Fuel-Based Flight Inefficiency Quantification

A Multi-Level Post-Operational Analysis

Lilien Madi

Delft University of Technology

TU Delft

Fuel-Based Flight Inefficiency Quantification

A Multi-Level Post-Operational Analysis

Thesis report

by

Lilien Madi
4836626

to obtain the degree of Master of Science
at the Delft University of Technology
to be defended publicly on December 15th, 2025, at 09:00

Thesis committee:	
Chair:	Dr. Marta Ribeiro
Supervisors:	Dr. Junzi Sun
	Ir. Jeanette Derks (LVNL)
External examiner:	Dr. Alessandro Bombelli
Project Duration:	September, 2024 - December, 2025
Faculty:	Faculty of Aerospace Engineering, Delft

TU Delft

Contents

I

Scientific Article

Fuel-Based Flight Inefficiency Quantification: A Multi-Level Post-Operational Analysis

Lilien Madi*
Supervised by Junzi Sun*, Aidana Tassanbi*, Jeanette Derks†

*Faculty of Aerospace Engineering, Delft University of Technology, the Netherlands
†LVNL - Luchtverkeersleiding Nederland, the Netherlands

Abstract

Improvements in air traffic management (ATM) operations offer a faster pathway to reducing aviation emissions compared to long-term technological solutions. Existing ATM performance frameworks rely on proxy indicators that fail to account for wind effects and vertical profile optimization. This research developed a fuel-based flight inefficiency quantification framework utilizing the open-source OpenAP aircraft performance model and OpenAP.top trajectory optimizer, incorporating wind effects and individual initial mass estimation. The methodology was applied to over 200,000 commercial operations at Amsterdam Airport Schiphol in 2024, comparing executed trajectories against wind-optimal references at city-pair and Flight Information Region (FIR) levels. Results show that strategic inefficiency consistently represents the larger component (7-25% for most routes), while tactical inefficiency is predominantly negative, indicating execution improvements through ATC interventions and wind exploitation. FIR analysis reveals distinct phase characteristics: horizontal inefficiency accounts for approximately 90% of climb inefficiency, while descent operations show vertical inefficiency contributing the larger share with greater seasonal variation in horizontal components. Total relative inefficiencies differ between phases (climb: 9-14%; descent: 24-30%), with the magnitude difference amplified by lower baseline fuel consumption during descent. Wind integration proved essential; for an example flight, omitting wind nearly doubled city-pair inefficiency while underestimating descent inefficiency. Initial mass estimation uncertainty propagates through the analysis and reported values represent conservative upper bounds with actual relative total inefficiencies potentially 15-30% lower. The framework provides researchers and air navigation service providers with tools for environmental performance monitoring and quantitative baselines for evaluating ATM initiatives.

I. Introduction

Commercial aviation faces the dual challenge of reducing environmental impact while accommodating projected traffic growth. In 2022, aviation accounted for approximately 4% of European Union greenhouse gas emissions [1], with absolute emissions continuing to increase despite regulatory targets for climate neutrality in the upcoming decades. While technological advancement has resulted in improvements in aircraft and engine-level fuel efficiency, operational inefficiencies embedded within air traffic management (ATM) systems offset them; it is estimated that 10% of aviation's emissions in Europe are influenced by ATM [2].

These inefficiencies arise from multiple operational factors. Economic considerations, such as airspace charges, influence route selection independent of fuel optimization. Congestion at major airports extends taxi times and incentivizes fuel tankering, a practice conducted by 30% of flights in Europe [3], whereby an aircraft loads more fuel than required to avoid refueling at destination airports or take advantage of favorable fuel prices. ATM constraints, such as airspace structure and flow management practices, result in suboptimal cruise altitudes, non-continuous climbs and descents, and lateral route extensions. The combined fuel penalty from these sources represents a significant fraction of total consumption, motivating research into measurement and reduction strategies.

Existing performance measurement frameworks within the Single European Sky (SES) Performance Scheme rely predominantly on proxy indicators that compare executed trajectories to geometric references [4]. The implicit assumption underlying these metrics has been shown to fail under operational conditions where wind optimization, vertical profile selection, and speed management influence fuel burn [5].

Fuel-based performance indicators overcome limitations of proxy indicators by directly quantifying fuel consumption. At a European level, advancements have been made to develop and implement such metrics, however, technical requirements such as aircraft performance models, four-dimensional trajectory optimization, weather data integration, and initial mass estimation have constrained its adoption. Existing implementations rely on proprietary tools with limited methodological transparency which limit accessibility, reproducibility and comparison across different organizations.

This research develops a framework for fuel-based flight inefficiency quantification that explores current implementation barriers and utilizes OpenAP aircraft performance model [6] and OpenAP.top trajectory optimizer [7]. The methodology is applied through analysis of commercial flight operations at Amsterdam Airport Schiphol (EHAM) throughout 2024, at a city-pair and Flight Information Region (FIR) level.

Section II reviews existing literature on indicators utilized for performance assessment and trajectory reconstruction and generation methods. The main research question and subquestions are also introduced in this section. Section III describes data sources utilized. Section IV presents the fuel-based performance indicators employed in this study. Section V presents the proposed methodological framework. Section VI reports results including city-pair and FIR-level performance. Section VII discusses findings, and section VIII provides conclusions.

II. Literature Review & Research Definition

A. Proxy Indicators

Early efforts to quantify flight inefficiency relied on geometric indicators, which were favored for their simplicity and ability to be calculated from available surveillance data [8]. The most established of these is the concept of Horizontal Flight Efficiency (HFE), which is found by comparison of the length of a flown trajectory to the shortest possible distance (great-circle/orthodromic distance) between its endpoints [5]. This general concept is formally implemented in the SES Performance Scheme through two specific Key Performance Indicators: KEP and KEA. These two indicators are directly related: KEP (Key performance Environment indicator based on last filed flight Plan) measures the efficiency of the planned route, while KEA (Key performance Environment indicator based on Actual trajectory) measures the efficiency of the executed path the aircraft using ground track data [9]. The focus of these indicators is strictly on the en-route portion of the flight to remove the influence of terminal procedures [9]. To achieve this, the areas within a 40 nmi radius of the departure and arrival airports are excluded [9]. This principle of measuring horizontal efficiency is often tailored by individual Air Navigation Service Providers (ANSPs) for their specific monitoring needs. For example, Dutch ANSP Luchtverkeersleiding Nederland (LVNL) reports an ”Internal en-route HFE,” a metric designed to measure the additional distance flown specifically within the Amsterdam FIR excluding the 40 nmi radius around Schiphol. Similarly, the UK’s ANSP, NATS, employs an HFE component within its 3Di score that is measured relative to its own airspace boundaries [10]. However, a critical limitation of HFE is its assumption that the shortest distance is the most efficient. Multiple studies revealed this is often not the case from a fuel perspective [11] [8] [12], as a flight with low lateral inefficiency (a good HFE score) can simultaneously exhibit ”relatively poor fuel performance” due to factors such as a suboptimal altitude or speed profile [8]. In a focused study, it was found that there is ”no total correlation” between horizontal route extension and fuel efficiency, with the two metrics often acting in opposition for a considerable number of flights [5]. Other contributing factors

include wind; for many routes, the true fuel-optimal path deviates from the great-circle track to take advantage of strong tailwinds or avoid headwinds, meaning a longer ground track can burn less fuel [10][13].

Given that HFE does not account for inefficiencies in the vertical profile, it is conventionally complemented by Vertical Flight Efficiency (VFE), which measures the deviation of a vertical profile of a trajectory from its optimal vertical path. Defining an optimal vertical path is inherently challenging, as it depends on dynamic factors such as weather, aircraft weight, and operational constraints [14]. This challenge is further complicated by the phase-dependent definition of an efficient vertical profile; a sustained level flight is the goal of cruise, whereas it is a primary source of inefficiency during climb or descent. For this reason, VFE assessment is typically separated into two main areas: en-route vertical efficiency, which focuses on the cruise segment, and the efficiency of the climb and descent portions of the flight.

Monitoring of the en-route phase involves assessing the optimality of the cruise altitude. Methodologies often compare a flight's actual cruise level against a benchmark, which can be derived from either unconstrained flights on routes of a similar length [14][4] or from a statistical 'best performer' profile based on historical data [15][4].

The assessment of the climb and descent phases, as developed by EUROCONTROL's Performance Review Unit (PRU), focuses on identifying periods of level flight during climb and descent [4]. The benchmark for this method is an ideal Continuous Climb/Descent Operation (CCO/CDO), where the aircraft flies a smooth profile with minimal level segments, starting from the ideal top of descent. The specific parameters for detecting a level-off, however, can differ based on the monitoring framework. PRU cites a vertical speed below 300 fpm for over 20 seconds [14], whereas other studies have adopted slightly different thresholds, such as vertical speed below 300 fpm for over 30 seconds[16][17]. The implementation by LVNL for its CCO/CDO monitoring also uses a duration of more than 20 seconds at a vertical rate of 300 fpm or less. LVNL's analysis also segments climbs and descents into different altitude bands and derives specific indicators like the "percentage of flights without level segments" and the "average time in level flight per flight". Other approaches have also been implemented, such as the VFE component within the NATS 3Di score which defines vertical inefficiency by level flight occurring below the requested cruise altitude [10].

Beyond detecting level segments, other studies quantify vertical inefficiency by measuring the deviation between the executed and an ideal altitude profile, expressed in units like ft x minutes [18], or ft x hour [15]. A newer approach, proposed in [14], redefines VFE as a fuel-based metric, expressing VFE as a percentage difference in fuel burn compared to an optimal reference trajectory. Research has also focused on the specific fuel impact of VFE drivers, such as the fuel-based penalty incurred from early descents [19].

Temporal metrics which measure additional time are also common; LVNL actively monitors ASMA (Additional Time in Arrival Sequencing & Metering Area) to capture extra flight time within the 40 nmi radius of Schiphol and complements this with "Taxi-In and Taxi-Out Additional Time" for ground operations. In these specific terminal and ground contexts, additional time is generally a strong proxy for inefficiency, as it is primarily driven by congestion, holding, and vectoring that lead to increased fuel burn. However, when the concept of additional time is applied to the en-route phase of flight, it can be misleading, as it may penalize fuel-efficient strategies like speed reduction which intentionally increase flight duration [20].

B. Fuel Based Indicators

To overcome limitations in capturing the environmental performance of complete flight trajectories, there has been a shift towards fuel-based indicators, as fuel consumption integrates all dimensions of a flight's trajectory into a single metric [8]. The most direct approach is to calculate the absolute or relative fuel penalty of an executed flight compared to an ideal reference [4].

A prominent example of this approach is EUROCONTROL's eXcess Fuel Burn (XFB) indicator, a network-level metric that quantifies additional fuel consumed compared to a reference based on historical performance. Rather than using a theoretical optimum, XFB employs a statistical "best performer" benchmark, defined as the 5th or 10th percentile of all fuel burn observations for the same airport pair and aircraft type over a given period. While this data-driven approach provides practical benchmarking, it has recognized limitations: it remains a relative measure that can be influenced by external factors beyond stakeholder control, such as wind conditions. Additionally, the methodology lacks transparency, and EUROCONTROL advises against using this indicator at the ANSP-level. Similarly, ANSPs like MUAC have developed their own fuel-based horizontal flight efficiency indicators, as well as vertical flight efficiency metrics that quantify excessive fuel burn of non-optimal vertical profiles using BADA models [4].

Other European stakeholders are developing alternative fuel-based indicators to address some of these limitations [4]. The airline Vueling, supported by Spanish ANSP ENAIRE, has proposed the KEO concept, which replaces the great-circle reference of traditional efficiency metrics with dynamically calculated fuel-optimal trajectories that account for aircraft weight and weather conditions. France's DSNA is developing ACROPOLE indicators that leverage machine learning to estimate fuel burn from radar tracks and calculate "wind corrected fuel scores" against statistical benchmarks. These emerging initiatives show a commitment to improved fuel-based monitoring, though they currently face challenges including limited methodological transparency, validation scopes confined to the developer's operational context, and varying degrees of maturity. However, stakeholders acknowledge these limitations, and ongoing development efforts aim to address these gaps [4].

In parallel, the research community has also utilized fuel-based metrics to better understand the root causes of flight inefficiency. Numerous studies have used excess fuel burn as the primary metric to quantify the impact of ATM; it has been used to evaluate the fuel penalty of specific operational procedures like early descents [19] and inefficient terminal area flows [21], to quantify the potential benefits of ideal, conflict-free operations [22], and to compare the systemic performance of different regional ATM systems [14]. In [23], econometric techniques are used to deconstruct airline-reported fuel burn data, statistically attributing consumption to different forms of delay and terminal inefficiency.

In an effort to gain a better overview of the causes of flight inefficiencies and their respective contributions, a separate line of research has emerged with a focus on development of hierarchal frameworks capable of deconstructing the total inefficiency into its constituent sources. A conceptual basis for this is presented in [10], where a cascade of references is proposed to separate inefficiencies into procedural, pre-tactical, and tactical layers. Similarly, in a study for the SESAR AURORA project [11], a multi-layered set of trajectories are used to calculate a suite of user-centric cost-based indicators. An alternative, data-driven approach to creating reference layers is shown in [15], which uses statistical benchmarks like the "Best Performer" (best observed on a specific route) and the "Ideal Performer" (best observed over similar distances) to distinguish between operational and structural inefficiencies.

A particularly extensive implementation of this hierarchical approach, and the one whose indicators are adopted for this thesis, is presented in [13]. This work implements the concept with a detailed family of nine fuel-based indicators. The framework is specifically designed to decouple total fuel inefficiency into the contributions from the tactical and strategic ATM layers, and further separates these into their horizontal and vertical components. This is achieved by comparing the executed trajectory against a hierarchy of reference trajectories, including the filed flight plan and various cost-optimal trajectories.

The implementation of these types of frameworks requires a two-stage computational process: (1) trajectory reconstruction, which uses flight data to determine the performance of the executed flight path in order to calculate fuel burn, and (2) trajectory generation, which computes the optimal reference trajectories needed for comparison.

C. Trajectory Reconstruction

The calculation of a trajectory's fuel consumption is dependent on the use of aircraft performance models. These are computational tools that simulate an aircraft's performance, allowing for the estimation of fuel burn for any given trajectory, provided the aircraft type is supported. The foundation of most aircraft performance models is a point-mass model, which simplifies the complex dynamics of an aircraft to a single point in space. These models then relate the forces of thrust, drag, lift, and weight to the aircraft's motion. The most prominent and widely cited of these is Eurocontrol's Base of Aircraft Data (BADA).

BADA consists of two core models: the kinetic-based Aircraft Performance Model (APM) and the Airline Procedure Model (ARPM). The APM is organized into four sub-components: Actions (computing aerodynamic, propulsive, and gravitational forces), Motion (calculating trajectories via ordinary differential equations ODEs), Operations (defining real-life flight procedures), and Limitations (maintaining safe operational boundaries). BADA's core models utilize aircraft-specific coefficients which are derived from reference performance data. [24]

In practice, as detailed in [16], BADA is used to first calculate the aerodynamic drag for a given flight segment. The Total Energy Model is then solved to determine the engine thrust required. Finally, this thrust value is translated to a coefficient which is fed into the corresponding fuel consumption model. Due to its comprehensive nature, BADA has been the engine of numerous flight efficiency studies [25][16][17][11]. However, as a propriety tool restricting the open sharing of data and code, it has motivated the development of alternatives.

This thesis utilizes OpenAP, an open-source aircraft performance model developed in [6]. OpenAP is founded on the same point-mass principle but distinguishes itself by using open-source data. It comprises four main components: (1) a database of aircraft and engine properties, (2) a data-driven kinematic model (WRAP) derived from statistical analysis of ADS-B surveillance data, which describes typical flight parameters like speed and vertical rate (3) a dynamic performance model for calculating forces and fuel consumption based on the ICAO aircraft engine emissions databank and literature models, and (4) utility libraries for aircraft performance calculations [6]. OpenAP calculates fuel consumption by using its dynamic model to determine required engine thrust from aircraft motion, then applies engine-specific fuel flow coefficients to convert thrust into fuel burn rate based on altitude.

Beyond general-purpose models, the literature also highlights the use of more specialized or alternative modeling approaches. For instance, a study presented in [19] employs the Airbus Performance Engineering Programs (PEP) suite, a propriety tool which leverages manufacturer-specific data to provide higher accuracy for the supported aircraft types. Other models have been developed to address the particular needs of specific organizations. Examples include the Australian ANSP's Dalí trajectory modeler, used in [10], and TU Dresden's SOPHIA model, used in [14], for flight efficiency assessment.

Alternative modeling approaches include machine learning-based and reduced models to derive fuel burn. A study presented in [26] used bagged decision trees for regression and classification, and hierarchical cluster analysis for anomaly detection. These models were trained on flight data to create aircraft-specific fuel consumption predictions that account for real-world variability such as engine degradation, route characteristics, and weather. An example of a reduced model is FEAT (Fuel Estimation in Air Transportation), presented in [27]. FEAT works by pre-computing equations that relate flight distance to fuel burn for each aircraft type, allowing for extremely fast, large-scale estimations based on aggregated flight schedules rather than detailed individual trajectories. Reduced models like FEAT offer a trade-off between accuracy and speed. The various models presented offer different advantages and limitations; thus, model choice depends on the user's research objectives and accessibility to tools.

A critical input parameter for aircraft performance models and a significant source of uncertainty in the trajectory reconstruction and generation phases, is the initial aircraft mass. This data is commercially sensitive and not readily available. To resolve this, a variety of approaches are taken, from using the actual take-off weight when provided

directly by an airline [23][25], or using proprietary tools such as BR&TE's PERCEPT [11], to more mathematical approaches such as the application of Bayesian inference for initial mass estimation [28]. When such data or methods are unavailable or a large-scale studies are performed, heuristic assumptions are common. These include using a fixed percentage (e.g., 90%) of the maximum take-off weight [29][17][13], testing a range of weights to capture uncertainty [5][27], utilizing a standard reference mass for the aircraft type [22], or an average of its empty and maximum weights [21].

A second significant challenge in accurately reconstructing flight performance is the incorporation of meteorological data, particularly wind. Wind vectors are essential for converting an aircraft's observed ground speed and track into its true airspeed (TAS), which is required for the accurate calculation of aerodynamic forces, required thrust, and fuel consumption. The most robust approach and one that is demonstrated in numerous studies, is to integrate 4D gridded data from numerical weather prediction models, such as the ERA5 reanalysis dataset [25] or the Global Forecast System (GFS) from NOAA [11][13].

This method requires multi-dimensional linear interpolation to align the high-resolution trajectory points with the coarser grid of the weather data [25]. Where such comprehensive datasets are not employed, simplified methods are used; in a study analyzing arrivals, wind is estimated using the the power law wind profile equation and average wind speed and direction information from an airport database [21]. For some analyses focused on specific flight phases or where a lower fidelity is acceptable, wind effects are sometimes omitted entirely for simplicity [17].

