79 interventions** across unique flights approaching AMS via the designated IAFs. Table 2 presents the aggregate intervention characteristics. All aircraft were subject to at most one cruise speed adjustment, consistent with the one-time intervention constraint described in Sec. IV.

Table 2 Summary statistics of speed control interventions

Metric	Value
Total no. of Transatlantic Flights	1,953
Total Interventions	79
Study Period (days)	30
Mean Time to Arrival (hours)	5.16 ± 1.14
Mean Mach Reduction	0.018 ± 0.012
Mean Delay per Intervention (minutes)	8.72 ± 6.44
Mean Probability Reduction (%)	34.96 ± 19.64

The geographical distribution of interventions was highly concentrated: **76 interventions (96.2%)** targeted flights approaching via **SUGOL**, while only **3 interventions (3.8%)** involved the **RIVER**. This distribution reflects both the traffic volume patterns for transatlantic arrivals and the capacity constraints at each fix. The time distribution across the study period exhibited variability consistent with day-to-day fluctuations in transatlantic traffic demand.

B. Intervention Timing Characteristics

Figure 1 shows the distribution of interventions relative to the predicted arrival time at the target IAFs over time. Speed adjustments were typically issued about **5.16 ± 1.14 hours** before the originally scheduled IAF crossing (median: 5.25 hours), with observed values ranging from 1.33 to 7.82 hours. Most interventions are concentrated in a narrow window between roughly 4.5 and 5.8 hours before arrival, indicating that the system generally intervenes at a consistent mid-range look-ahead time.

This timing distribution reflects the strategic nature of the approach, with interventions predominantly occurring while aircraft were still in cruise phase over the Atlantic, well before reaching European airspace with the majority having **4–6 hours before arrival**, which had 54 interventions (68.4%). This indicates that the threshold-based trigger logic (Equation 12) predominantly identified capacity exceedance risks at this time horizon, where uncertainty distributions remain sufficiently constrained to enable effective targeting while providing adequate lead time for meaningful speed adjustments.

C. Speed Reduction Parameters and Induced Delays

The speed control system employed conservative Mach number reductions to minimize operational impact while achieving probability reduction objectives. Figure 2 presents a comprehensive analysis of the speed adjustment parameters and their relationship to the caused delays. Aircraft were subject to cruise Mach reductions ranging from **0.01 to 0.05**, with a mean reduction of **0.018 ± 0.012** and median of 0.01. The distribution of applied reductions demonstrates a strongly conservative approach, with $\Delta M = 0.01$ accounting for 47 interventions (59.5%) and $\Delta M = 0.02$

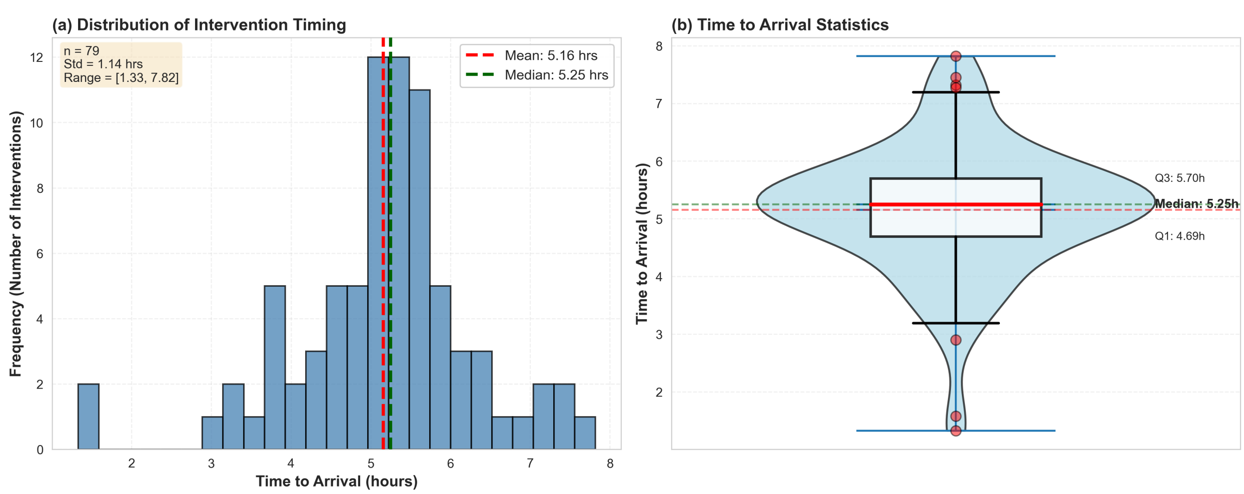

Fig. 1 Distribution of intervention timing relative to arrival. Panel (a) shows the frequency distribution with mean and median markers. Panel (b) shows a box plot overlaid with a violin plot, showing the interquartile range, median, and potential outliers.