D. Trajectory Generation

The generation of reference trajectories required for flight inefficiency analysis is computationally intensive and relies on dedicated optimization engines, and the literature showcases a variety of specific algorithms and tools used for this purpose. In [17], CDO is defined as the benchmark for arrival efficiency, and is generated by formulating and solving it as an optimal control problem. In a study of arrivals into Tokyo [22], a Dynamic Programming (DP) algorithm is employed to find the fuel- and time-optimal path for individual flights. Some of the aforementioned trajectory reconstruction services also provide trajectory generation utilities; in a comparison of European and Chinese airspace [14], SOPHIA is utilized to generate fuel-optimal references. BR&TE'S PERCEPT is employed in [11] to generate reference trajectories. Tools have also emerged from efforts made in educational institutions, such as DYNAMO, a 4D trajectory prediction and optimization engine, which is employed in [13].

Founded on OpenAP, OpenAP.top was developed as an open-source 4D trajectory optimizer capable of incorporating all flight phases into the optimization [7]. It implements a nonlinear optimal control approach and uses the CasADi symbolic framework with the IPOPT solver to compute trajectories that are optimal with respect to fuel, time, cost, or environmental objectives such as global warming and temperature potential. The optimization can be applied to complete trajectories or individual flight phases while considering wind effects, making it suitable for general flight inefficiency analysis and specialized applications such as ANSP performance monitoring.

E. Research Gap & Questions

Despite advances in methodologies for flight inefficiency analysis, the implementation of fuel-based metrics remains fragmented across different stakeholders, from the definitions of reference trajectories to the methodologies and tools utilized. The predominance of closed-source tools also creates barriers to reproducible research and widespread adoption. Furthermore, complexities in implementation of these indicators such as obtaining initial mass, incorporating wind effects, and computational demand, further prevent the adoption of these metrics.

To explore these limitations and aid standardization efforts, this research develops a framework for fuel-based flight inefficiency analysis, combining trajectory reconstruction through OpenAP with trajectory optimization using

OpenAP.top. The approach is applied to analyze flights inbound and outbound Amsterdam Airport Schiphol for the year 2024, incorporating multiple flight data sources, initial mass estimation, and wind data integration. A case study on the Amsterdam FIR is also conducted for the purpose of testing the methodology for ANSP performance monitoring.

This research objective is structured through the following main research question and supporting sub-questions:

How can environmental flight inefficiencies of commercial flights be modeled and quantified?

- How do executed and planned trajectories compare to optimal trajectories in terms of flight inefficiency?
- What insights can be drawn from the comparison of inefficiency across different city pairs and airlines?
- How do horizontal and vertical flight inefficiencies compare in their contribution to total inefficiency in the Amsterdam FIR?
- How do different flight phases contribute to flight inefficiency of a flight?
- How do input parameters such as initial mass and true airspeed affect the quantification of flight inefficiencies?
- What are the limitations in the proposed methodology for quantification of flight inefficiencies?

The following sections present the data, key performance indicators, and methodology developed to address these research questions.

III. Data Sources

A. Automatic Dependent Surveillance-Broadcast (ADS-B) Data

Executed flight trajectory data is obtained from the OpenSky Network, a global community-based sensor system that collects and archives Automatic Dependent Surveillance-Broadcast (ADS-B) data [30]. ADS-B is a surveillance technology in which aircraft automatically broadcast their state vectors including position, altitude, and velocity, derived from onboard Global Navigation Satellite System (GNSS) receivers. The OpenSky Network aggregates these transmissions from a distributed network of ground receivers and provides historical flight datasets for academic research. The detailed structure of the data is presented in Table I.

Importantly, beyond ADS-B data, the OpenSky Network also archives Mode S Comm-B messages transmitted by aircraft. These messages are utilized in this study to validate the TAS derivation.

TABLE I: Structure of raw ADS-B data from the OpenSky Network.

Field	**Type**	**Unit**
time	timestamp	-
icao24	string	-
lat	float	deg
lon	float	deg
velocity	float	m/s
heading	float	deg
vertrate	float	m/s
callsign	string	-
onground	bool	-
alert	bool	-
spi	bool	-
squawk	string	-
baroaltitude	float	m
geoaltitude	float	m
lastposupdate	float	s
lastcontact	float	s
serials	object	-
hour	timestamp	-

B. Demand Data Repository (DDR2) Data

Data from EUROCONTROL's Demand Data Repository (DDR2) is utilized to provide planned trajectories and an additional source of executed trajectories. DDR2 data is provided in the form of SO6-format files. This study uses two types of DDR2 SO6 files:

- **Model 1 (M1):** contain flight data as per the last filed flight plan, representing the planned trajectory of the flight.
- **Model 3 (M3):** represent the executed trajectory, which is derived from the last filed flight plan updated when flights deviate by more than 5 minutes, 7 FL, or 20 nmi.

Both M1 and M3 data are structured as a series of connected flight segments, where each segment is defined by a start and end longitude, latitude, altitude, and time. The detailed structure of the SO6 data format is presented in Table II.

TABLE II: Structure of DDR2 M1 and M3 flight data.

Field	Type	Unit
segment identifier	string	-
origin of flight	string	-
destination of flight	string	-
aircraft type	string	-
time begin segment	string	-
time end segment	string	-
FL begin segment	integer	FL
FL end segment	integer	FL
status	integer	-
callsign	string	-
date begin segment	string	-
date end segment	string	-
lat begin segment	float	arcmin
lon begin segment	float	arcmin
lat end segment	float	arcmin
lon end segment	float	arcmin
flight identifier	integer	-
sequence	integer	-
segment length	float	nmi
segment parity/color	integer	-

C. ECMWF ERA5 Reanalysis Data

The meteorological data used in this framework is sourced from the ERA5 reanalysis dataset, provided by the European Centre for Medium-Range Weather Forecasts (ECMWF)[31]. ERA5 offers global coverage with a horizontal resolution of $0.25° \times 0.25°$ and a temporal resolution of one hour. The wind data is structured as a four-dimensional grid consisting of latitude, longitude, pressure level (altitude), and time. The dataset provides the eastward (u-component) and northward (v-component) wind velocities in meters per second. For any given point in space and time within a flight's trajectory, multi-linear interpolation on this 4D grid provides an estimate of the wind conditions experienced by the aircraft.

IV. Key Performance Indicators

Five of the nine fuel-based KPIs developed in [13] are utilized in this study and illustrated in Figure 1. These KPIs support a city-pair flight inefficiency analysis and a case study within the Amsterdam FIR.

In the following KPIs, the Reference Business Trajectory (RBT) - a component of ICAO's Trajectory Based Operations (TBO) concept - can be defined as a dynamic agreement of the Airspace User's intended flight path approved by relevant stakeholders through coordination with the Network Manager. In this work, the RBT is equivalent to the planned trajectory and the DDR2 M1 files will be employed to represent them.

A. City-pair Analysis

The total fuel inefficiency is defined as

$$\Delta F_T = \hat{F}_e - F^*, \tag{1}$$

where ΔF_T represents the discrepancy between the fuel consumption of the executed flight ($\hat{F}_e$) and the corresponding optimal flight (F^*).

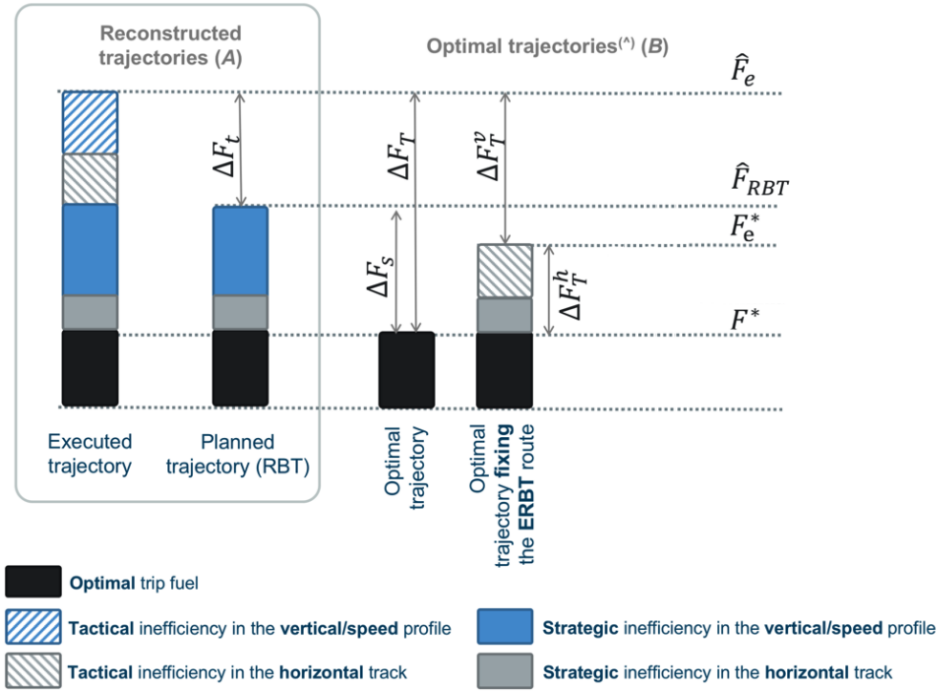

Fig. 1: Framework of the five fuel-based indicators used in this study [13]. The figure has been adapted to exclude four more granular indicators beyond the scope of this research.

The total strategic inefficiency is given by

$$\Delta F_s = \hat{F}_{\text{RBT}} - F^*, \tag{2}$$

which quantifies the difference between the fuel consumption of the planned ($\hat{F}_{\text{RBT}}$) and optimal (F^*) trajectories.

The total tactical inefficiency is expressed as

$$\Delta F_t = \hat{F}_e - \hat{F}_{\text{RBT}}, \tag{3}$$

representing the fuel difference between the executed ($\hat{F}_e$) and the planned ($\hat{F}_{\text{RBT}}$) trajectories .

B. FIR Analysis

The total horizontal inefficiency is defined as

$$\Delta F_T^h = F_e^* - F^*, \tag{4}$$

where F_e^* denotes the fuel consumption of an optimal trajectory that constrains the executed route's horizontal path while optimizing the vertical/speed profile. The difference between this constrained optimal trajectory and the optimal trajectory (F^*) yields the horizontal inefficiency component.

The total vertical inefficiency is calculated as the difference between the total and total horizontal inefficiencies:

$$\Delta F_T^v = \Delta F_T - \Delta F_T^h. \tag{5}$$

Relative inefficiencies are also analyzed by normalizing (1), (2), (3), and (4) by the fuel consumption of the optimal trajectory and expressing them as percentages. The relative total vertical inefficiency is found as the difference between the relative total and relative total horizontal inefficiencies. A negative value indicates improved flight efficiency compared to the reference.

V. Methodology

The following data processing framework, presented in Figure 2, is designed as a multi-stage pipeline, capable of processing one month of data at a time with optional granularity for specific days or hours. The benefit of this is two-fold: it makes the processing computationally manageable, and enables targeted analysis such as evaluating the environmental impact of new ATM procedures on a smaller scale. The methodology is organized into four phases: (1) Data Collection (2) Trajectory Reconstruction (3) Trajectory Generation, and (4) Fuel Consumption Calculation & Data Aggregation.

After the initial data collection phase, 210,000 flights were processed using an M4 Pro chip parallelized across 10 cores, requiring approximately 72 hours (1.2 seconds per flight).

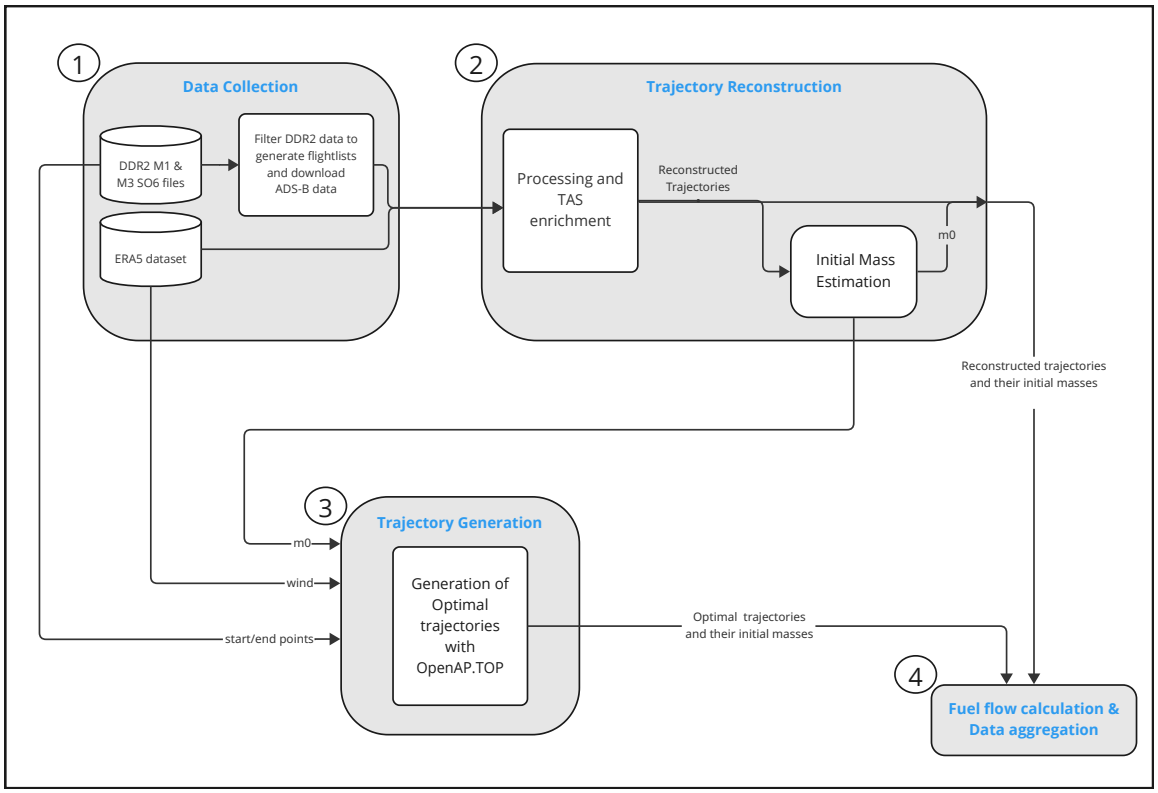

Fig. 2: Proposed data processing framework.

1. Data Collection

Raw DDR2 M1 and M3 data for flights arriving and departing EHAM are obtained for each month in 2024 through the repository's Filtered Traffic service. The subsequent step is the first of many filtering stages which are incorporated to prevent the processing of unused data in computationally intensive stages. In the first stage, DDR2 data is filtered by the origin and destination airports having ICAO codes starting with 'E' or 'L', as the scope of the study is limited to European airspace. Data is also filtered by aircraft type, retaining only those supported by OpenAP.top and the fuel flow model in OpenAP. This step inherently removes most non-commercial flights, such as those operated on general aviation aircraft and helicopters. M1 and M3 files are also cross-referenced and any flights not presented in both datasets are discarded. Table III represents the aircraft types used for this analysis and the percentage of total European flights covered.

TABLE III: Aircraft types and the percentage of total European flights covered.

Typecode	**Flights**
A319 A320 A321 A332 A333 B737 B738 B739 E190	65%

Filtered DDR2 data serve as the basis for generating flight lists, which are summary files that evolve throughout the pipeline, progressively accumulating metadata for each flight. Flight lists also allow dataset-specific information, such as flight identifiers and aircraft type from DDR2, or ICAO24 codes from ADS-B, to be exchanged across sources. This completes datasets without dependence on external aircraft databases. The initial version of the flight list is utilized to query trajectory data from the OpenSky network through the OpenSky Trino database using the `pyopensky` library. The ADS-B data retrieval process includes several fallback query strategies to maximize the data acquisition success rate. Before being saved, ADS-B trajectories are checked for existing coverage at or below an altitude of 10,000 ft; although ADS-B coverage over Europe is relatively good, there are areas where coverage is incomplete at lower altitudes. Furthermore, 10,000 ft is chosen as the threshold as fuel optimization is not generally prioritized in operations below this altitude and is excluded from this study. Flights with incomplete ADS-B coverage are removed and filtered from other corresponding datasets.

The final step, which can be performed concurrently with other stages, is pre-syncing the wind data from the Google ARCO ERA5 store using the `fastmeteo` library [32], developed to facilitate the retrieval and interpolation of meteorological data for flight trajectory enrichment.

2. Trajectory Reconstruction

The DDR2 data is first transformed from a segment-based representation into continuous time series, which allows the derivation of ground speed, vertical rate, and track. DDR2 and ADS-B trajectories are subsequently filtered and uniformly resampled at 10-second intervals using the `traffic` library [33]. Missing values are removed, and all quantities are expressed in consistent units required for the fuel burn calculation stage. The trajectories are then enriched with wind data using the `fastmeteo` library for rapid interpolation of wind components at each trajectory point. This interpolation enables the calculation of the TAS, which is determined by the vector subtraction of the interpolated wind vector from the aircraft's ground speed vector:

$$\vec{V}_{TAS} = \vec{V}_{GS} - \vec{W} \tag{6}$$

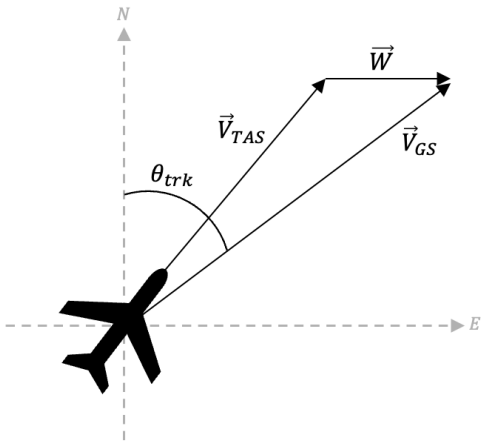

Fig. 3: Relationship between aircraft true airspeed, ground speed, and wind velocity.

Validation of the TAS calculation was performed against TAS values decoded from Mode S Enhanced Surveillance (EHS) messages with the `pyModeS` library [34] for one day of flights, resulting in a Mean Absolute Error (MAE) of approximately 2 kts.

Trajectory reconstruction is continued with the estimation of the initial mass from ADS-B trajectories using a three-feature multi-linear regression model with aircraft-specific coefficients developed in [35]. The estimation uses mean cruise altitude, mean cruise TAS, and total flight distance as inputs. To find the mean cruise altitude and TAS, cruise segments are identified using `FlightPhase`, an OpenAP utility that applies fuzzy logic to classify flight phases based on altitude, vertical rate, and ground speed. For instance, cruise is characterized by high altitude, high speed, and minimal vertical rate [36].

A large proportion of the executed flights are operated below 30,000 ft mean cruise altitude, predominantly short-haul operations where optimal cruise altitude is not reached as a significant portion of the flight is spent in the climb and descent phase, and where ATC constraints may restrict flights from reaching optimal cruise altitudes. Since the estimation model was trained on a dataset of optimal trajectories with mean cruise altitudes above 30,000 ft and without considering wind, the model exhibits higher uncertainty for the executed flights with a general overestimation (often above 1.0) of the ratio of initial mass to maximum takeoff weight (MTOW). Realistic mass estimates are necessary for accurate fuel burn calculations. Furthermore, trajectory generation in subsequent stages was found to fail above a mass ratio of approximately 0.85 for those flights. Mass ratio overestimates are therefore constrained to this value.

3. Trajectory Generation

For the selected KPIs, the city-pair analysis requires complete optimal trajectories, whereas the FIR analysis requires both unconstrained and constrained optimal trajectories within the FIR.

1) Complete Optimal: Complete wind-optimal trajectories are generated for each flight using the OpenAP.top optimizer, which utilizes non-linear optimal control approach to solve the optimization problem and generate optimal trajectories. Although it is capable of generating trajectories for different objectives, this study focuses solely on fuel-optimal trajectories.

The trajectory optimization is formulated as a non-linear optimal control problem, utilizing a point-mass aircraft performance model. The state vector at time t is defined as:

$$\mathbf{x}_t = [x_t, y_t, h_t, m_t] \tag{7}$$

where x_t and y_t represent the aircraft's position coordinates, h_t is the altitude, and m_t is the aircraft's mass. The control vector is given by:

$$\mathbf{u}_t = [M_t, \mathrm{vs}_t, \psi_t] \tag{8}$$

where M_t is the Mach number, vs_t is the vertical speed, and ψ_t is the heading of the aircraft. The system's dynamics are governed by the following ordinary differential equations:

$$\frac{dx}{dt} = v_t \sin(\psi_t) \cos(\gamma_t) + w_{x,t} \tag{9}$$

$$\frac{dy}{dt} = v_t \cos(\psi_t) \cos(\gamma_t) + w_{y,t} \tag{10}$$

$$\frac{dh}{dt} = vs_t \tag{11}$$

$$\frac{dm}{dt} = -\mathrm{ff}_t(m, v, h) \tag{12}$$

where v_t is the TAS, γ_t is the flight path angle, $w_{x,t}$ and $w_{y,t}$ are the wind components, and ff_t is the fuel flow rate as a function of mass, TAS and altitude. The objective function for the optimization is to minimize the total fuel consumption over the flight, from time t_0 to t_f, using the following generalized form of an objective function (J):

$$J(\mathbf{x}, \mathbf{u}, t_0, t_f) = \int_{t_0}^{t_f} \mathrm{ff}_t(\mathbf{x}, \mathbf{u}, t_0, t_f)\, dt \tag{13}$$

The following constraints are considered in the minimization of the objective function, where $\dot{\mathbf{x}}$ represents the first-order system dynamics given by the equations defined earlier, $\mathbf{h}(\cdot)$ represents the path constraints, and $\mathbf{e}(\cdot)$ represents the endpoint constraints.:

$$\dot{\mathbf{x}}_t = \mathbf{f}(\mathbf{x}_t, \mathbf{u}_t) \tag{14}$$

$$\mathbf{h}(\mathbf{x}_t, \mathbf{u}_t) < 0 \tag{15}$$

$$\mathbf{e}(t_0, t_f, \mathbf{x}_{t_0}, \mathbf{u}_{t_0}, \mathbf{x}_{t_f}, \mathbf{u}_{t_f}) = 0 \tag{16}$$

The optimal control problem is solved numerically using the direct collocation method. This approach divides the continuous problem into discrete segments and time intervals, approximating states with polynomials at collocation points. OpenAP.top employs CasADi, an open-source symbolic optimization framework, which calls on the IPOPT solver to solve the resulting nonlinear programming problem.

For each flight in the dataset, the optimizer is initialized with the flight's origin, destination, aircraft type, estimated initial mass ratio, and the relevant 4D wind field for the flight's duration. The optimizer returns a complete wind and fuel-optimal trajectory.

2) FIR-Optimal: The FIR-Optimal corresponds to the unconstrained trajectory, representing the theoretically optimal climb or descent profile. In addition to its ability to generate complete flight trajectories, OpenAP.top provides the functionality to independently calculate phase-specific optimal trajectories. This is achieved through

modification of endpoint constraints, path constraints for states, path constraints for aircraft performance, and the generation of a pseudo-optimal cruise to establish the top of climb (TOC) and top of descent (TOD) waypoints for the climb and descent modules, respectively. In this study, these climb and descent optimizers have been adapted to generate the necessary reference optimal trajectories.

Considering first the climb phase, the following endpoint constraints are enforced:

$$(x_0, y_0, h_0, m_0) = (x_{\text{start}}, y_{\text{start}}, h_{\text{start}}, m_{\text{start}}) \tag{17}$$

$$(x_f, y_f, h_f) = (x_{\text{end}}, y_{\text{end}}, h_{\text{end}}) \tag{18}$$

$$m_{\text{oew}} < m_f < m_{\text{start}} \tag{19}$$

where the subscript 'start' denotes the aircraft state at 10,000 ft, and 'end' corresponds to the point at which the FIR boundary is crossed.

Path constraints on state and control variables remain consistent with the original climb optimizer, with the exception of stricter heading control constraints necessitated by the relatively shorter segment lengths. Path constraints for aircraft performance are maintained without modification.

Path constraints ensuring smooth control variable transitions have been refined to accommodate the increased node density and reduced segment distances. These bounds have been established through a combination of observed rate of change of variables from executed flight segments and appropriate scaling of the original constraints:

$$-0.02 < M_{k+1} - M_k < 0.02 \tag{20}$$

$$-200 < vs_{k+1} - vs_k < 200 \quad \text{(fpm)} \tag{21}$$

$$-2.5 < \psi_{k+1} - \psi_k < 2.5 \quad \text{(deg)} \tag{22}$$

An important distinction between the original climb optimizer and the modified formulation lies in the a priori knowledge of entry and exit points. Consequently, the FIR crossing point is treated as the TOC, eliminating the need for a pseudo-optimal cruise or additional constraints regarding alignment of the final position with the optimal cruise trajectory, as described in [7].

The FIR descent optimizer implements the inverse trajectory optimization problem. The boundary condition formulation mirrors that of the climb optimizer, with reversed endpoint constraints: 'start' corresponds to the state at the FIR boundary crossing point (designated as the TOD), while 'end' represents the state of the aircraft at 10,000 ft. Constraints on state and control variables are similar to the original descent optimizer, and path constraints for aircraft performance remain unmodified.

Both modified optimizers incorporate atmospheric wind effects using wind field data with the resolution of ERA5 reanalysis data.

3) Constrained Optimal: The second reference trajectory required for the FIR inefficiency analysis is an optimal trajectory fixing the lateral path of the executed route. This is referred to as the constrained optimal trajectory. Unlike the OpenAP climb and descent optimizers that compute complete 4D trajectories, the constrained optimizer focuses exclusively on vertical profile optimization. Due to convergence difficulties in simultaneous speed and vertical rate optimization, the velocity profile is fixed based on observed flight data, with optimization limited to the vertical rate. This allows a reduced-order problem formulation with the following state vector:

$$\mathbf{x}_t = [h_t, m_t]^T \tag{23}$$

where h_t represents altitude and m_t represents aircraft mass. The control vector reduces to a single variable representing vertical speed:

$$\mathbf{u}_t = vs_t \tag{24}$$

The system dynamics are governed by simplified ordinary differential equations:

$$\frac{dh}{dt} = vs_t \tag{25}$$

$$\frac{dm}{dt} = -\text{ff}_t(m, v, h) \tag{26}$$

where ff_t is the fuel flow rate as a function of mass, TAS, and altitude. In this formulation, TAS is an externally provided function defined as the piecewise linear interpolant of the executed flight data points, which constrains the optimizer to follow the observed speed profile.

The objective function for the optimization is to minimize the total fuel consumption over the flight using the generalized form of an objective function presented in Equation 13.

Endpoint constraints are extracted from observed trajectory data, fixing initial and final altitude and initial mass:

$$(h_0, m_0) = (h_{\text{start}}, m_{\text{start}}) \tag{27}$$

$$h_f = h_{\text{end}} \tag{28}$$

Considering first the climb optimizer, the following path constraints for state and control variables are adapted in this optimizer:

$$h_{\text{start}} < h_k < h_{\text{end}} \tag{29}$$

$$m_{\text{oew}} < m_k < m_{\text{start}} \tag{30}$$

$$0 < vs_k < 2500 \text{ (fpm)} \tag{31}$$

Path constraints for aircraft performance ensure physical feasibility at each point in the trajectory, and have been adopted from the complete optimizer in OpenAP.top.

The first constraint enforces that available thrust (maximum climb thrust) exceeds drag with a 5% safety margin:

$$T_{\max,k} \cdot 0.95 - D_k > 0 \tag{32}$$

where drag is computed from the drag polar:

$$D_k = \left(C_{D_0} + k \left(\frac{m_k g}{\frac{1}{2}\rho_k v_k^2 S} \right)^2 \right) \frac{1}{2}\rho_k v_k^2 S \tag{33}$$

The second constraint ensures adequate lift generation with a 20% margin, to prevent stalls:

$$L_{\max,k} \cdot 0.8 - m_k \cdot g > 0 \tag{34}$$

where maximum lift is determined from the maximum lift coefficient:

$$L_{\max,\,k} = C_{L,\max} \cdot \frac{1}{2}\rho_k v_k^2 S \tag{35}$$

with $C_{L,\max}$ set to a default value of 1.4.

The third constraint validates that excess power is sufficient for the required increase in kinetic and potential energy:

$$\frac{T_{\max,k} - D_k}{m_k} - \frac{g}{v_k}\frac{dh_k}{dt} - \frac{dv_k}{dt} > 0 \tag{36}$$

All aircraft model coefficients for thrust, drag polar, and geometric properties are obtained from the OpenAP database. Control smoothness constraints ensure flyable profiles and have more permissive bounds than the FIR optimizers to accommodate reduced control authority:

$$-500 < vs_{k+1} - vs_k < 500 \quad \text{(fpm)} \tag{37}$$

Similar to the FIR optimizers, the boundary condition formulation of the constrained descent optimizer mirrors that of the constrained climb optimizer, with reversed endpoint constraints: 'start' corresponds to the state at the FIR boundary crossing point, while 'end' represents the state of the aircraft at 10,000 ft. Additionally, the descent optimizer constrains vertical speeds to $vs_k \in [-2000, 0]$ fpm, control smoothness to ± 1000 fpm, and utilizes cruise thrust settings in the path constraints for aircraft performance.

4. Fuel Flow Calculation & Aggregation

The final phase of the pipeline is the fuel consumption calculation, performed for three trajectory types: the executed trajectory derived from ADS-B data, DDR2 M3 data, and the planned trajectory from DDR2 M1 data. The `openap.FuelFlow` model is employed for this task; the model iterates through each point in a trajectory, calculating the instantaneous fuel flow rate based on the aircraft's current mass, TAS, altitude, and vertical rate. The fuel consumed during each time step is then calculated and subtracted from the aircraft's mass for the subsequent step. The total fuel consumption is computed as the summation of fuel burned during each flight segment at altitudes above 10,000 ft, to establish a consistent comparison across all datasets. For the optimal trajectories, fuel consumption is calculated using the same `openap.FuelFlow` model, but integrated directly within the optimization framework and later extracted from the optimal solution.

VI. Results

A. Fuel Flow Validation

The credibility of the results are dependent on the accuracy of the OpenAP FuelFlow model. Figure 4a presents a validation case comparing OpenAP fuel consumption against Quick Access Recorder (QAR) data for a flight operated on a B738 obtained from a large European Airline. The cumulative fuel consumption curves show close agreement between all three trajectory sources (ADS-B, DDR2 M3, and QAR reference) throughout most of the flight, with final total consumption estimates differences ranging from 0.5% for ADS-B data to 5.5 % for DDR2 M3 data. The fuel flow rate comparison in Figure 4 reveals more detailed model behavior; during climb and cruise phases, the OpenAP model based on ADS-B as well as DDR2 M3 inputs captures the general trends and magnitude of fuel flow more accurately relative to the descent phase, where differences are larger. Despite these differences, OpenAP FuelFlow model exhibited good performance compared to the QAR fuel flow.

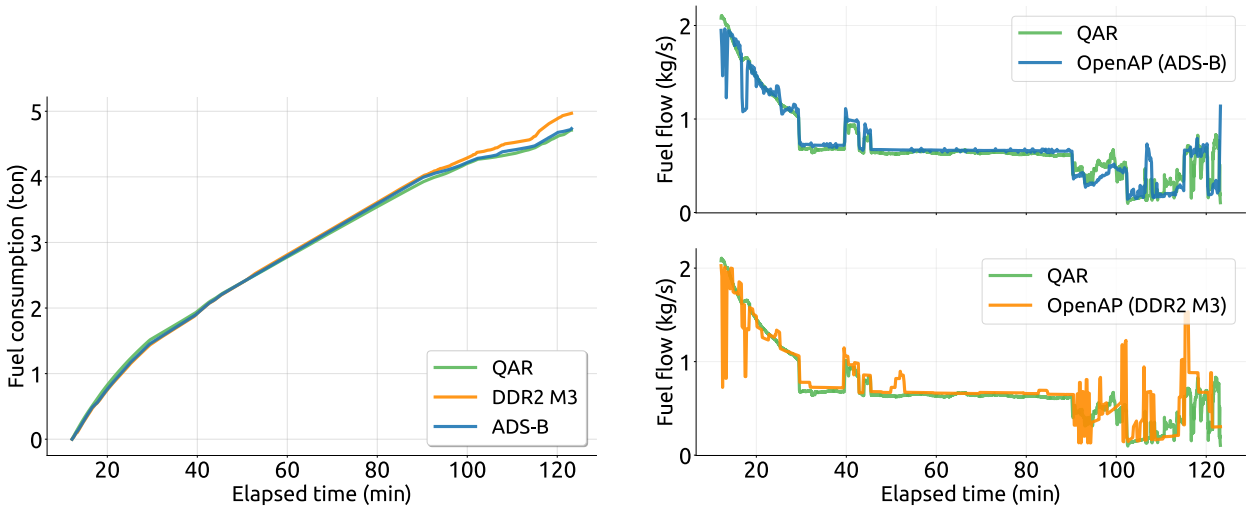

(a) Cumulative fuel consumption.

(b) Fuel flow rate comparison.

Fig. 4: Fuel flow validation with QAR data obtained from a European airline.

B. General Traffic Characteristics

This study included approximately 210,000 flights to and from Amsterdam Airport Schiphol (EHAM). Figure 5 presents the distribution of flights for 2024 across the 30 city pairs with the highest flight frequencies, revealing Barcelona (LEBL) as the most frequent connection with over 4000 flights in each direction, followed by London City (EGLC) and Dublin (EIDW) with approximately 3,800 and 3,600 flights respectively. The volume of inbound and outbound flights are nearly symmetrical across all city pairs, which is consistent with typical short to medium haul flight operations in Europe where aircraft usually execute same-day turnarounds.

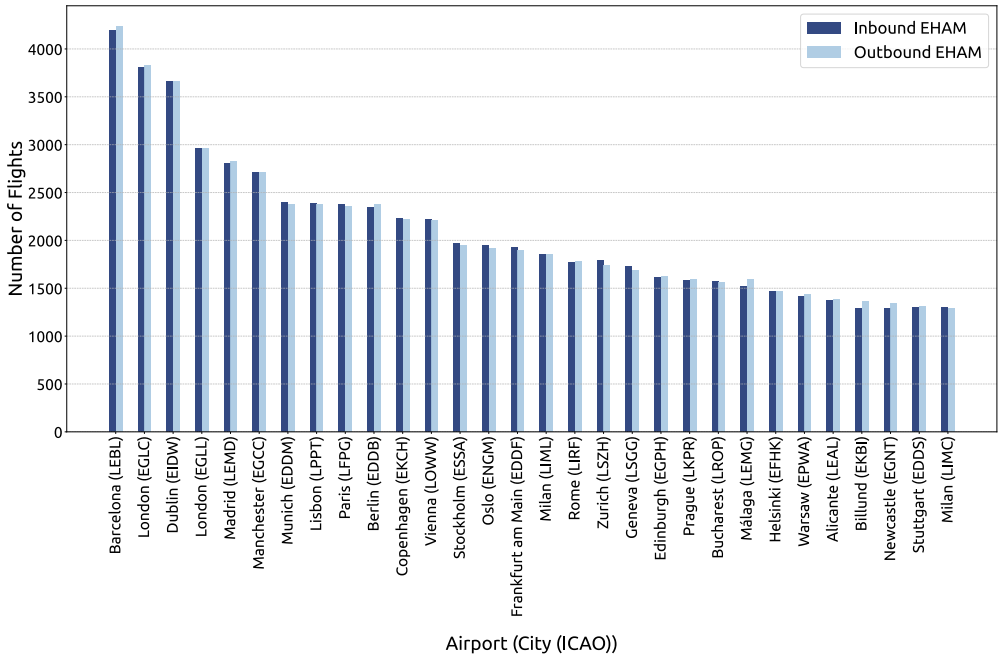

Fig. 5: Flight frequency distribution for the 30 busiest city-pairs in 2024.

Seasonal variations in the flight network are illustrated in Figure 6a and Figure 6b, comparing February and August 2024 respectively. The winter month shows concentrated traffic on core routes, with the United Kingdom, Spain, and Germany representing the dominant flow patterns. The average daily flight frequency for these primary routes ranges between 10-20 flights per day. In contrast, the summer month exhibits a significant network expansion both in geographic spread and intensity. Apart from the core routes, Mediterranean destinations such as Spain, Italy, and Greece, experience increases in service frequency, with several routes exceeding 10 flights per day.