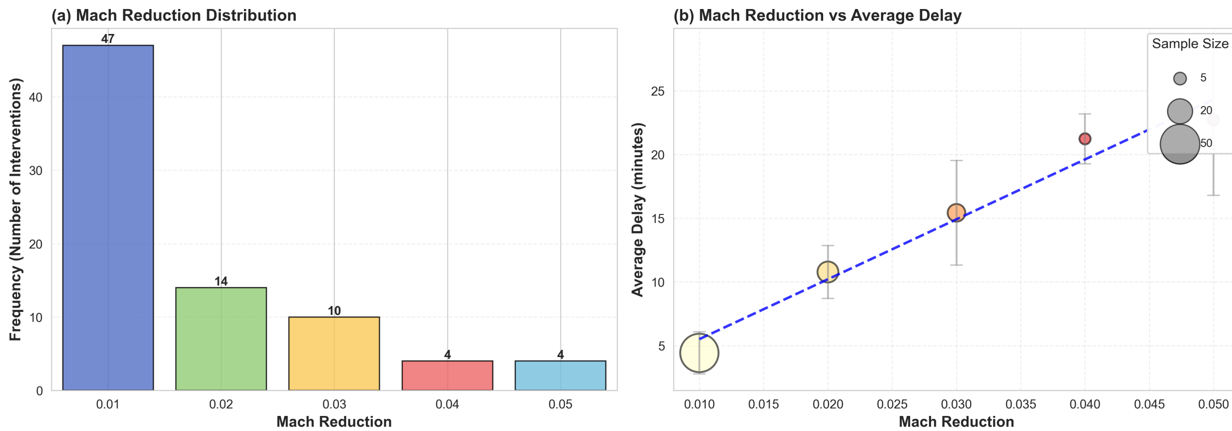

Fig. 2 Mach number reduction analysis. Panel (a) shows the distribution of applied reductions. Panel (b) illustrates the relationship between mach reduction magnitude and average induced delay with standard deviation bars.

for 14 interventions (17.7%). Panel (b) links these settings to their operational effect, showing that larger Mach reductions are associated with higher average induced delays, increasing from roughly 5 minutes at $\Delta M = 0.01$ to over 20 minutes for $\Delta M \geq 0.04$. The bubble sizes correspond to the frequencies in Panel (a), highlighting that the observed trend is primarily driven by the numerous small Mach reductions.

For the interventions, original cruise Mach numbers averaged **0.836** ± **0.025**, reduced to **0.818** ± **0.027** post-intervention. This modest adjustment strategy aligns with operational guidelines for strategic speed management [17] and minimizes fuel penalties while maintaining schedule predictability.

1. Time Delays

The induced time delays, computed via wind-corrected trajectory recalculation (Sec. III), exhibited the following characteristics:

- **Mean delay:** 8.72 ± 6.44 minutes
- **Median delay:** 6.29 minutes
- **Range:** 1.00 to 27.11 minutes
- **Total accumulated delay:** 688.8 minutes (11.5 hours across all interventions)

The delay distribution reveals operational impact concentration: **38.0%** of interventions resulted in delays under 5 minutes, **29.1%** induced 5–10 minute delays, and **32.9%** exceeded 10 minutes. The strong positive correlation between Mach reduction magnitude and induced delay ($r = 0.920$, Figure 2b) confirms the expected linear relationship between speed adjustment and time absorption, validating the trajectory recalculation methodology.

Notably, achieving an average delay of 8.72 minutes through cruise speed adjustment demonstrates the efficiency of strategic intervention. This result supports the viability of cruise speed management as a primary tool for arrival flow smoothing in transatlantic operations.

D. Uncertainty Decomposition and Probability Reduction

A big contribution of this methodology is the quantification of capacity exceedance probability and its reduction through strategic speed control. Figure 3 presents a comprehensive exceedance probability analysis. Prior to intervention, as can be seen from panel (b), the **baseline capacity exceedance probability** averaged approximately 80%. The **net capacity exceedance probability reduction**, which is defined as the improvement in capacity exceedance probability from pre-intervention to post-intervention per flight, averaged approximately 35% (median: 32.75%). Positive values indicate successful reduction in the capacity exceedance probability. The distribution exhibited that **96.2%** of interventions (76 of 79) achieved measurable capacity exceedance probability reduction and **78.5%** achieved reductions exceeding 20%, and, as can be seen from panel (a), the majority fall between 20% and 60%.

E. Intervention Efficiency Analysis

To quantify the operational cost-effectiveness of speed control, Figure 4 presents a comprehensive efficiency analysis and time evolution of interventions. Panel (a) reveals a moderate positive correlation between induced delay and net probability reduction ($r = 0.389$), confirming that larger delays generally produce greater probability reductions. However, the substantial scatter indicates that delay magnitude alone does not determine effectiveness as timing, traffic context, and aircraft positioning play critical roles. Furthermore, panel (b) shows that there much more