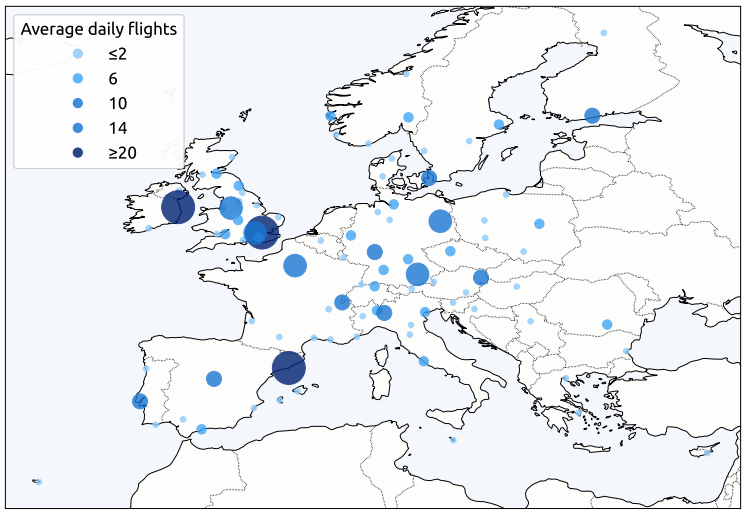

(a) February 2024

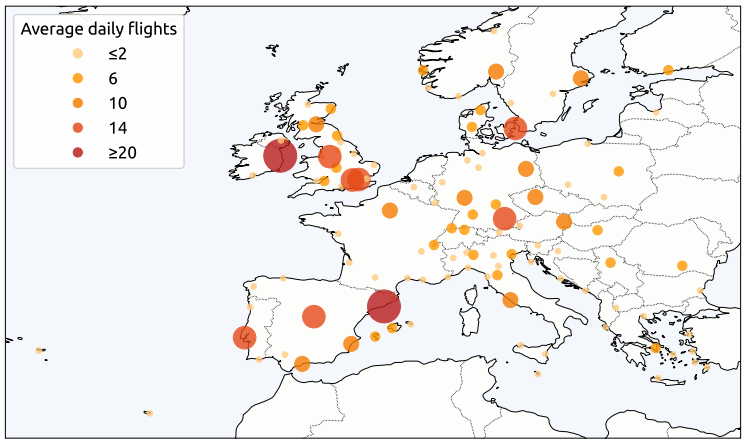

(b) August 2024

Fig. 6: Seasonal flight network variations illustrated through comparison of February and August traffic in 2024.

Figure 7 presents the monthly variation in average fuel consumption per flight throughout 2024, comparing four trajectory types: planned trajectories (DDR2 M1), executed trajectories (DDR2 M3 and ADS-B), and wind-optimal reference trajectories. Seasonal patterns are consistent across all data sources and may be attributed to the increased frequency in flights, especially to southern Europe.

Wind-optimal trajectories consistently demonstrate the lowest fuel consumption across all months, with values ranging between 2,200-2,550 kg per flight. On the contrary, planned trajectories exhibit the highest fuel consumption, exceeding the wind-optimal average by 20%, suggesting substantial inefficiency embedded in the strategic plan-

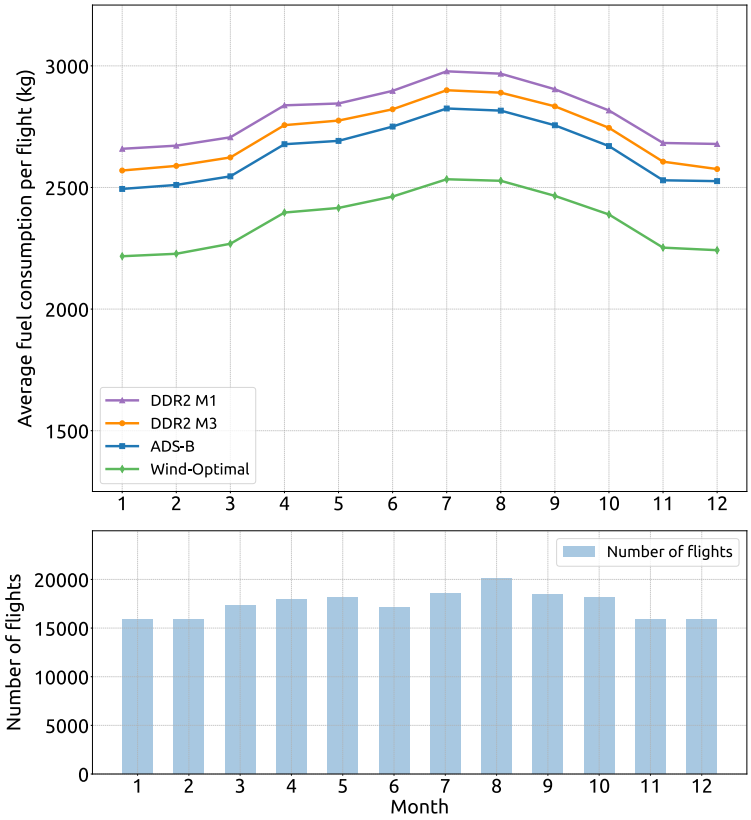

Fig. 7: Monthly flight frequency and average fuel consumption per flight.

ning phase. Executed trajectories consistently fall between these two extremes, with ADS-B trajectories showing marginally lower consumption than DDR2 M3 trajectories.

Figure 8 presents the distribution of fuel consumption across the three main flight phases for the chosen dataset. The analysis reveals that cruise operations have the highest percentage of fuel consumption, accounting for 47.5% of the total fuel burned. Climb operations consume 39.0% of total fuel, while descent accounts for the remaining 13.5%. Climb and descents segments collectively make up over half of the fuel consumption of a flight. This distribution remained relatively constant across all months analyzed, and provides important context for interpreting the inefficiency results. Since cruise represents the highest fuel consumption phase, even small percentage improvements in cruise efficiency translate to large absolute fuel savings. The relatively large proportion consumed during climb also suggests that continuous climbs to optimal altitudes could yield significant efficiency gains.

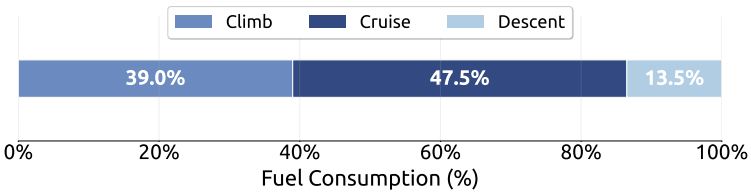

Fig. 8: Fuel consumption distribution across climb, cruise, and descent phases.

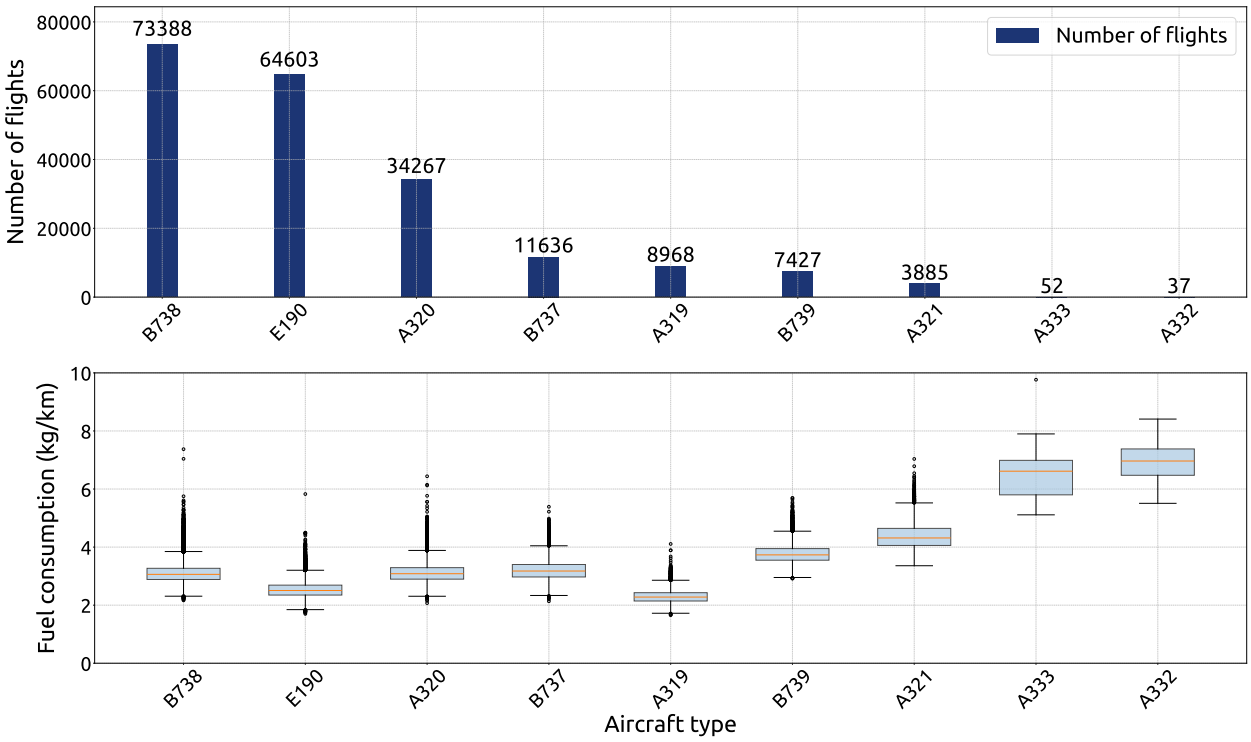

Fig. 9: Distribution of flight operations and fuel consumption by aircraft type.

The data show significant variation in operational frequency and fuel consumption across aircraft types, as presented in Figure 9. The Boeing 737-800 (B738) accounts for 73,388 flights, followed by the Embraer E190 with 64,603 flights, collectively representing approximately 68% of the total filtered dataset. The E190, Airbus A320 family (A319, A320, A321) and Boeing 737 variants (B737, B738, B739) constitute the majority of the fleet, while wide-body A330 variants (A332, A333) account for only 89 flights combined. Median fuel consumption, normalized with great circle distances, varies across aircraft types and ranges from 2.28 kg/km for the A319 to 6.96 kg/km for the A332. Within the narrow-body segment, smaller aircraft including the E190 (2.50 kg/km) and A319 (2.28 kg/km) show the lowest consumption, mid-sized aircraft including the B738, A320, and B737 cluster around 3.1-3.2 kg/km, while larger variants such as the B739 (3.73 kg/km) and A321 (4.31 kg/km) show elevated consumption. The relatively narrow interquartile ranges across most aircraft types indicate consistent operational performance within each fleet.

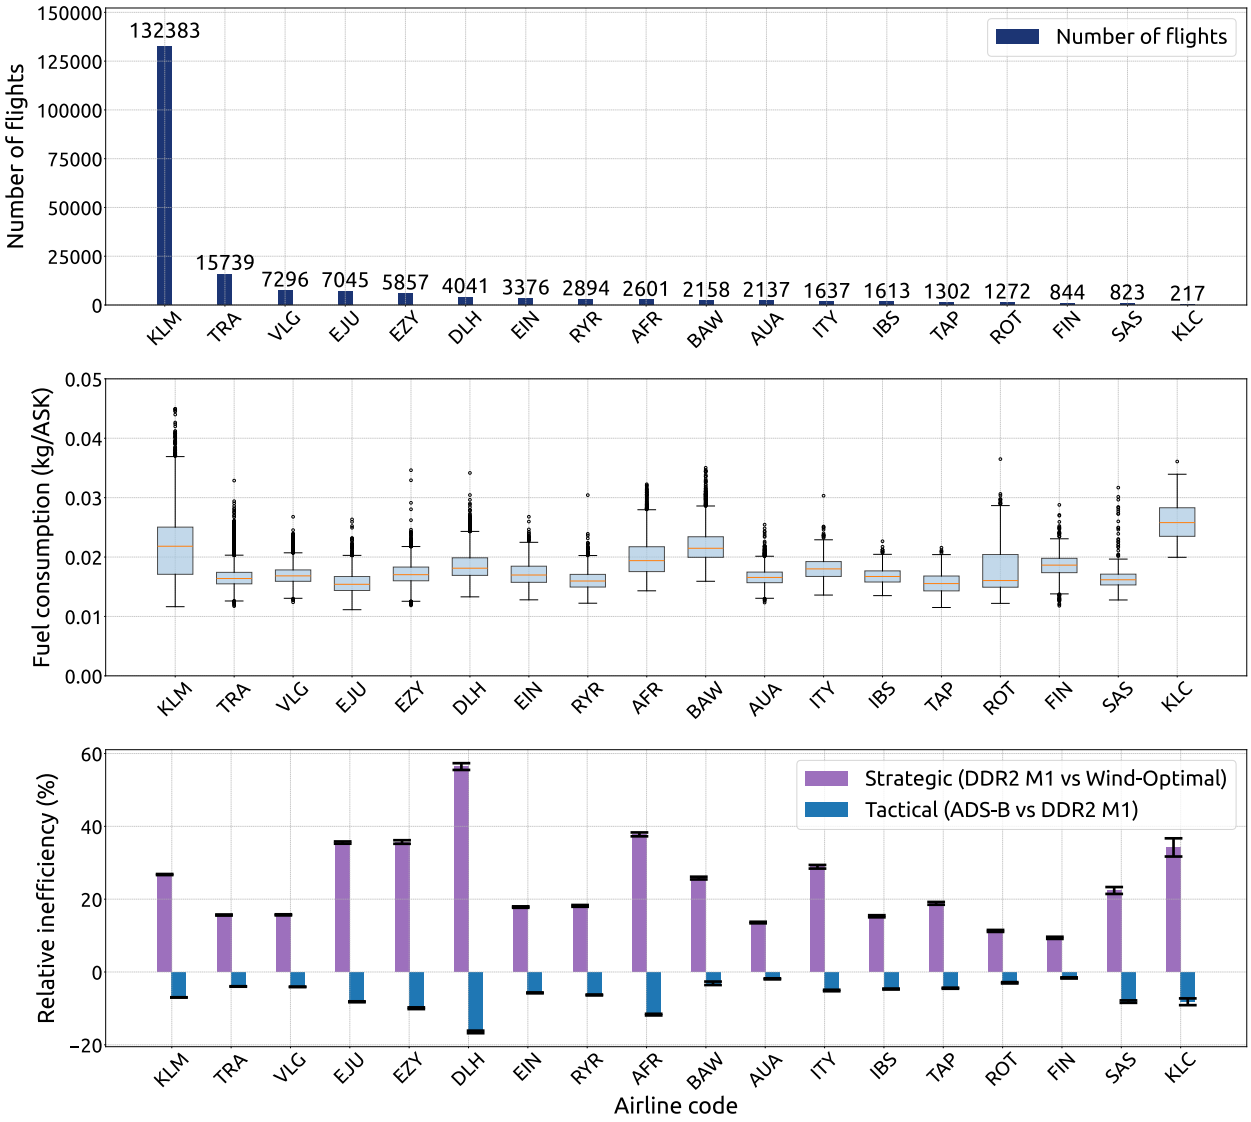

Fig. 10: Airline-specific fuel consumption and relative inefficiencies.

Airline-specific analysis (Figure 10) shows variations in operational scale and normalized fuel consumption, where Available Seat Kilometer (ASK) is computed from great-circle distances and airline–aircraft seating configurations sourced from fleet databases. KLM accounts for 132,383 flights (68.5% of the dataset), followed by Transavia with 15,739 flights. Other major operators, including easyJet (EJU, EZY), Vueling (VLG), Lufthansa (DLH), and Aer Lingus (EIN), each contribute 3,000–7,000 flights. The median fuel consumption varies across airlines, ranging from 0.015 kg/ASK (easyJet Europe) to 0.026 kg/ASK (KLC). Low-cost carriers cluster around 0.015–0.017 kg/ASK, indicating relatively low and consistent fuel use, while network carriers show broader variations, with median values of 0.016–0.022 kg/ASK and wider interquartile ranges. The mean value of strategic inefficiency (the planned versus optimal fuel penalty) ranges from about 10% to 57%, reflecting large differences across airlines. Lufthansa exhibits the highest mean at 56.4% followed by easyJet variants (EZY at 35.7%, EJU at 35.5%), while KLM's mean is 26.8% and other low-cost carriers such as Vueling, Transavia, and Ryanair lie around 16–18%. Tactical inefficiency is negative across all airlines, indicating that executed flight trajectories consume less fuel than their planned counterparts.

C. City-pair analysis

Figure 11 quantifies the relative total fuel inefficiency for the 20 busiest city pairs, comparing executed trajectories from ADS-B and DDR2 M3 sources against wind-optimal references. The relative inefficiency, expressed as a percentage of the optimal fuel consumption, reveals differences across routes. For inbound flights, the majority of routes exhibit inefficiencies between 7-20% for ADS-B trajectories. Relatively longer routes such as flights from LPPT, LIRF, LEMD, ENGM, show low to moderate inefficiencies in the range of 7-13%, while short-haul connections display a greater variation with flights from EIDW having one of the lowest relative inefficiencies of 8%, to 39% from EDDF.

Outbound flights from EHAM exhibit generally higher inefficiencies compared to their inbound counterparts. Similar to the inbound, flights to EDDF indicate high relative inefficiency of approximately 28%. Other significant inefficiencies are observed for flights to LSZH and LFPG. Flights to EGLL show inefficiency of 26% in the ADS-B dataset, compared to approximately 9% for nearby EGLC. Notably, the DDR2 M3 dataset indicates 47% inefficiency for the same EGLL route, a substantial discrepancy that requires further investigation.

The 95% confidence intervals, shown as error bars, are relatively narrow for most routes due to large sample sizes of high-frequency routes, which reduce uncertainty in the estimates despite flight to flight variability.

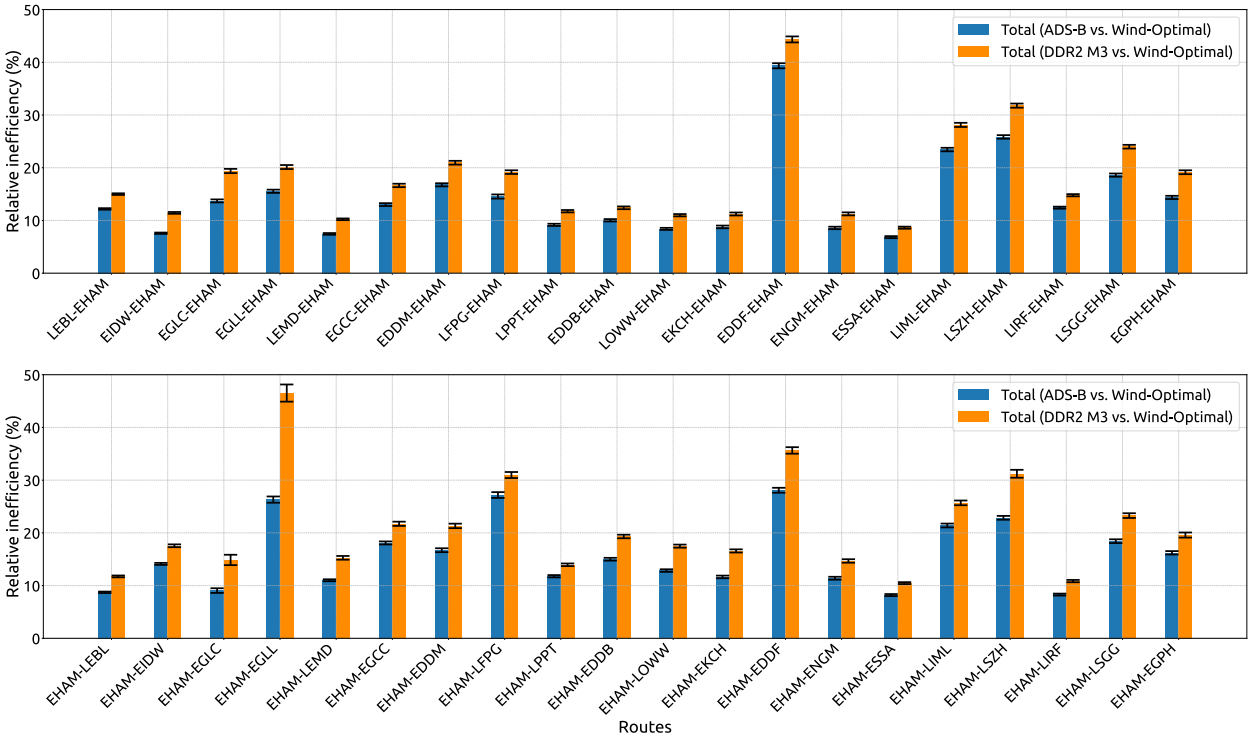

Fig. 11: Relative total fuel inefficiencies across high frequency routes for flights inbound and outbound EHAM.

Figure 12 presents the decomposition of the total inefficiency into strategic and tactical components. For inbound operations, strategic inefficiency is consistently higher than tactical, ranging from 8% to 61%. The highest strategic inefficiency is observed for flights from EDDF, reaching 61%, while the least inefficient route is from ESSA at 8%. In contrast, tactical inefficiency shows predominantly negative values for most routes, with a minimum of -17%. A representative inbound flight from EDDF is presented in Figure 13.

For outbound operations, strategic inefficiency similarly dominates for the majority of city-pairs, with values ranging from 7% (to EGLC) to 100% (to EDDF). The EDDF route represents an extreme case where the fuel penalty in the strategic phase approaches the reference fuel consumption. The tactical component predominantly exhibits negative values (17 out of 20 routes), ranging from -72% to -20% for most routes, again indicating execution improvements. However, three routes demonstrate positive tactical inefficiency: flights to EGLL (5.04%), EGLC (1.87%), and EDDB (0.23%), suggesting that actual flight operations on these routes negatively deviate from planned trajectories leading to increased fuel consumption. An example flight to EGLL is presented in Figure 14.

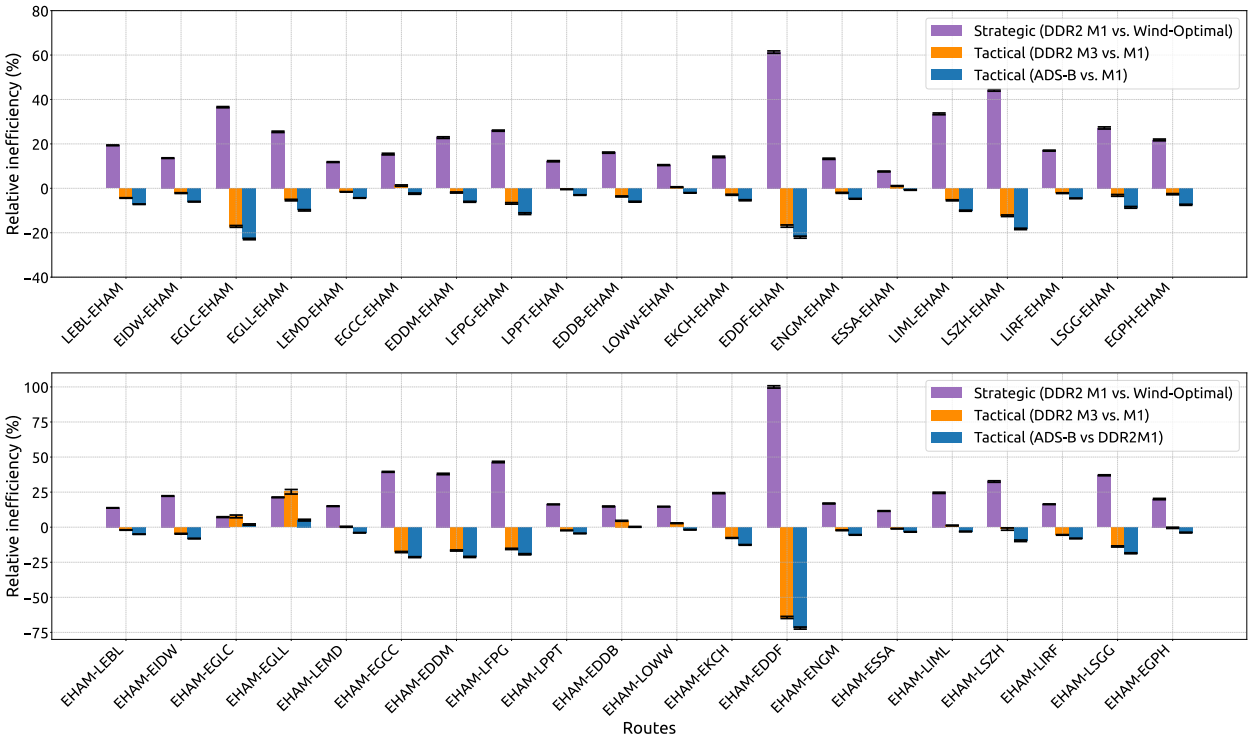

Fig. 12: Decomposition of relative total inefficiency into strategic and tactical components across high frequency city-pairs.

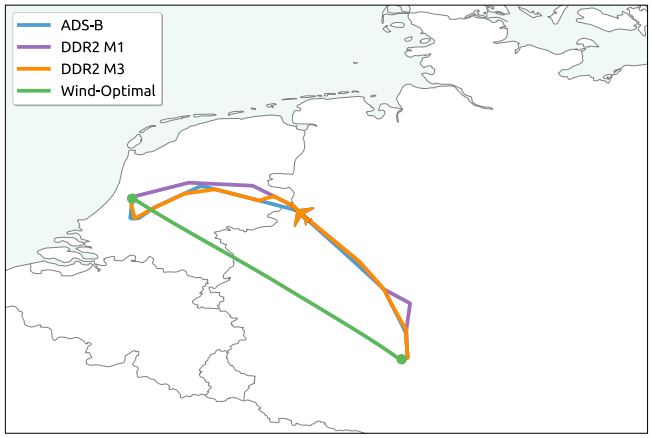

(a) Comparison of executed, planned, and optimal flight paths.

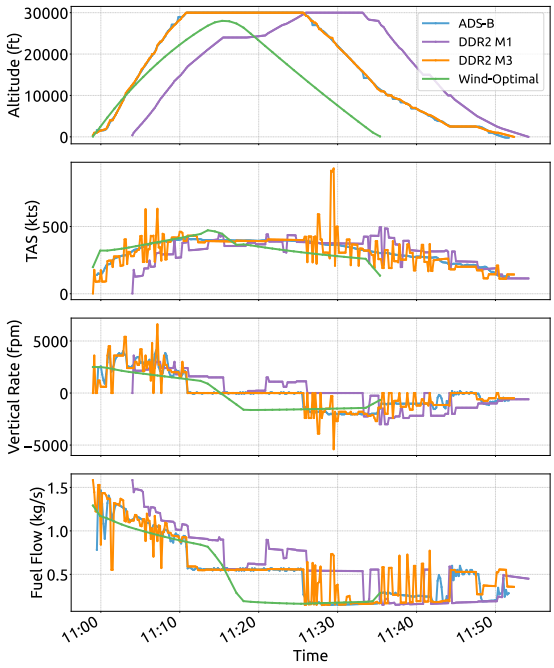

(b) Altitude, speed, vertical rate and fuel flow profiles.

Fig. 13: An example inbound flight from EDDF with a relative total inefficiency of 32%, strategic inefficiency of 54% and tactical inefficiency of -22%.

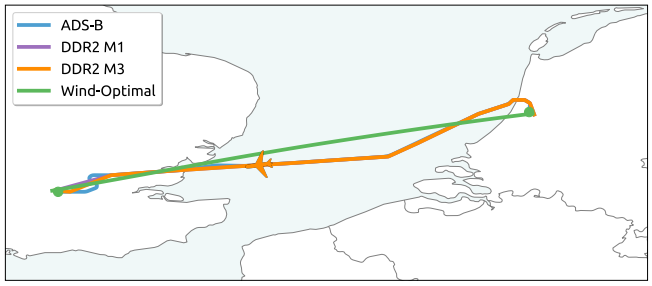

(a) Comparison of executed, planned, and optimal flight paths.

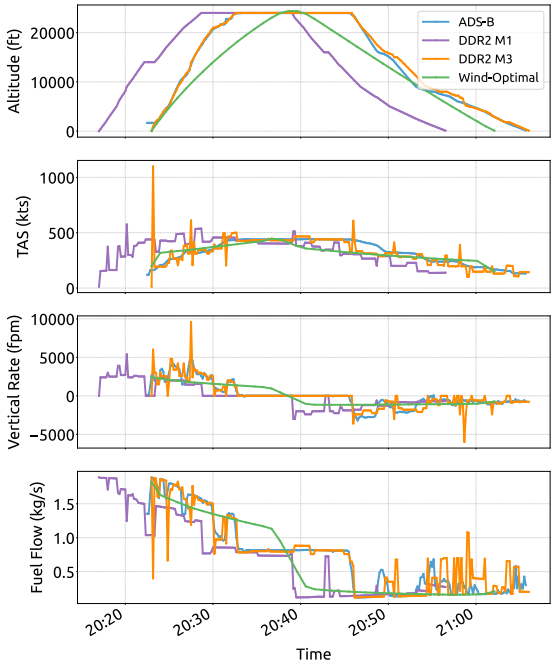

(b) Altitude, speed, vertical rate and fuel flow profiles.

Fig. 14: An example outbound flight to EGLL with a relative total inefficiency of 21.2%, strategic inefficiency of 18.5% and tactical inefficiency of 3%.

The absolute fuel penalty represented in Figure 15 offers a more direct visualization of the environmental impact. Specifically, it displays a theoretical benefit pool if planned trajectories followed optimal routes. For inbound flights, strategic inefficiency accounts for an average of 200-720 kg of excess fuel per flight, while tactical execution typically recovers up to 300 kg through efficient real-time operations. Similar to the relative inefficiency, outbound flights display generally higher strategic inefficiencies, with flights to EDDF for instance accounting for 852 kg of excess fuel per flight. The most notable exception where tactical interventions increase fuel consumption are observed for outbound flights to EGLL where fuel penalties of up to 55 kg in the ADS-B dataset beyond the filed plan are incurred.

Longer routes exhibit larger absolute fuel penalties even when relative inefficiency percentages remain moderate, highlighting how moderate deviations from optimal trajectories accumulate into substantial absolute fuel waste. An example of this is presented by flights to and from LPPT and LIRF.

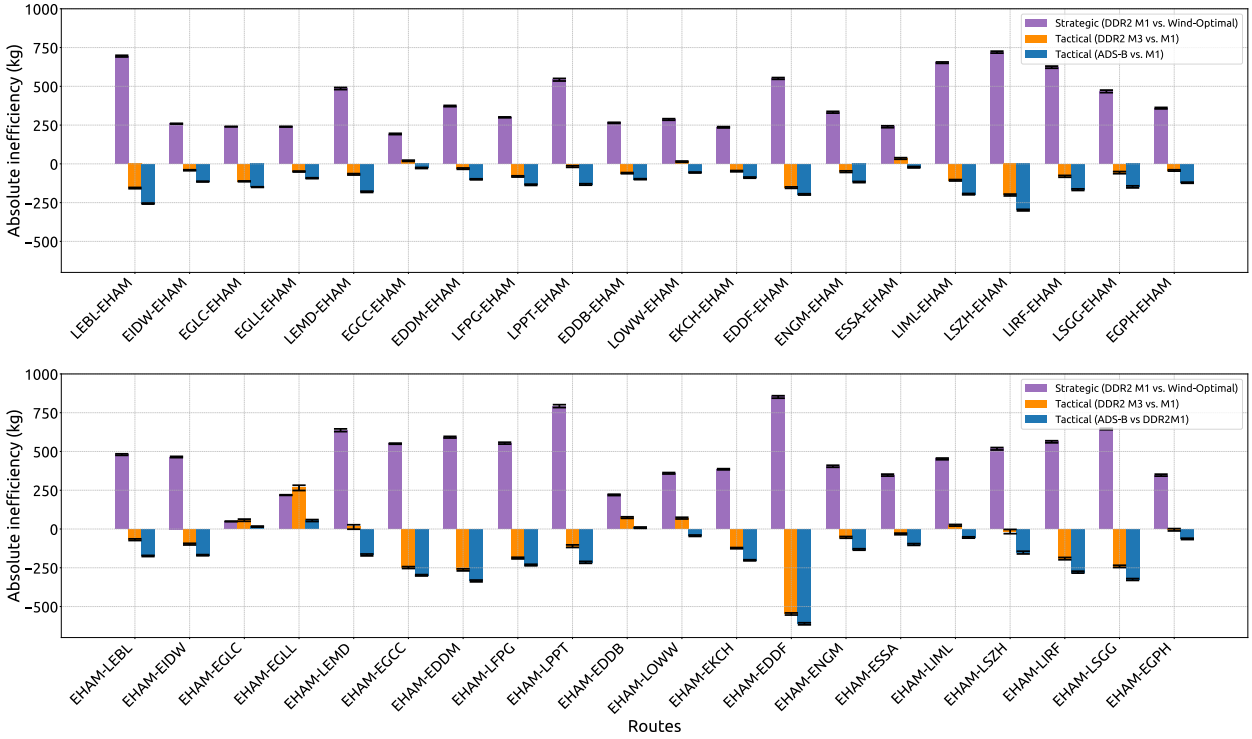

Fig. 15: Decomposition of absolute total inefficiency into strategic and tactical components across high frequency city-pairs.

D. FIR Analysis

Inefficiencies present within the Amsterdam FIR are analyzed by decomposing total inefficiency into total horizontal and total vertical components for climb and descent phases. This analysis focuses on flight segments between 10,000 ft and the FIR boundary crossing point, providing insights for ANSP performance monitoring. For this FIR analysis, ADS-B data was selected to represent executed trajectories, as the combination of validated fuel consumption performance against QAR data, good coverage above 10,000 ft in the Netherlands, and continuous trajectory representation without DDR2's segment-based update criteria provides more suitable data for the detailed profile analysis required to isolate horizontal and vertical inefficiency components.

1) Climb: The climb phase analysis reveals moderate inefficiencies throughout 2024, with average total relative inefficiencies ranging from 9.13% in January to 13.8% in July (median: 8.27% to 13.1%), corresponding to absolute fuel penalties of 57-84 kg per flight (median: 55-83 kg). A seasonal pattern is observed, with winter months having lower inefficiencies compared to summer months. Horizontal inefficiency accounts for approximately 90% of total climb inefficiency throughout the year. The average relative horizontal component ranges from 8.43% to 11.9% (median: 7.72% to 11.2%), translating to absolute fuel penalties of 54-73 kg per flight (median: 51-71 kg). Vertical inefficiency remains modest, with average values of 0.70-1.97% (median: 1.14-2.15%) or 3-12 kg per flight (median: 8-13 kg), with increased inefficiency observed during summer months.

Figure 16 and Figure 17 demonstrates a representative example for an outbound EHAM-LPPT climb. Figure 16 compares executed (760 kg) against the optimal (687 kg), revealing 10.7% (73 kg) total inefficiency. The flight path in Figure 16a shows minimal lateral deviation, with both paths appearing nearly collinear. The comparison of profiles in Figure 16b shows both climbing from 10,000 to approximately 25,000 ft over different durations, with the optimal climb maintaining a consistent vertical rate of 2,500 fpm and a higher TAS compared to the executed climb. Figure 17 presents the constrained optimal (747 kg), isolating vertical inefficiency at 1.9% (13.2 kg). Horizontal paths and TAS profiles match by design (Figure 17a). The vertical rate comparison (Figure 17b) shows the constrained optimal achieving a consistent climb rate of 2,500 fpm for majority of the climb.

TABLE IV: Average climb inefficiencies decomposed into their total horizontal and total vertical components.

(a) Relative climb inefficiencies.

Month	ΔF_T^h **(%)**	ΔF_T^v**(%)**	ΔF_T **(%)**
January	8.43	0.70	9.13
February	8.81	0.95	9.76
March	9.87	0.98	10.8
April	10.6	1.28	11.9
May	10.8	1.58	12.4
June	11.2	1.97	13.1
July	11.9	1.84	13.8
August	10.8	1.69	12.5
September	11.6	1.53	13.1
October	10.5	1.18	11.7
November	9.47	0.95	10.4
December	9.33	1.27	10.6

(b) Absolute climb inefficiencies.

Month	ΔF_T^h **(kg)**	ΔF_T^v **(kg)**	ΔF_T **(kg)**
January	53.9	3.19	57.1
February	56.3	4.76	61.0
March	63.2	4.89	68.0
April	66.8	7.25	74.1
May	68.6	9.73	78.4
June	69.7	11.9	81.5
July	73.2	11.0	84.3
August	66.2	10.0	76.2
September	73.0	8.90	81.9
October	65.5	6.67	72.2
November	60.0	5.05	65.0
December	59.5	7.57	67.1

TABLE V: Median climb inefficiencies decomposed into their total horizontal and total vertical components.

(a) Relative climb inefficiencies

Month	ΔF_T^h **(%)**	ΔF_T^v**(%)**	ΔF_T **(%)**
January	7.72	1.14	8.27
February	8.14	1.50	8.85
March	9.22	1.50	10.2
April	10.1	1.71	11.2
May	10.1	1.84	11.6
June	10.7	2.04	12.7
July	11.2	2.10	13.1
August	10.0	2.15	11.8
September	11.0	2.00	12.7
October	9.85	1.67	11.0
November	8.88	1.41	9.79
December	8.59	1.47	9.75

(b) Absolute climb inefficiencies

Month	ΔF_T^h **(kg)**	ΔF_T^v **(kg)**	ΔF_T **(kg)**
January	51.1	7.51	55.4
February	54.3	9.47	59.5
March	61.6	9.60	68.4
April	65.4	11.0	73.7
May	66.1	11.7	77.4
June	68.2	12.5	81.1
July	71.3	12.8	82.6
August	64.2	13.0	76.2
September	71.2	12.5	82.3
October	64.4	10.4	72.2
November	58.7	9.05	65.9
December	58.2	9.59	66.0

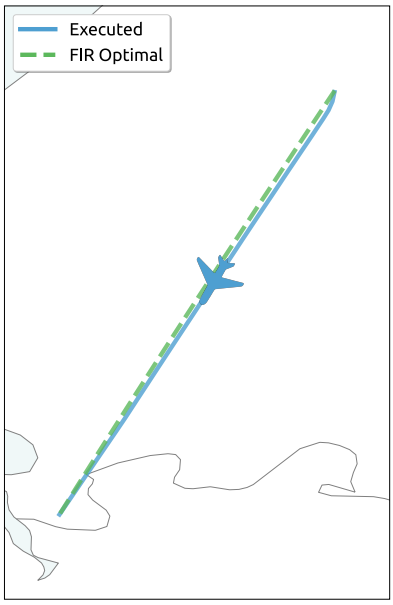

(a) Comparison of executed and optimal flight paths within the FIR.