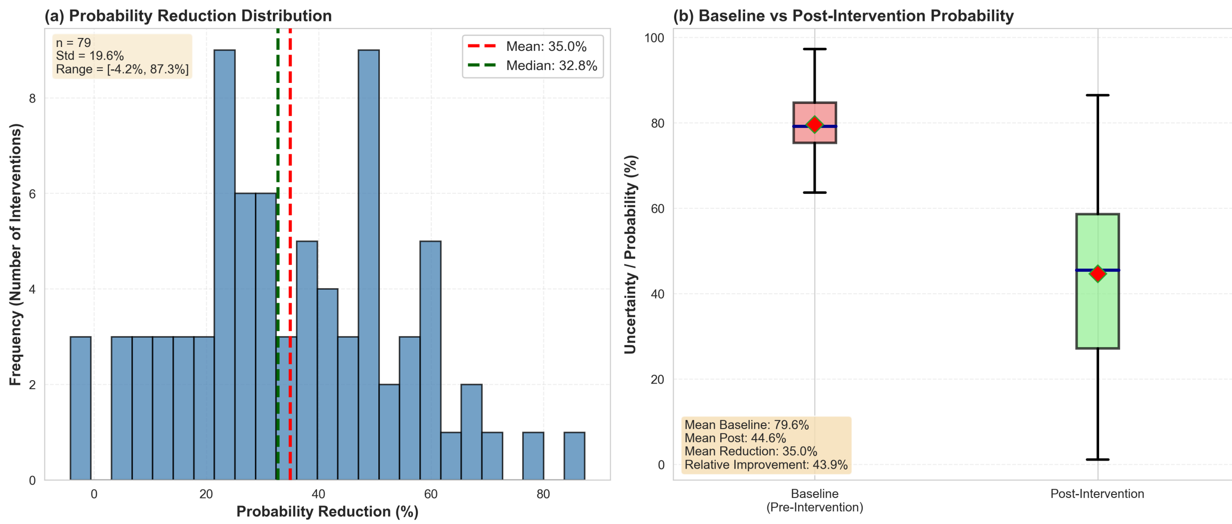

Fig. 3 Capacity exceedance probability analysis. Panel (a) shows the distribution of net capacity exceedance probability reduction achieved across all interventions (positive values indicate successful reduction). Panel (b) compares baseline uncertainty to post-intervention uncertainty, with box plots to see the distribution of before and after intervention.

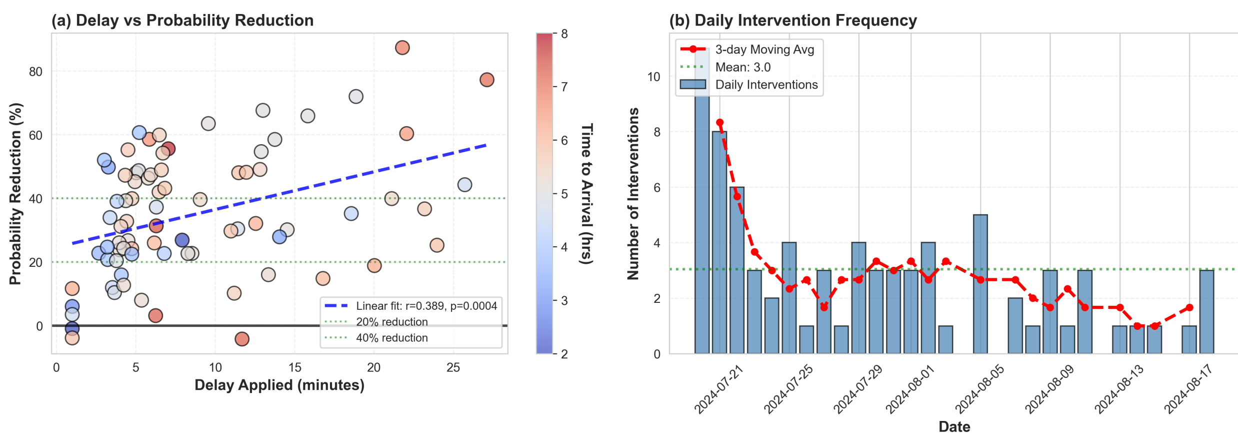

Fig. 4 Intervention effectiveness metrics. Panel (a) shows the relationship between induced delay and net probability reduction, colored by time to arrival, with regression line and reference thresholds. Panel (b) displays the daily frequency of interventions throughout the study period with 3-day moving average.

interventions in the first 3 days of the simulation which shows higher congestion levels in the B2B data for them.

F. Arrival Horizon Effects

To assess how intervention timing relative to arrival influenced effectiveness, the dataset was analyzed by time-to-arrival categories. Table 3 and Figure 5 summarize the results, and the following patterns emerge from this stratification:

- **Near-term interventions (0–2h):** Achieved delays (4.46 min) but achieved correspondingly modest probability reductions (approximately 13%), and in addition, higher mach reductions were required.
- **Optimal window (2–4h):** Achieved substantial probability reductions (approximately 30%) with moderate delays (4.45 min). This window appears to balance predictability improvements with sufficient lead time for speed adjustments.
- **Standard window (4–6h):** Contained the majority of interventions (68.4%) and achieved strong probability reductions (36.31%) with reasonable delays (8.72 min). The concentration suggests this horizon represents the primary operational sweet spot where capacity exceedance risks become sufficiently clear

to trigger intervention, and it does not require a very large mach reduction.

- **Extended horizon (6–8h):** Achieved larger delays (12.66 min) with a slightly higher mach reduction and achieved similar probability reductions (36.36%) as the 4–6h window. The diminishing returns likely reflect greater trajectory uncertainty at extended horizons and reduced precision in targeting specific arrival bins.