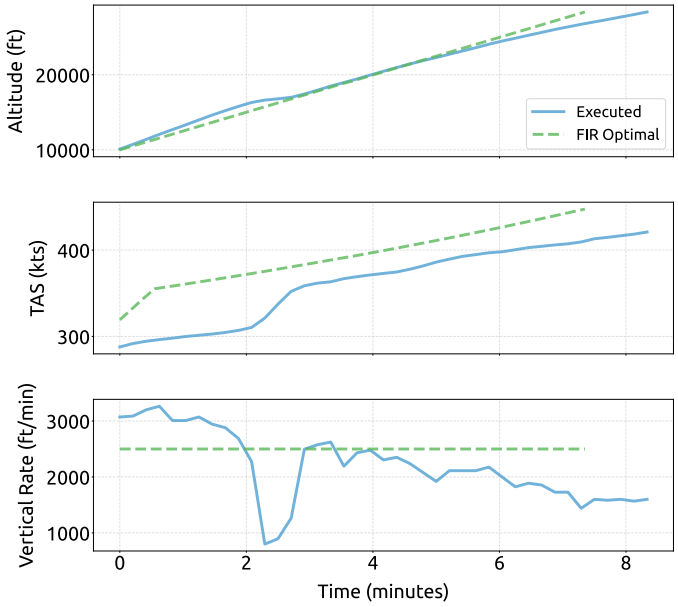

(b) Altitude, TAS and vertical rate profiles.

Fig. 16: An example outbound flight to LPPT with a relative total inefficiency of 10.7%.

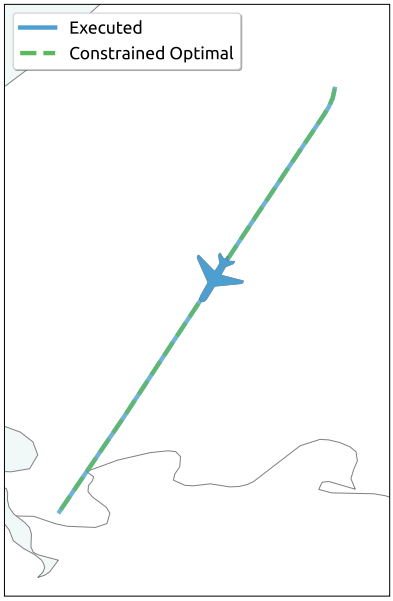

(a) Comparison of executed and optimal flight paths within the FIR.

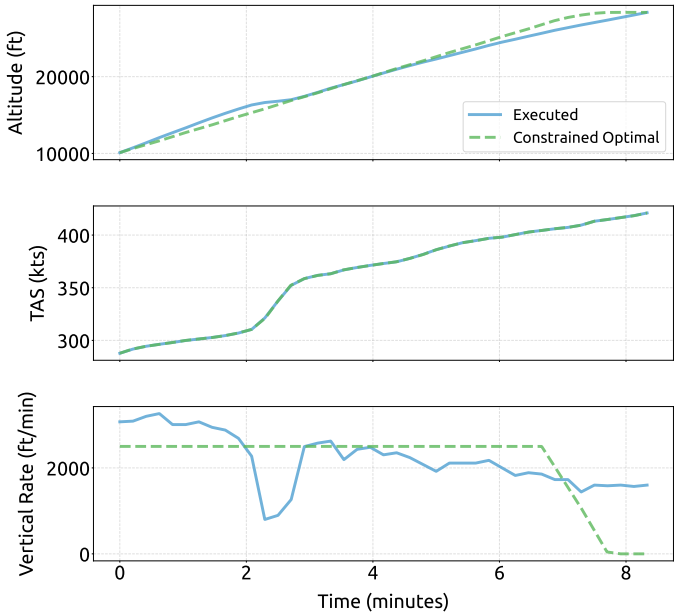

(b) Altitude, TAS and vertical rate profiles.

Fig. 17: An example outbound flight to LPPT with a relative horizontal inefficiency of 8.8% and vertical inefficiency of 1.9%.

2) Descent: The descent phase exhibits higher relative inefficiencies than climb, with average total values ranging from 24.3% in January to 30.2% in August (median: 23.1% to 28.5%), corresponding to 19.0-26.5 kg per flight (median: 20.9-25.4 kg). In contrast to the climb phase, vertical inefficiency contributes more to total inefficiency than horizontal for the majority of months. The average vertical component remains relatively stable throughout the

year at approximately 14.5-16.0% (median: 12.6-13.8%) or 14.4-16.7 kg per flight (median: 12.8-14.7 kg), while the horizontal component shows greater seasonal variation, ranging from 9.18% in March to 15.7% in August (median: 7.83% to 14.0%) or 3.62-12.2 kg per flight (median: 7.28-12.6 kg), representing approximately 52% of the total descent inefficiency.

Figure 18 and Figure 19 demonstrate the decomposition for an inbound LFPG-EHAM descent. Figure 18 compares executed (272) kg) against the optimal (145 kg), revealing 87.5% (127 kg) total inefficiency. Lateral path deviation from the optimal trajectory is observed in Figure 18a. The profile comparison (Figure 18b) shows both aircraft descending from 20,000 to 10,000 ft over different durations, with the optimal maintaining a consistent descent rate of 1,500 fpm and a higher TAS compared to the executed descent. Figure 19 presents the constrained optimal (255 kg), isolating vertical inefficiency at 11.7% (17.4 kg). Horizontal paths match by design (Figure 19a). The profile comparison (Figure 19b) shows the constrained optimal achieving a consistent descent rate of 1000 fpm throughout the descent.

TABLE VI: Average descent inefficiencies decomposed into their total horizontal and total vertical components.

(a) Relative descent inefficiencies

Month	ΔF_T^h **(%)**	ΔF_T^v**(%)**	ΔF_T **(%)**
January	9.79	14.5	24.3
February	10.5	14.6	25.1
March	9.18	15.6	24.8
April	10.1	16.0	26.0
May	12.6	15.2	27.8
June	13.9	14.5	28.3
July	13.4	14.6	28.0
August	15.7	14.5	30.2
September	13.5	15.1	28.6
October	11.2	14.6	25.8
November	11.3	15.2	26.5
December	13.0	14.8	27.8

(b) Absolute descent inefficiencies

Month	ΔF_T^h **(kg)**	ΔF_T^v **(kg)**	ΔF_T **(kg)**
January	4.24	14.8	19.0
February	6.66	14.9	21.5
March	3.62	15.5	19.1
April	5.45	16.7	22.1
May	9.48	14.9	24.4
June	7.99	14.4	22.4
July	9.98	14.8	24.8
August	12.2	14.4	26.5
September	9.94	14.9	24.9
October	7.51	15.0	22.5
November	6.43	15.0	21.4
December	8.25	14.6	22.8

TABLE VII: Median descent inefficiencies decomposed into their total horizontal and total vertical components.

(a) Relative descent inefficiencies

Month	ΔF_T^h **(%)**	ΔF_T^v**(%)**	ΔF_T **(%)**
January	8.89	13.2	23.1
February	8.95	13.1	23.8
March	7.83	13.8	23.7
April	8.61	13.7	24.2
May	11.2	13.3	26.2
June	12.8	12.8	27.7
July	11.8	12.9	26.4
August	14.0	12.7	28.5
September	11.7	13.1	26.7
October	9.74	12.6	24.0
November	10.3	12.8	24.7
December	11.8	13.1	26.8

(b) Absolute descent inefficiencies

Month	ΔF_T^h **(kg)**	ΔF_T^v **(kg)**	ΔF_T **(kg)**
January	8.15	13.6	20.9
February	8.15	13.8	21.5
March	7.28	14.1	21.1
April	8.28	14.7	22.6
May	10.4	13.4	23.9
June	11.7	13.1	24.9
July	10.9	13.2	24.3
August	12.6	12.8	25.4
September	10.8	13.6	25.0
October	8.90	13.4	22.3
November	9.66	13.3	23.0
December	10.9	13.0	24.5

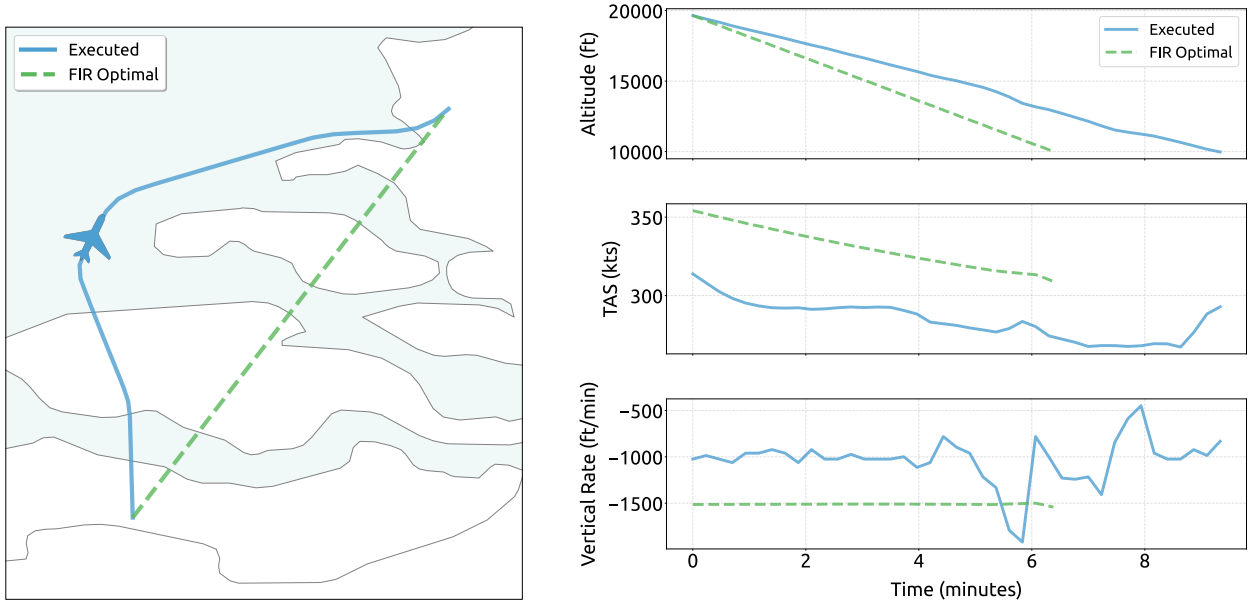

(a) Comparison of executed and optimal flight paths within the FIR.

(b) Altitude, TAS and vertical rate profiles

Fig. 18: An example inbound flight from LFPG with a relative total inefficiency of 87.5%.

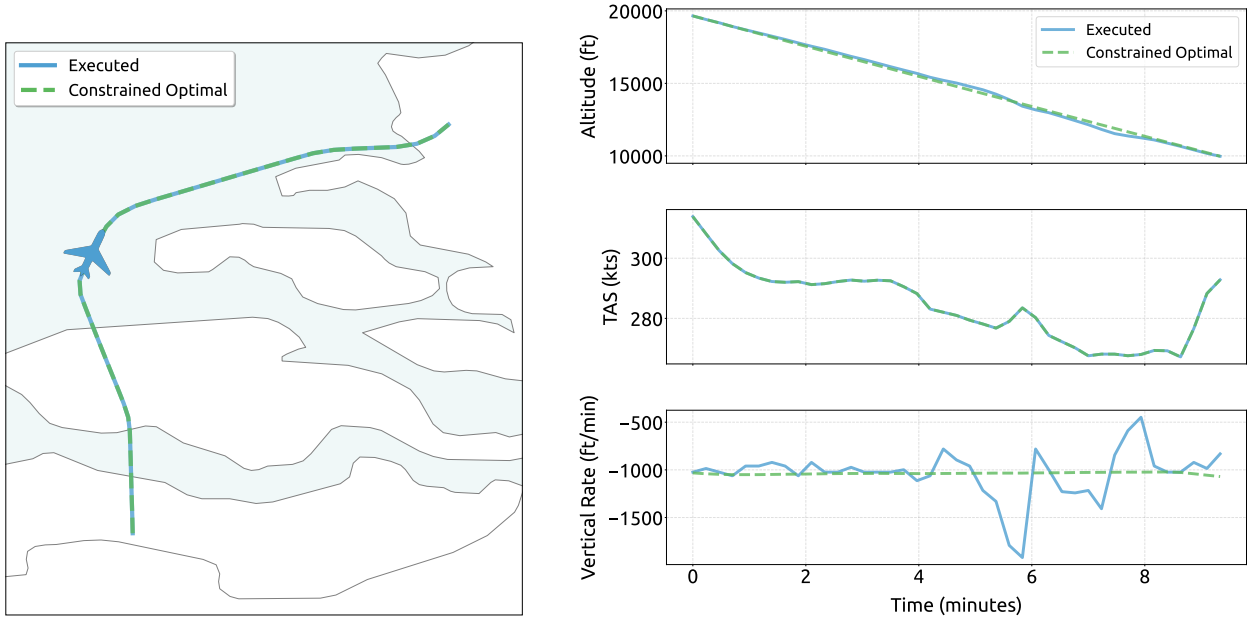

(a) Comparison of executed and optimal flight paths within the FIR.

(b) Altitude, TAS and vertical rate profiles.

Fig. 19: An example inbound flight from LFPG with a relative horizontal inefficiency of 75.9% and vertical inefficiency of 11.7%.

VII. Discussion

A. City-pair and FIR Inefficiency

The city-pair analysis shows that strategic inefficiency consistently dominates total inefficiency for most routes (ranging from 7% for EGLC to 100% for EDDF), indicating that filed flight plans systematically deviate from wind-optimal trajectories due to pre-tactical constraints such as airspace restrictions. The predominantly negative tactical inefficiency across most routes (ranging from -72% to -20%) indicates execution improvements through ATC interventions and real-time wind exploitation. However, three outbound routes (EGLL, EGLC, EDDB) exhibit positive tactical inefficiency of 5.04%, 1.87%, and 0.23% respectively.

The EHAM-EDDF route represents an extreme case where outbound strategic inefficiency reaches 100% of the reference fuel consumption, translating to an average 852 kg of excess fuel per flight. This anomaly stems from planned trajectories that include extended level flight segments during climb and descent phases. In actual operations, flights execute continuous climbs without the planned level segments, resulting in large negative tactical inefficiency values. The low total fuel consumption on this short route also amplifies the relative impact of these inefficient vertical profile choices.

The discrepancy observed in the relative total inefficiencies for EHAM-EGLL outbound flights (26% inefficiency in ADS-B versus 47% in DDR2 M3) is caused by differences in data completeness between the two sources. The update criteria of DDR2 M3 data are consistently unmet for flights to EGLL, resulting in incomplete trajectory information and artificially inflated fuel flow values when vectoring, tromboning, or holding patterns are not captured, as illustrated in Figure 20. In extreme cases, this data incompleteness can lead to inefficiency values exceeding 350% in the DDR2 M3 dataset (Figure 21). The ADS-B data, which provides continuous position updates, presents a more accurate representation of executed flight trajectories for this city pair.

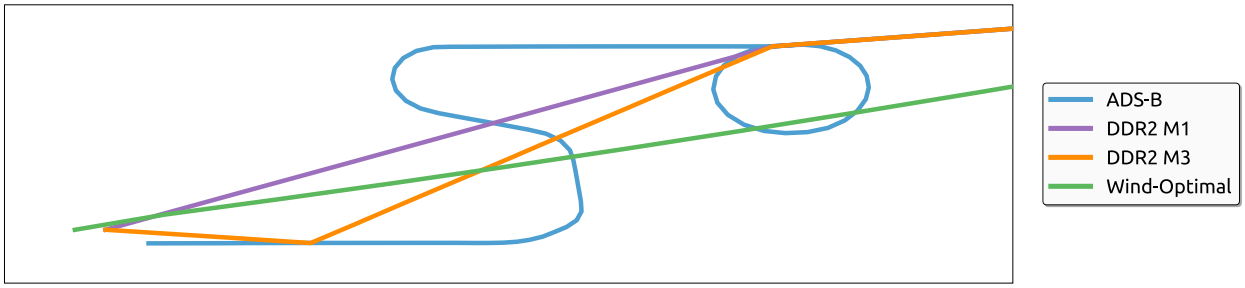

Fig. 20: Flight trajectory for representative EHAM-EGLL flight showing final descent phase into EGLL. Relative total inefficiency: 25.7% (ADS-B) versus 46.5% (DDR2 M3).

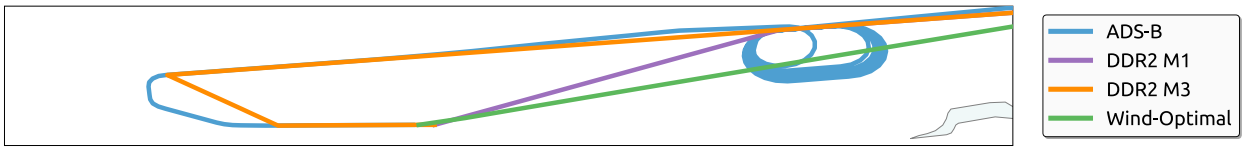

Fig. 21: Flight trajectory for extreme case EHAM-EGLL flight showing final descent phase into EGLL. Relative total inefficiency: 114% (ADS-B) versus 353% (DDR2 M3).

The FIR analysis reveals differences in inefficiency between climb and descent phases. In climb operations, horizontal inefficiency accounts for approximately 90% of total inefficiency throughout the year (8.43-11.9%), indicating that lateral routing constraints are the primary sources of climb inefficiency, while vertical inefficiency remains modest (0.70-1.97%). Descent operations present a contrasting pattern: vertical inefficiency contributes

a larger share at approximately 48-52% during summer months and 53-60% during winter, with the horizontal component showing greater seasonal variation from 9.18% in March to 15.7% in August. The magnitude difference between climb (9.13-13.8%) and descent (24.3-30.2%) total relative inefficiencies is amplified by the lower baseline fuel consumption during descent operations; the higher relative percentages reflect deviations from a smaller reference value rather than proportionally larger absolute fuel penalties.

B. Uncertainty from Initial Mass Estimation

Initial mass estimation generally represents one of the largest sources of uncertainty in trajectory optimization and fuel consumption modeling. Due to overestimation in the initial mass estimation stage, approximately 50% of the flights in the dataset were constrained to an initial mass ratio of 0.85. Figure 22 quantifies the sensitivity of fuel consumption to the initial mass through two flights of different aircraft types and durations. For the B738 flight from EPWA to EHAM (Figure 22a), variations in initial mass ratios from 0.70 to 1.00 of the MTOW produces fuel consumption deviations ranging from -13.5% to +13.3%, corresponding to absolute differences of -570 kg to +563 kg. The E190 short-haul flight from EHAM to EDDW (Figure 22b) demonstrates a smaller variation of -6.6% to +6.8% relative to the 0.85 baseline. Across the broader sample, maximum variations of up to 15% relative to the baseline were observed.