Table 3 Intervention outcomes by time-to-arrival category

Time to Arrival	Count	Avg Delay (min)	Mach Red.
0–2 hours	2	4.46 ± 4.89	0.020 ± 0.014
2–4 hours	10	4.45 ± 3.57	0.014 ± 0.013
4–6 hours	54	8.72 ± 5.96	0.018 ± 0.011
6–8 hours	13	12.66 ± 8.11	0.022 ± 0.013

G. Waypoint-Specific Performance

The two IAFs exhibited distinct intervention characteristics, summarized in Table 4 and visualized in Figure 6.

Table 4 Intervention performance by IAF

IAF	Count	Avg Delay (min)	Certainty Red. (%)	Avg TTA (h)
SUGOL	76	8.49 ± 6.31	34.07 ± 19.46	5.10 ± 1.11
RIVER	3	14.45 ± 9.63	57.67 ± 26.83	6.61 ± 1.44

As can be seen from panel (a), **SUGOL** is the primary transatlantic arrival fix for Schiphol, which accounted for 96.2% of interventions (76 of 79). The large sample size provides statistical confidence in the observed performance metrics. Furthermore, it can be seen from panel (b) that average delays of 8.5 minutes yielded probability reductions of 34.07%.

Furthermore, **RIVER**, with only 3 interventions from panel (a), exhibited notably different characteristics, which can be seen from panel (b): larger average delays (14.45 min), substantially greater probability reductions (57.67%). While the small sample size is inconclusive, the results may suggest that RIVER-bound traffic may face higher baseline uncertainty, necessitating more aggressive interventions. Furthermore, from Table 4, the longer average time-to-arrival at intervention for RIVER in comparison to SUGOL (6.61 vs 5.10 hours) indicates these adjustments were applied earlier in cruise.

VI. Discussion

The 79 speed control interventions over 30 days demonstrates the operational viability of uncertainty-aware arrival management. Overall, the results indicate that a simple, one-time cruise speed control scheme can reduce the probability of capacity exceedance at Schiphol's relevant IAF while inducing only modest additional delay per affected flight.

A first key observation is that in 96.2% of cases, the Mach reduction produced a measurable reduction in capacity exceedance probability, with an average improvement of 34.96 %. This suggests that, when combined with an uncertainty-aware thresholding scheme, relatively small speed changes are sufficient to meaningfully reshape the arrival flow distribution. Rather than fine tuning each trajectory, the system benefits from a robust intervention logic.

From an operational perspective, the induced delays appear manageable. The mean delay of 8.72 minutes, absorbed entirely in the cruise phase represents a modest burden, particularly given that 38.0% of interventions only incurred less than 5 minutes delay. These delays are even smaller than, typical airborne holding or vectoring costs near the TMA, but are applied much earlier and in a more predictable fashion, especially given that most interventions involved conservative Mach reductions of only 0.01–0.02.

Because predicted demand-capacity imbalances at the Schiphol's IAF are targeted, these cruise-phase minutes can be interpreted as an upstream buffer that would otherwise have to be absorbed tactically closer to the TMA. While the present study does not explicitly model TMA sequencing or quantify the one-to-one substitution of holding minutes, the observed reduction in capacity exceedance indicates that a non-trivial part of the induced cruise delay is likely shifting, rather than adding, terminal-area delay.

The time to arrival pattern shows the strategic nature of the approach. The concentration of interventions 4-6 hours before arrival (68.4% of cases) shows that early-phase adjustments are both practical and effective. At the same time, the absence of a significant performance difference between earlier and later interventions within the 2-8 hour window indicates that controllers retain flexibility in choosing the exact intervention time, which is important for integration with workload and other operational constraints.

Importantly, the benefits are not limited to single flights. Post-intervention capacity exceedance probability for a flight decreased from 70.63% to 44.60% on average, and this reduction means that potentially some other flights no longer need ATFM delay, indicating a genuine system-wide benefit. This suggests that strategic cruise speed control can enhance overall arrival flow predictability and reduce the likelihood of overload at the IAF, thereby complementing existing ATFM and AMAN tools.

At the same time, several limitations should be acknowledged. The analysis is restricted to a single airport and a 30-day period of transatlantic traffic, which may limit effect in other traffic mixes, seasonal patterns, or conditions. The 30 days period is for the most busy period

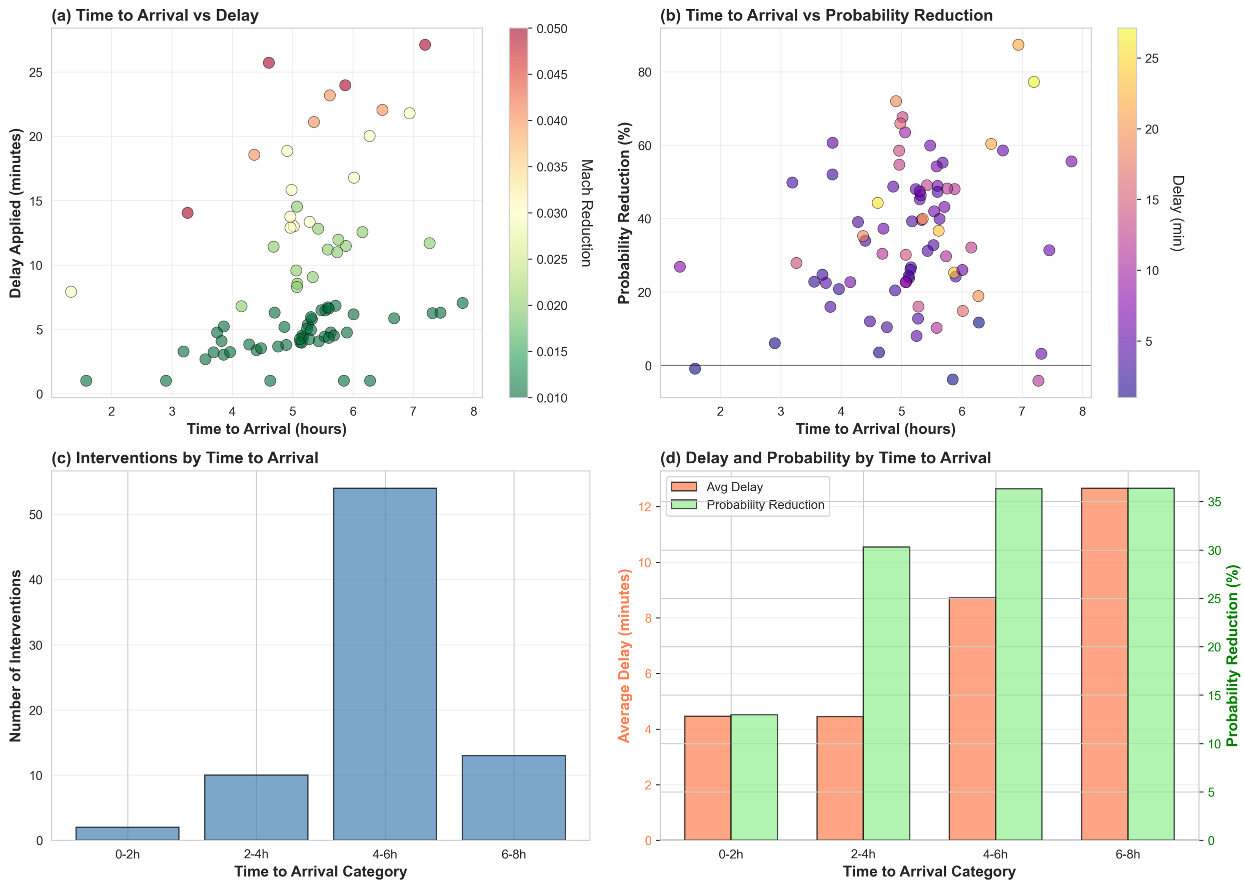

Fig. 5 Time-to-arrival impact analysis. Panel (a) shows the relationship between time to arrival and induced delay, colored by Mach reduction magnitude. Panel (b) presents time to arrival versus net probability reduction, colored by delay. Panel (c) displays intervention frequency by time-to-arrival category. Panel (d) shows a dual-axis comparison of average delay and net probability reduction across categories.

of the year at Schiphol, but it is also required to see how it would function in other periods. Furthermore, the study focuses on a one-time Mach reduction per flight and does not consider a more complex multi-constraint optimization, nor does it explicitly quantify fuel or emissions impacts.

A direct way to assess the hypothesis is to examine whether flights become below the selected speed-control threshold after intervention. **86.1% of flights end below the 70% threshold**, indicating that the proposed uncertainty-aware cruise speed control achieves threshold-compliance for the majority of cases. This supports the claim that a one-time, early cruise-phase Mach reduction can reduce capacity exceedance risk below the threshold. The remaining flights that do not fall below the threshold likely correspond to situations with higher demand peaks.

Taken together, these results indicate that uncertainty-aware, one-time cruise speed control has strong potential as a practical tool for managing transatlantic arrival flows, achieving reductions in capacity exceedance risk with limited additional delay per affected flight, and thereby supporting the hypothesis of this study. It is also important to emphasize that this work evaluates the feasibility of the concept rather than an optimized operational system. Several components of the methodology, such as the uncertainty modeling, capacity representation, and intervention threshold selection, could be refined in future work to make it a practical and effective system.

VII. Recommendations

Building on these findings, several directions emerge for research and implementation. First, the methodology should be extended beyond transatlantic arrivals to other long-haul flights and to mixed traffic streams that include short- and medium-haul continental flights. Since continental traffic constitutes the majority of arrivals at many European hubs and exhibit different uncertainty characteristics and shorter cruise phases, adapting the uncertainty models and intervention timing to these flights is essential for getting system-wide benefits.