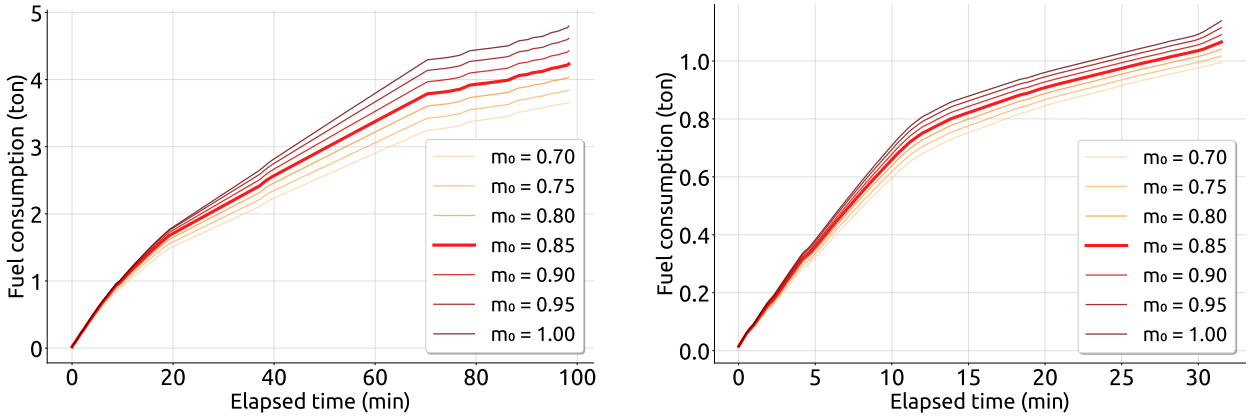

(a) B738 flight from EPWA to EHAM.

(b) E190 flight from EHAM to EDDW.

Fig. 22: Variations in fuel consumption comparison due to different initial mass ratios.

Uncertainties in initial mass also propagate to the optimized trajectories, as lighter aircraft achieve higher cruise altitudes and consequently improved fuel performance. The magnitude of this propagation to the aggregate inefficiency metrics is quantified in Table VIII for the same B738 flight from EPWA to EHAM across initial mass ratios from 0.70 to 0.85. The city-pair analysis reveals that total inefficiency decreases from 8.5% to 4.0% as the mass ratio increases, primarily due to variations in strategic inefficiency which declines from 12.1% to 7.2%. Similarly, the FIR analysis demonstrates that total inefficiency decreases from 25.4% to 17.6%. This reduction is attributable to horizontal inefficiency, which decreases from 21.8% to 10.6%, while vertical inefficiency increases from 3.6% to 7.0%.

This inverse relationship reflects the mass-dependent nature of inefficiency results for descents. Higher initial mass assumptions decrease strategic and horizontal inefficiency while increasing tactical and vertical inefficiency. Since true initial masses are predicted to exceed the constrained 0.85 ratio, reported inefficiencies represent conservative upper bounds, with actual total relative inefficiencies potentially 15-30% lower than the values presented in this analysis.

TABLE VIII: Influence of initial mass estimates on inefficiency metrics, by city-pair and FIR analysis.

(a) City-pair analysis

	m0			
Metric	**0.7**	**0.75**	**0.8**	**0.85**
$\hat{F}_e$ (kg)	3077	3238	3407	3581
F^* (kg)	2836	3015	3222	3445
$\hat{F}_{\text{RBT}}$ (kg)	3180	3341	3512	3692
ΔF_s (%)	12.1	10.8	9.0	7.2
ΔF_t (%)	-3.6	-3.4	-3.3	-3.2
ΔF_T (%)	8.5	7.4	5.7	4.0

(b) FIR analysis

	m0			
Metric	**0.7**	**0.75**	**0.8**	**0.85**
$\hat{F}_e$ (kg)	84.4	81.5	79.3	77.6
F^* (kg)	67.3	65.9	65.9	66.0
F_e^* (kg)	82.0	79.0	76.0	73.0
ΔF_T^h (%)	21.8	19.9	15.3	10.6
ΔF_T^v (%)	3.6	3.8	5.0	7.0
ΔF_T (%)	25.4	23.7	20.3	17.6

C. *Influence of Wind on Fuel Consumption*

The influence of wind integration on fuel flow calculations for a flight operated from ENBR to EHAM experiencing tailwind is presented in Figure 23. The comparison between fuel flow computed with and without wind effects reveals the greatest differences during the cruise phase, with variations of up to 0.3 kg/s. The aggregate of these differences can amount to several hundred kilograms of fuel per flight, comparable to or exceeding the inefficiency margins quantified in the city-pair analysis. In the represented flight, the inclusion of tailwind effects yielded a reduction in fuel consumption of 12.6% (520 kg). This finding reinforces the limitation of proxy indicators that do not account for meteorological effects or fuel-based inefficiency studies that omit them.

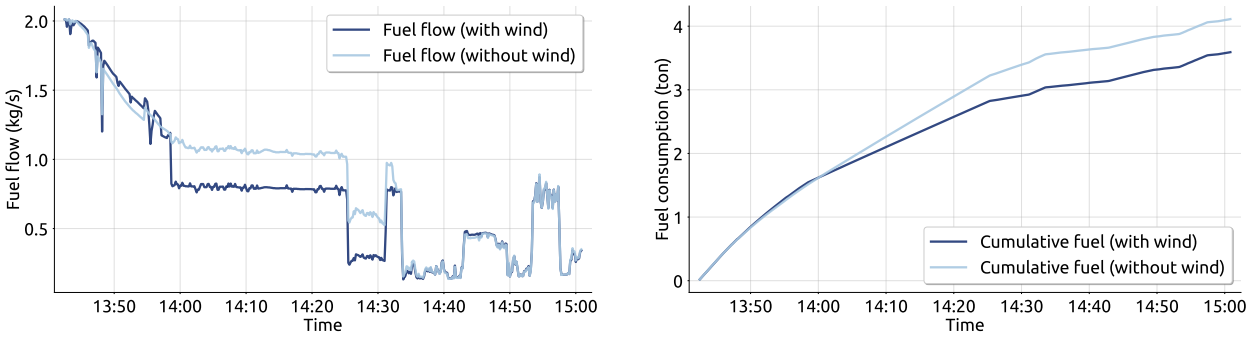

(a) Fuel flow

(b) Cumulative fuel consumption

Fig. 23: Influence of tailwind on fuel flow and total fuel consumption for a flight from ENBR to EHAM on February 2nd 2024.

Table IX quantifies the impact of wind integration on inefficiency metrics for the same ENBR to EHAM flight. The city-pair analysis reveals that omitting wind effects increases total inefficiency from 15.8% to 30.3%, nearly doubling the calculated value. Strategic inefficiency increases from 21.1% to 29.3%, while tactical inefficiency increases from -5.3% to 1.0%. In the descent to EHAM, the relative total inefficiency is underestimated when wind is excluded: total inefficiency decreases from 14.3% to 5.3%, a reduction of 63%. The contrasting behavior between city-pair and FIR analysis reflects the different wind conditions perceived during different flight phases, yet both demonstrate that wind integration is necessary for accurate fuel-based analysis.

TABLE IX: Influence of wind on inefficiency metrics, by city-pair and FIR analysis.

(a) City-pair analysis

Metric	With wind	Without wind
$\hat{F}_e$ (kg)	2560	3203
F^* (kg)	2211	2458
$\hat{F}_{\text{RBT}}$ (kg)	2678	3179
ΔF_s (%)	21.1	29.3
ΔF_t (%)	-5.3	1.0
ΔF_T (%)	15.8	30.3

(b) FIR analysis

Metric	With wind	Without wind
$\hat{F}_e$ (kg)	248	276
F^* (kg)	217	262
F_e^* (kg)	226	272
ΔF_T^h (%)	4.1	3.8
ΔF_T^v (%)	10.1	1.5
ΔF_T (%)	14.3	5.3

D. Data Quality and Dependency Limitations

Several data quality issues affected the completeness of the dataset. Missing DDR2 files for June 18 resulted in gaps in the dataset for that period. ADS-B contained discontinuities, with unusable data observed in May 10, May 21 and June 26, and missing data in May 20, resulting in gaps that potentially introduce sampling bias if coinciding with atypical operational conditions.

A key limitation of DDR2-based trajectory analysis stems from inherent and artificially level segments during climb and descent which produces stepped altitude profiles that do not reflect the continuous executed trajectories. When vertical rates are derived, the resulting time series exhibits discontinuities that propagate into fuel flow calculations, which contributes to the elevation of DDR2 M3 fuel consumption estimates relative to ADS-B observations throughout the analysis.

The framework's interdependency between DDR2 and OpenSky data limits accessibility. DDR2 files are required to query ADS-B data while ADS-B metadata is required to process DDR2 data. This mutual dependency prevents implementation by stakeholders or researchers with access to only one source.

E. Trajectory Optimization Limitations

The city-pair analysis employed airport-to-airport trajectory generation followed by fuel aggregation above 10,000 ft, to simplify implementation. To evaluate the effect of boundary selection on wind-optimal fuel consumption, altitude-constrained optimization (10,000 ft to 10,000 ft) was performed for approximately 100 flights using measured aircraft mass and state parameters at the actual crossing points as boundary conditions. This altitude-based approach yields optimal fuel estimates 8.64% higher than the conventional airport-based method, demonstrating that airport-to-airport optimization produces measurable and overly optimistic fuel estimates.

The constrained optimizer developed for decomposing horizontal and vertical inefficiencies in the FIR analysis is subject to an important technical limitation. It optimizes only the vertical rate profile, keeping the executed horizontal trajectory and speed profile fixed. Ideally, vertical rate and TAS would be simultaneously optimized along the fixed horizontal path to isolate purely horizontal inefficiency. However, achieving stable convergence when optimizing both control variables simultaneously proved technically challenging, necessitating that the speed profile remain fixed to the executed trajectory.

This limitation has important implications for interpreting the FIR analysis results. The close agreement between executed and constrained-optimal speed profiles shown in Figure 17 and Figure 19 is by design, not evidence that executed speed management was already optimal. In reality, for a given horizontal trajectory, there exists an optimal combination of vertical rate and speed that minimizes fuel consumption. By constraining speed to the executed profile, the optimizer cannot achieve the true fuel-optimal vertical profile for that climb or descent. Consequently, the calculated vertical inefficiency represents a lower bound, and the actual value would likely be higher if speed optimization were included. Conversely, the horizontal inefficiency is likely overestimated, as it

absorbs lateral deviations as well as speed-related inefficiencies present in the executed trajectory. This consideration is particularly relevant when interpreting the horizontal inefficiency results for climbs.

The FIR-Optimal climb and descent optimizers exhibit route-specific limitations in the FIR analysis. The methodology assumes purely climb or descent phases within the FIR, but flights to proximate airports such as Brussels, Düsseldorf or regional flights often climb, cruise and begin their descent within the FIR. Imposing climb constraints on these trajectories produced unrealistically high 'optimal' fuel consumption, necessitating exclusion of approximately 1.5% of flights. Additionally, the optimizer does not implement aircraft-specific vertical rate limits, which has resulted in executed flights consuming less fuel than optimal ones for specific cases.

F. Recommendations for Future Work

The limitations identified in this study indicate several areas for improvement in data processing, modeling, and framework architecture. Recommendations progress from near-term to long-term implementations.

Addressing the data representation challenges inherent in DDR2-based analysis represents an immediate opportunity for improving fuel consumption accuracy. Application of smoothing techniques such as Savitzky-Golay filtering, as implemented in [13], could mitigate the discontinuities that artificially elevate fuel consumption estimates relative to ADS-B data. However, care must be taken to preserve genuine features such as level-off segments during climb and descent. Beyond data processing, the framework's interdependency between DDR2 and OpenSky datasets limits its accessibility to stakeholders with access to only one source. Developing modular architecture that enables independent data source processing would facilitate broader ANSP adoption and support fully open-source research.

The sensitivity analysis presented in this section demonstrates that initial mass uncertainty propagates significantly through fuel consumption calculations and inefficiency metrics, with variations in fuel consumption reaching 15% across the operational mass range. This motivates prioritizing refinement of the coefficients of the mass estimation model, particularly for short-haul operations where the current approach systematically overestimates mass ratios. Given current industry momentum toward TBO, improved transparency in initial mass data may become increasingly feasible in the future.

Complementing improved mass estimation, validation of the OpenAP fuel flow model against Quick Access Recorder data for multiple aircraft types would establish uncertainty bounds for the quantified inefficiencies. The single validation case presented, while demonstrating good agreement, cannot define model performance across the heterogeneous dataset analyzed; access to multiple representative flights per aircraft type would strengthen confidence in the results. This validation effort would be particularly valuable given OpenAP's growing adoption.

The technical limitations in the constrained optimizer reveal opportunities for methodological refinement. Development of an optimizer capable of simultaneously optimizing vertical rate and TAS along fixed horizontal paths would enable accurate decomposition of horizontal and vertical inefficiency components. Related to this, incorporation of route-specific climb and descent constraints would eliminate the need to exclude flights experiencing mixed vertical profiles within the FIR. Additionally, implementation of aircraft-specific vertical rate limits would prevent generation of reference trajectories that can be outperformed by executed flights, a phenomenon observed in some flights in the analysis. Lastly, future studies implementing complete trajectory optimization from 10,000 ft would generate more representative wind-optimal baselines for the city-pair analysis.

As OpenAP's fuel flow models and OpenAP.top's support extend to additional aircraft types, expanding the methodology's type coverage represents a natural extension of this work. Incorporation of more aircraft would allow comprehensive network-wide inefficiency analyses to be conducted.

In the longer term, integration of operational constraints into the trajectory optimizers would improve the practical relevance of the quantified inefficiencies. The current wind-optimal references represent theoretical upper bounds that

assume perfect meteorological knowledge, absence of airspace restrictions, unlimited routing flexibility, and freedom from tactical separation requirements. While valuable for comparative assessment and trend analysis, these idealized references cannot directly inform operational decision-making. Incorporating constraints from Eurocontrol's Route Availability Document or published airspace restrictions would generate operationally feasible reference trajectories that more accurately define the environmental benefit pool genuinely accessible through ATM improvements.

VIII. Conclusion

This research developed a framework for quantifying flight inefficiency through fuel-based metrics. Applied to over 200,000 commercial flights operated at Amsterdam Airport Schiphol throughout 2024, the analysis reveals that strategic inefficiency consistently dominates total inefficiency across most routes, ranging from 7% for operations to EGLC to 100% for the anomalous EDDF route, with the majority of routes exhibiting values between 8-25%. This indicates that systematic deviations originate primarily in flight planning due to pre-tactical constraints such as airspace restrictions rather than execution. Tactical inefficiency is predominantly negative across most routes (ranging from -72% to -20%), demonstrating that real-time operations generally improve planned trajectories through ATC interventions and wind exploitation.

Flight Information Region analysis reveals different inefficiency characteristics between climb and descent phases. In climb operations, horizontal inefficiency accounts for approximately 90% of total inefficiency (8.43-11.9% throughout 2024), while vertical inefficiency remains modest (0.70-1.97%). Descent operations present a contrasting pattern: vertical inefficiency contributes 48-60% of total inefficiency, remaining relatively stable at 14.5-16.0%, while the horizontal component exhibits greater seasonal variation from 9.18% in March to 15.7% in August. The magnitude difference between climb (9.13-13.8%) and descent (24.3-30.2%) total relative inefficiencies is amplified by the lower baseline fuel consumption during descent operations. Seasonal patterns indicate increased inefficiencies during summer months for both phases, suggesting opportunities for seasonal optimization strategies. Phase-specific analysis establishes that cruise represents the highest percentage of fuel consumption, followed closely by climb. Improvements in these phases would positively reflect on inefficiency results.

Wind integration proves essential for accurate fuel-based analysis, with effects exceeding operational inefficiencies. A representative flight experiencing tailwind conditions demonstrated that omitting wind effects nearly doubled calculated city-pair inefficiency (from 15.8% to 30.3%) while underestimating FIR descent inefficiency (from 14.3% to 5.3%). Initial mass estimation uncertainty also propagates through calculations, with variations in fuel consumption reaching up to 15% across the operational mass range. This propagation produces inverse relationships in reported inefficiencies; higher initial mass estimations decrease strategic and horizontal inefficiency while increasing tactical and vertical inefficiency. Since true initial masses are predicted to exceed the constrained 0.85 ratio applied in this analysis, reported inefficiencies represent conservative upper bounds, with actual total relative inefficiencies potentially 15-30% lower.

The framework's limitations indicate several improvement pathways. Development of optimizers capable of simultaneously optimizing vertical rate and TAS along fixed horizontal paths would eliminate any bias in component attribution and provide more accurate decomposition of horizontal and vertical inefficiency. Incorporation of aircraft-specific vertical rate limits and route-specific constraints would prevent generation of unrealistic reference trajectories. Refinement of initial mass estimation coefficients would reduce uncertainty in aggregated results. Long-term integration of operational constraints from airspace restrictions would generate operationally feasible reference trajectories that more accurately characterize the environmental benefit pool genuinely accessible through air traffic management improvements.

The developed framework provides researchers and air navigation service providers with tools for performance monitoring and comparative analysis, enabling data-driven decision-making towards an enhanced ATM network.

DATA AVAILABILITY

The code supporting this study is available at https://github.com/Ltm-01/Quant

REFERENCES

[1] European Commission. (2024) Reducing emissions from aviation. Accessed: Nov. 14, 2024. [Online]. Available: https://climate.ec.europa.eu/eu-action/transport/reducing-emissions-aviation_en

[2] EUROCONTROL, "Environmental assessment: European ATM network fuel inefficiency study," EUROCONTROL, Brussels, Belgium, Environmental Assessment 01-01, Dec. 2020. [Online]. Available: https://www.eurocontrol.int/sites/default/files/2020-12/eurocontrol-european-atm-network-fuel-inefficiency-study.pdf

[3] EUROCONTROL Aviation Intelligence Unit, "Fuel tankering: Economic benefits and environmental impact," EUROCONTROL, Brussels, Belgium, Think Paper 1, Jun. 2019. [Online]. Available: https://www.eurocontrol.int/sites/default/files/2020-01/eurocontrol-think-paper-1-fuel-tankering.pdf

[4] European Union Aviation Safety Agency (EASA) and EUROCONTROL, "Critical review of ATM/ANS environmental performance measurements," EASA and EUROCONTROL, Tech. Rep., January 2023. [Online]. Available: https://ana.gouvernement.lu/dam-assets/publications/eurocontrol/eurocontrol-atm-ans-env-performance-measurements.pdf

[5] E. Calvo, J. Cordero, L. D'Alto, J. López-Leonés, M. Vilaplana, and M. La Civita, "A new method to validate the route extension metric against fuel efficiency," in *Proceedings of the 11th USA/Europe Air Traffic Management Research and Development Seminar ATM 2015*, 2015.