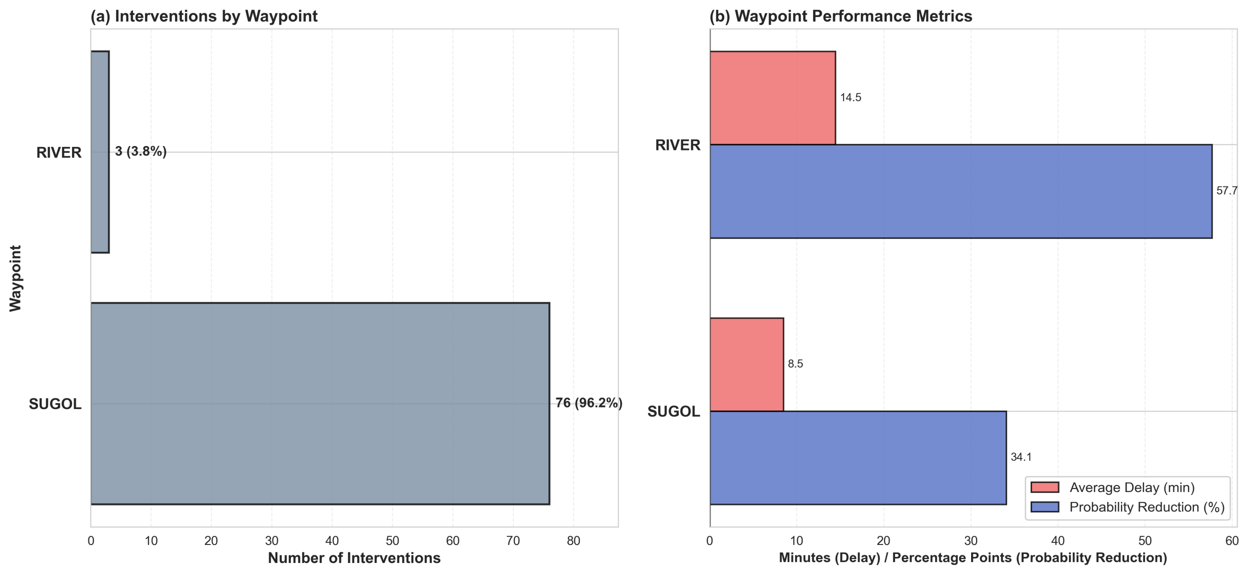

Fig. 6 Waypoint-specific performance analysis. Panel (a) shows intervention frequency by IAF. Panel (b) presents a comparison of average delay and probability reduction for each IAF.

Second, the fixed exceedance probability threshold used in this study could be replaced or complemented by adaptive, data-driven thresholds. Machine learning methods that take into account observed traffic patterns, weather forecasts, seasonal effects, and real-time performance could dynamically change interventions, reducing unnecessary actions during low-uncertainty periods while becoming more conservative when conditions are worse.

Third, incorporating variable and forecast-based capacity constraints would enhance the demand–capacity balancing framework. Arrival capacity at IAFs depends on weather, runway configuration, controller workload, and airspace complexity. Integrating dynamic capacity predictions, for example derived from weather models and historical capacity–weather relationships, would allow the system to tighten or relax intervention criteria in response to predicted capacity changes.

Fourth, the uncertainty quantification component itself can be further refined. Future work should consider additional explanatory factors such as wind forecast uncertainty, airline- and aircraft-specific performance characteristics, and temporal correlation in prediction errors, as well as explore advanced data-driven approaches and potentially hybrid models that blend physics-based prediction with machine-learning uncertainty corrections. Also, integrating speed measurements via satellite datalink may enhance real-time monitoring and support more precise speed advisories.

Finally, a comprehensive assessment of economic, human-factor, and environmental impacts is recommended. Detailed cost-benefit analysis, including fuel consumption, emissions under evolving carbon pricing, and passenger delay costs, would support context-dependent threshold design. Simulator studies and operational trials with air traffic controllers and flight crews are needed to evaluate usability, workload, and acceptance of probabilistic decision support. In parallel, a more detailed quantification of climate benefits would strengthen the sustainability case for long-range, uncertainty-aware cruise speed management.

VIII. Conclusion

This paper addressed a key operational challenge at major hub airports: managing arrival demand peaks that exceed available capacity under substantial uncertainty in long-range air traffic flow management focusing on transatlantic arrivals to Amsterdam Schiphol Airport. Unlike traditional reactive measures such as ground delay programs and terminal holding, the proposed strategy intervenes earlier in the trajectory by adjusting cruise speed, thereby exploiting long-haul cruise segments to absorb delay more efficiently.

The central contribution is an uncertainty-aware, one-time cruise speed control strategy that explicitly integrates probabilistic demand–capacity forecasting into the intervention decision. Rather than assuming perfect information or attempting to eliminate uncertainty, the framework treats uncertainty as part of the decision process. It continuously evaluates the probability that arrival demand at an IAF will exceed available capacity and triggers speed interventions only when this exceedance probability surpasses a predefined threshold. This threshold-based logic provides a safeguard against intervention when forecasts remain highly uncertain and distinguishes the concept from prior long-range ATFM approaches.

Furthermore, by fitting Johnson distributions to historical prediction errors conditioned on flight status and time horizon, the approach quantifies how much of the uncertainty is attributable to individual flights and is therefore potentially controllable via speed adjustments. This enables targeted interventions on those aircraft that can most effectively reduce demand–capacity imbalance.

Using 30 days of simulated transatlantic operations demonstrates that the proposed strategy is operationally realistic. Across 79 interventions, the system achieved an average reduction of 34.96 %. This was done by a mean delay of 8.28 minutes per intervention, implemented through conservative cruise Mach reductions typically between 0.01 and 0.02.

Also, 64% of interventions occurred in the 4–6 hour look-ahead window suggest that the threshold logic successfully identifies risk at horizons where forecasts are sufficiently reliable while leaving enough time to absorb delay through speed adjustment alone.

From an implementation perspective, the approach is intentionally conservative and compatible with current operations. Relying on a single speed intervention per flight and small Mach reductions, it aligns with existing cockpit procedures and controller practice, avoiding the complexity of multiple, tightly coordinated adjustments. The ability to shift 8–10 minutes of delay from low-altitude holding and vectoring to high-altitude cruise supports both operational efficiency and environmental performance, in line with benefits reported in Extended AMAN implementations at other European hubs. Moreover, the concept is designed to work with currently available B2B trajectory data and standard speed control instructions.

Overall, the results show that uncertainty-aware cruise speed control is a viable, practical, and effective strategy for managing long-haul arrival flows into capacity-constrained hub airports. This framework provides a building block for future long-range ATFM that balances operational efficiency, environmental sustainability, and economic viability, while embracing the fundamental uncertainties in long-horizon planning.

References

[1] Commission, P. R., "Performance Review Report 2024: An Assessment of Air Traffic Management in Europe," Prr 2024, EUROCONTROL, Brussels, Belgium, Mar. 2025. URL https://www.eurocontrol.int/sites/default/files/2025-03/eurocontrol-performance-review-report-2024.pdf, published March 2025.

[2] Undertaking, S. J., "Extended Arrival Management (AMAN) horizon," , 2025. URL https://www.sesarju.eu/sesar-solutions/extended-arrival-management-aman-horizon, sESAR Solution #05 — Extended Arrival Management (E-AMAN) horizon.

[3] Zetsche, F., and FABEC XMAN Project, "XMAN: Implementing Extended Cross-Border Arrival Management," Technical presentation WATM2016–XMAN, FABEC, Madrid, Spain, Mar. 2016. URL https://www.fabec.eu/images/newsletter-files/2016-03/XMAN_WATM_Live_ohne_video.pdf, world ATM Congress, 8 March 2016.

[4] Ayhan, S., and Samet, H., "Predicting Estimated Time of Arrival for Commercial Flights," *Transportation Research Part C: Emerging Technologies*, Vol. 146, 2023, p. 103712. https://doi.org/10.1016/j.trc.2022.103712.

[5] Enea, G., and Bronsvoort, J., "Trajectory-Based Operations to Improve Long-Range Air Traffic Flow Management," *Proceedings of the 33rd Congress of the International Council of the Aeronautical Sciences (ICAS 2022)*, Stockholm, Sweden, 2022. URL https://www.icas.org/icas_archive/ICAS2022/data/papers/ICAS2022_0087_paper.pdf.

[6] Vanwelsenaere, A., Ellerbroek, J., Hoekstra, J. M., and Westerveld, E., "Effect of Popup Flights on the Extended Arrival Manager," *Journal of Air Transportation*, Vol. 26, No. 2, 2018, pp. 60–69. https://doi.org/10.2514/1.D0060, URL https://pure.tudelft.nl/ws/portalfiles/portal/47547889/AIAAPaper_Final.pdf.

[7] Hoekstra, J. M., and Ellerbroek, J., "BlueSky ATC Simulator Project: An Open Data and Open Source Approach," *Proceedings of the 7th International Conference on Research in Air Transportation (ICRAT 2016)*, Philadelphia, Pennsylvania, USA, 2016. URL https://pure.tudelft.nl/ws/files/10083831/Hoekstra_BlueSky_project.pdf, conference paper, 20–24 June 2016.

[8] Schultz, M., Lubig, D., Rosenow, J., Itoh, E., Athota, S., and Duong, V. N., "Concept of a Long-Range Air Traffic Flow Management," *Proceedings of the 40th Digital Avionics Systems Conference (DASC)*, IEEE/AIAA, San Antonio, TX, USA, 2020. https://doi.org/10.1109/DASC50938.2020.9256749.

[9] "EHAM AD 2 — Amsterdam/Schiphol Aerodrome," , Aug. 2024. URL https://eaip.lvnl.nl/web/2024-08-22-AIRAC/html/eAIP/EH-AD-2.EHAM-en-GB.html, aIP Netherlands, AIRAC effective 22 Aug 2024.

[10] "Arrival Management at Long Distance," , Dec. 2013. URL https://kdc-mainport.nl/2013/12/01/arrivalmanagement-en-gb-2/, airspace Design & Innovation research article.

[11] Dhief, I., Lim, Z. J., Goh, S. K., Pham, D.-T., Alam, S., and Schultz, M., "Speed Control Strategies for E-AMAN using Holding Detection–Delay Prediction Model," *Proceedings of the 10th SESAR Innovation Days (SIDs 2020)*, Virtual conference, 2020, p. Paper 69. URL https://www.sesarju.eu/sites/default/files/documents/

sid/2020/papers/SIDs_2020_paper_69red.pdf, paper 69, SIDs 2020 technical session.

[12] Jones, J. C., Lovell, D. J., and Ball, M. O., "Stochastic Optimization Models for Transferring Delay Along Flight Trajectories to Reduce Fuel Usage," *Transportation Science*, Vol. 52, No. 1, 2018, pp. 134–149. https://doi.org/10.1287/trsc.2016.0689, URL https://pubsonline.informs.org/doi/epdf/10.1287/trsc.2016.0689.

[13] De Falco, P., and Karaarslan, M., "Probabilistic Constraints Prediction with Uncertainty Quantification in Trajectory-Based Operations," *Proceedings of the SESAR Innovation Days*, 2024.

[14] Tielrooij, M., "Air Traffic Control: How to Deal with Uncertainty?" , Jun. 2023. URL https://to70.com/air-traffic-control-how-to-deal-with-uncertainty/, blog post, published 7 Jun 2023; last updated 3 Jun 2024.

[15] Tielrooij, M., Borst, C., Van Paassen, M. M., and Mulder, M., "Predicting Arrival Time Uncertainty from Actual Flight Information," *Proceedings of the 11th USA/Europe Air Traffic Management Research and Development Seminar (ATM 2015)*, FAA / EUROCONTROL, Lisbon, Portugal, 2015. URL https://www.researchgate.net/profile/Clark-Borst/publication/283861651_Predicting_arrival_time_uncertainty_from_actual_flight_information/links/5651c82508ae4988a7aeb72f/Predicting-arrival-time-uncertainty-from-actual-flight-information.pdf, conference paper, 23–26 June 2015.

[16] Oosterhof, E., Hoekstra, J., and Dijkstra, F., "Effect of Trajectory Prediction Uncertainty on a Probabilistic Debunching Concept for Inbound Air Traffic," Master's thesis, Delft University of Technology, Delft, The Netherlands, Mar. 2022. URL https://repository.tudelft.nl/file/File_31408351-2ea6-477f-bd47-2d835f4c012a, mSc thesis, defended 18 March 2022.

[17] Delgado, L., and Prats, X., "En Route Speed Reduction Concept for Absorbing Air Traffic Flow Management Delays," *Journal of Aircraft*, Vol. 49, No. 1, 2012, pp. 214–224. https://doi.org/10.2514/1.C031484, URL https://doi.org/10.2514/1.C031484.

[18] Matsuno, Y., and Andreeva-Mori, A., "Analysis of Achievable Airborne Delay and Compliance Rate by Speed Control: A Case Study of International Arrivals at Tokyo International Airport," *IEEE Access*, Vol. 8, 2020, pp. 90686–90697. https://doi.org/10.1109/ACCESS.2020.2994109, URL https://doi.org/10.1109/ACCESS.2020.2994109.

[19] Nikoleris, T., Chatterji, G. B., Almog, N., and Palopo, K., "Arrival Delay Absorption Using Extended Metering with Speed Control," *Proceedings of the 12th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, American Institute of Aeronautics and Astronautics, Indianapolis, Indiana, 2012, p. AIAA 2012–5640. https://doi.org/10.2514/6.2012-5640, URL https://aviationsystems.arc.nasa.gov/publications/2012/AIAA-2012-5640.pdf, paper presented 17–19 Sep 2012.

[20] EUROCONTROL Performance Review Unit, "Airport ATFM Delay," https://www.eurocontrol.int/prudata/dashboard/metadata/airport-atfm-delay/, 2025. Single European Sky Portal, metadata last updated 27 March 2025.

[21] Tielrooij, M., "Arrival Management Support in the Presence of Prediction Uncertainty," Ph.D. thesis, Delft University of Technology, Delft, The Netherlands, Oct. 2022. https://doi.org/10.4233/uuid:a403112b-48a2-40d7-9db3-a5b754f31eee, URL https://pure.tudelft.nl/ws/portalfiles/portal/132511938/Dissertation_MTielrooij.pdf, phD dissertation, defended 20 Oct 2022.

[22] Koolen, H., and Coliban, I., *Flight Progress Messages Document*, EUROCONTROL, Brussels, Belgium, 2nd ed., Jul. 2020. URL https://www.eurocontrol.int/sites/default/files/2020-06/flight-progress-msg-update-230620.pdf, document ref. APT/USD/MSG_INTF; edition issue date 01 Jul 2020.

[23] Lancia, C., and Lulli, G., "Predictive modeling of inbound demand at major European airports with Poisson and Pre-Scheduled Random Arrivals," *European Journal of Operational Research*, Vol. 280, No. 1, 2020, pp. 179–190. https://doi.org/10.1016/j.ejor.2019.06.056.

[24] Hersbach, H., and National Center for Atmospheric Research Staff, "The Climate Data Guide: ERA5 Atmospheric Reanalysis," , Sep. 2024. URL https://climatedataguide.ucar.edu/climate-data/era5-atmospheric-reanalysis, last modified 6 Sep 2024.