[6] J. Sun, J. M. Hoekstra, and J. Ellerbroek, "OpenAP: an open-source aircraft performance model for air transportation studies and simulations," *Aerospace*, vol. 7, no. 8, p. 104, 2020.

[7] J. Sun, "OpenAP.top: open flight trajectory optimization for air transport and sustainability research," *Aerospace*, vol. 9, no. 7, p. 383, 2022.

[8] T. Reynolds, "Air traffic management performance assessment using flight inefficiency metrics," *Transport Policy*, vol. 34, pp. 63–74, 2014.

[9] G. Guastalla, "Performance indicator - horizontal flight efficiency," EUROCONTROL Performance Review Unit (PRU), Performance Indicator Documentation 01-00, May 2014. [Online]. Available: https://ansperformance.eu/library/pru-hfe.pdf

[10] J. Bronsvoort, P. Zissermann, S. Barry, and G. McDonald, "A framework for assessing and managing the impact of ANSP actions on flight efficiency," in *Proceedings of the 11th USA/Europe Air Traffic Management Research and Development Seminar ATM 2015*, 2015.

[11] J. Leones, M. Morales, L. D'Alto, P. Escalonilla, D. Herrer, M. Bravo, F. Cámara, Mateo, B. MacNamee, S. Wang, A. Grover, and P. Plantholt, "Advanced flight efficiency key performance indicators to support air traffic analytics: assessment of European flight efficiency using ADS-B data," in *37th AIAA/IEEE Digital Avionics Systems Conference*, 2018, pp. 1–10.

[12] T. Reynolds, "Analysis of lateral flight inefficiency in global air traffic management," in *8th AIAA Aviation Technology Integration and Operations (ATIO) Conference*, 2008.

[13] X. Prats, C. Barrado, R. Herranz, D. Delgado, and R. G. M. López, "Identifying the sources of flight inefficiency from historical aircraft trajectories: a set of distance- and fuel-based performance indicators for post-operational analysis," in *Thirteenth USA/Europe Air Traffic Management Research and Development Seminar*, Vienna, Austria, 2019.

[14] J. Rosenow, G. Chen, H. Fricke, and Y. Wang, "Factors impacting Chinese and European vertical flight efficiency," *Aerospace*, vol. 9, no. 2, p. 76, 2022.

[15] P. Pasutto, K. Zeghal, and D. Brain, "Initial analysis of vertical flight efficiency in cruise for European city pairs," in *AIAA Aviation and Aeronautics Forum and Exposition (AIAA Aviation Forum)*, 2021.

[16] A. Lemetti, T. Polishchuk, R. Sáez, and X. Prats, "Analysis of weather impact on flight efficiency for stockholm arlanda airport arrivals," in *Air Traffic Management and Systems IV*, E. N. R. Institute, Ed. Singapore: Springer, 2021, pp. 77–92.

[17] A. Lemetti, T. Polishchuk, R. Saez, and X. Prats, "Evaluation of flight efficiency for Stockholm Arlanda Airport arrivals," in *AIAA IEEE Digital Avionics Systems Conference Proceedings*, vol. 2019-September, 2019.

[18] P. Pasutto, E. Hoffman, and K. Zeghal, "Flight inefficiency in descent: mapping where it happens," in *AIAA Aviation and Aeronautics Forum and Exposition (AIAA Aviation Forum)*, 2021.

[19] R. Dalmau, J. Sun, and X. Prats, "Fuel inefficiency characterisation and assessment due to early execution of top of descents: a case study for Amsterdam-Schiphol terminal airspace using ADS-B data," in *14th USA/Europe Air Traffic Management Research and Development Seminar (ATM)*, 2021.

[20] A. Lazarovski and R. Koelle, "Measuring operational air navigation system performance in the 21st century: arrival management," in *AIAA IEEE Digital Avionics Systems Conference Proceedings*, 2019, pp. 1–8.

[21] A. Belle and L. Sherry, "Estimated fuel burn performance for MDW arrivals," in *2013 Aviation Technology Integration and Operations Conference*, 2013.

[22] A. Harada, T. Kozuka, Y. Miyazawa, N. Wickramasinghe, M. Brown, and Y. Fukuda, "Analysis of air traffic efficiency using dynamic programming trajectory optimization," in *29th Congress of the International Council of the Aeronautical Sciences (ICAS)*, 2014.

[23] M. Ryerson, M. Hansen, and J. Bonn, "Time to burn: flight delay, terminal efficiency, and fuel consumption in the National Airspace System," *Transportation Research Part A: Policy and Practice*, vol. 69, pp. 286–298, 2014.

[24] A. Nuic, D. Poles, and V. Mouillet, "BADA: an advanced aircraft performance model for present and future ATM systems," *International Journal of Adaptive Control and Signal Processing*, vol. 24, no. 10, pp. 850–866, 2010.

[25] K. Krajček Nikolić, P. Papoči, D. Nikolić, and B. Antulov-Fantulin, "Fuel burn method assessment using Automatic Dependent Surveillance–Broadcast and European Reanalysis data: limited flight sample analysis," *Aerospace*, vol. 11, no. 2, p. 154, 2024.

[26] S. Baumann, "Using machine learning for data-based assessing of the aircraft fuel economy," in *IEEE Aerospace Conference Proceedings*, 2019, pp. 1–13.

[27] J. Sun, L. Basora, X. Olive, M. Strohmeier, M. Schäfer, I. Martinovic, and V. Lenders, "OpenSky Report 2022: evaluating aviation emissions using crowdsourced open flight data," in *IEEE/AIAA 41st Digital Avionics Systems Conference (DASC)*, 2022, pp. 1–8.

[28] J. Sun, J. Ellerbroek, and J. M. Hoekstra, "Aircraft initial mass estimation using Bayesian inference method," *Transportation Research Part C: Emerging Technologies*, vol. 90, pp. 59–73, 2018.

[29] X. Olive, J. Sun, L. Basora, and E. Spinielli, "Environmental inefficiencies for arrival flights at European airports," *PLOS ONE*, vol. 18, no. 6, pp. 1–20, 2023.

[30] M. Schäfer, M. Strohmeier, V. Lenders, I. Martinovic, and M. Wilhelm, "Bringing up OpenSky: A large-scale ADS-B sensor network for research," in *Proceedings of the 13th IEEE/ACM International Symposium on Information Processing in Sensor Networks (IPSN)*, Apr. 2014, pp. 83–94.

[31] H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, A. Simmons, C. Soci, S. Abdalla, X. Abellan, G. Balsamo, P. Bechtold, G. Biavati, J. Bidlot, M. Bonavita, G. De Chiara, P. Dahlgren, D. Dee, M. Diamantakis, R. Dragani, J. Flemming, R. Forbes, M. Fuentes, A. Geer, L. Haimberger, S. Healy, R. J. Hogan, E. Hólm, M. Janisková, S. Keeley, P. Laloyaux, P. Lopez, C. Lupu, G. Radnoti, P. de Rosnay, I. Rozum, F. Vamborg, S. Villaume, and J.-N. Thépaut, "The ERA5 global reanalysis," *Quarterly Journal of the Royal Meteorological Society*, vol. 146, no. 730, pp. 1999–2049, 2020.

[32] J. Sun and E. Roosenbrand, "Fast contrail estimation with opensky data," *Journal of Open Aviation Science*, vol. 1, no. 2, Nov. 2023.

[33] X. Olive, "traffic, a toolbox for processing and analysing air traffic data," *Journal of Open Source Software*, vol. 4, p. 1518, 2019.

[34] J. Sun, H. Vû, J. Ellerbroek, and J. M. Hoekstra, "pymodes: Decoding mode-s surveillance data for open air transportation research," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 7, pp. 2777–2786, 2020.

[35] A. Tassanbi, J. Sun, and J. Hoekstra, "Open loop aircraft take-off mass estimation: An optimal trajectory approach," in *Integrated Communications, Navigation and Surveillance Conference (ICNS)*, 2025, pp. 1–9.

[36] J. Sun, J. Ellerbroek, and J. Hoekstra, "Large-scale flight phase identification from ads-b data using machine learning methods," in *7th International Conference on Research in Air Transportation*, 06 2016.

II

Research Proposal

Research Proposal

Lilien Madi
Faculty of Aerospace Engineering, Delft University of Technology, the Netherlands

I. INTRODUCTION

Approximately 4% of all EU greenhouse gas emissions in 2022 came from aviation [1]. Although 4% may seem modest for a relatively energy-intensive sector, the full impact of aviation on the global atmosphere remains uncertain due to the complexity of atmospheric processes involved [2]. This uncertainty highlights the pressing need to reduce aviation-related emissions and enhance sustainability in the sector.

While European legislative bodies have set objectives for climate neutrality, the demand for air travel is expected to continue growing [3]. A multidisciplinary approach is essential to achieve sustainability in the industry. Solutions to minimize aviation's carbon footprint have been researched and, in some cases, developed—such as sustainable aviation fuel (SAF), efficient engines, and alternative propulsion methods like hydrogen and electric engines. Nonetheless, improvements in air traffic management (ATM) performance and operations can be implemented more quickly and deliver faster results, as they mostly include improving existing systems and procedures.

Operational inefficiencies begin months before flights are executed. Airlines may prioritize avoiding higher route charges over fuel efficiency when planning trajectories. On the day of flight execution, fuel tankering and congestion-induced long taxi times increase fuel consumption. En-route, factors like disruptive weather, ATC capacity constraints, and restricted airspace lead to excessive fuel burn [4]. While some dynamic factors (such as weather) are unavoidable, systemic issues, such as drastically differing route charges across European airspace, or route design can be addressed and improved in the future.

Advancements in ATM operations—such as the implementation of Free Route Airspace (FRA) and improved civil/military cooperation—have been proposed to improve flight efficiency. However, safety remains the top priority in aviation, and the wide-scale implementation of such concepts will be challenging if capacity remains limited in highly congested airspaces.

To improve European airspace harmonization and ATM performance, Eurocontrol launched the Performance Review Commission (PRC) which led to the development of Single European Sky (SES) and its technological arm, SESAR. The European ATM Master Plan now guides SESAR's modernization efforts, with updates to meet evolving needs. The PRC's approach also influenced ICAO's Global Air Navigation Plan (GANP) for worldwide ATM standards. Quantifying flight inefficiencies has been a recent topic of interest in these programs. Along with interested airlines, the PRC is investigating a number of approaches to improve and validate the measurement of ATM-related inefficiencies and link them to their root causes [5]. For the next reference period (RP4) in SES, ongoing research is taking into account new performance metrics that have a stronger correlation with actual CO2 emissions [6]. Additionally, proposals to incentivize top

environmental performers or neutralize route charges are under consideration, which would rely on accurate quantification to test potential outcomes through what-if scenarios.

Despite legislative and scientific efforts at optimizing ATM performance, ANSPs still utilize proxy indicators based on extra distance and time flown, such as Horizontal Flight Efficiency (HFE). However, additional distance and time flown may not necessarily correlate with increased fuel consumption particularly if the flight operates under more favorable conditions for fuel efficiency, such as optimal wind conditions, speed, and altitude, as demonstrated in [7]. In [8], results presented a negative correlation between horizontal efficiency and total fuel burn. Nonetheless, fuel-based performance indicators are not enforced as the complexity of fuel-based quantification remains a limiting factor in their implementation. One of the sources of complexity is the need for an aircraft performance model to quantify fuel flow to a certain degree of accuracy; Models such as EURCONTROL's Base of aircraft data (BADA) are sophisticated, but require a license to use. Another source of difficulty is the generation of a reference optimal trajectory. To address these challenges, this study utilizes OpenAP, an open-source aircraft performance model capable of calculating fuel flow and generating optimal trajectories.

The objective of this research is to quantify environmental inefficiencies in flights originating from or terminating at Amsterdam Airport Schiphol (EHAM) using OpenAP in conjunction with fuel-based performance indicators. By using an open-source framework, this work aims to establish a reproducible methodology for flight inefficiency quantification, with the hope of providing relevant stakeholders with a tool for evaluating and enhancing ATM performance.

II. RESEARCH QUESTION

This study is guided by the following main research question and corresponding sub-questions:

How can environmental flight inefficiencies of commercial flights be modeled and quantified?

1) How do executed and planned trajectories compare to optimal trajectories in terms of flight inefficiency?
2) What insights can be drawn from the comparison of inefficiency across different city pairs and airlines?
3) How do horizontal and vertical flight inefficiencies compare in their contribution to total inefficiency in Dutch airspace?
4) How do different flight phases influence flight inefficiency?
5) How do input parameters such as initial mass and true airspeed affect the quantification of flight inefficiencies?
6) What are the limitations in the proposed methodology for quantification of flight inefficiencies?

III. METHODOLOGY

A. Metrics & Data

Given that fuel consumption offers a more accurate representation of a flight's environmental inefficiency in comparison to additional distance or time flown, this study will focus on five of the nine fuel-based performance indicators (PIs) from [9]. These were developed with Trajectory-Based Operations (TBO) in mind, and shown in Figure 1. Four sets of trajectories are required to analyze these indicators: Executed (historical) trajectories, planned trajectories (flight plans), and two reference optimal trajectories.

Executed trajectories will be accessed from the OpenSky network [10]. State vector data from ADS-B messages will be combined with true airspeed (TAS) decoded from Mode-S enhanced surveillance (EHS) messages [11]. Compared to TAS derived from ground speed and wind data, TAS obtained directly from flight data provides better accuracy. Flight plans will be obtained from the Demand Data Repository (DDR2) provided by EUROCONTROL. Both datasets will be selected based on availability and relevance to LVNL's period of interest.

Two reference optimal trajectories will be generated using OpenAP's trajectory optimizer OpenAP.TOP. The first trajectory is unconstrained, representing the theoretically optimal path. The second trajectory, referred to as the optimal trajectory fixing the Executed Reference Business Trajectory (ERBT) route in [9], follows the same lateral path as the executed route but optimizes the vertical profile.

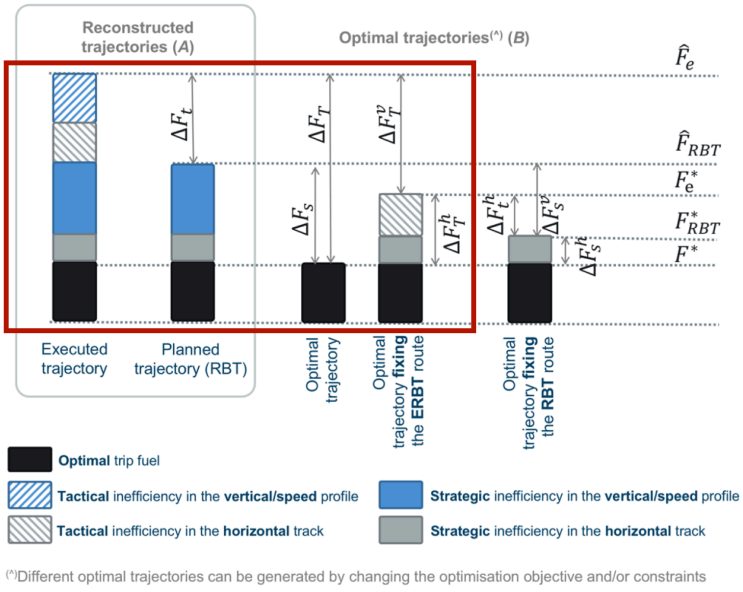

Fig. 1: Sources of fuel inefficiency obtained from [9].

$$\Delta F_T = \hat{F}_e - F^* \tag{1}$$

$$\Delta F_s = \hat{F}_{\mathrm{RBT}} - F^* \tag{2}$$

$$\Delta F_t = \hat{F}_e - \hat{F}_{\mathrm{RBT}} \tag{3}$$

$$\Delta F_T^h = F_e^* - F^* \tag{4}$$

$$\Delta F_T^v = \Delta F_T - \Delta F_T^h \tag{5}$$

Total fuel inefficiency (ΔF_T), is the discrepancy in the fuel consumption of the executed flight ($\hat{F}_e$) and the corresponding optimal flight (F^*). The total strategic inefficiency (ΔF_s) is the difference in the fuel of the planned- or Reference Business trajectory- ($\hat{F}_{\mathrm{RBT}}$) and F^*. The total tactical inefficiency (ΔF_t) is the difference in $\hat{F}_e$ and $\hat{F}_{\mathrm{RBT}}$. The total horizontal inefficiency (ΔF_T^h) is found by calculating the fuel of the optimal trajectory fixing the ERBT route (optimizing the vertical profile of the executed route) (F_e^*), and calculating the difference with the unconstrained optimal trajectory F^*. The total vertical (ΔF_T^v) inefficiency is the difference between the total inefficiency ΔF_T and the total horizontal inefficiency ΔF_T^h. To analyze the relative inefficiency expressed as a percentage, Equations (1), (2), (3), and (4) will be normalized, following the adaptation in [12] of the standard formula for efficiency. In Equation (6), I denotes inefficiency, F represents the fuel consumption of the flight under study, and $Fref$ denotes the fuel consumption of the reference flight. A negative value would indicate improved flight efficiency relative to the reference.

$$I(\%) = \left(\frac{F - F_{\mathrm{ref}}}{F_{\mathrm{ref}}} \right) \times 100 \tag{6}$$

B. OpenAP

Accurately quantifying fuel burn necessitates the use of aircraft performance models. OpenAP, an open-source model developed by Junzi Sun in [13], is well-suited for this analysis due to its multi-functionality. OpenAP is based on non-linear optimal control, and supports the characteristics of 36 aircraft types and over 400 engines. Kinematic performance, dynamic performance, aircraft and engine attributes, and utility libraries are the four main components of OpenAP [13]. In this study, OpenAP will be employed to reconstruct executed and planned flight trajectories (FlightGenerator), generate optimal trajectories (OpenAP.TOP), identify flight phases (FlightPhase), and calculate fuel consumption (FuelFlow).

A critical input for accurate fuel consumption calculations is the initial mass of an aircraft, which is proprietary and not readily available. To overcome this challenge, a machine learning model developed by Aidana Tassanbi at TU Delft will be used to estimate the initial mass based on historical flight data. Although developments in ATM such as ATS B2 (ADS-C), eFPL, and FF-ice may eventually provide direct access to aircraft mass information, prediction models offer a reliable interim solution for improving fuel consumption accuracy. Lastly, wind data from the ERA5 dataset will be used to generate optimal trajectories and derive True Air Speeds (TAS) where needed.

C. Approach

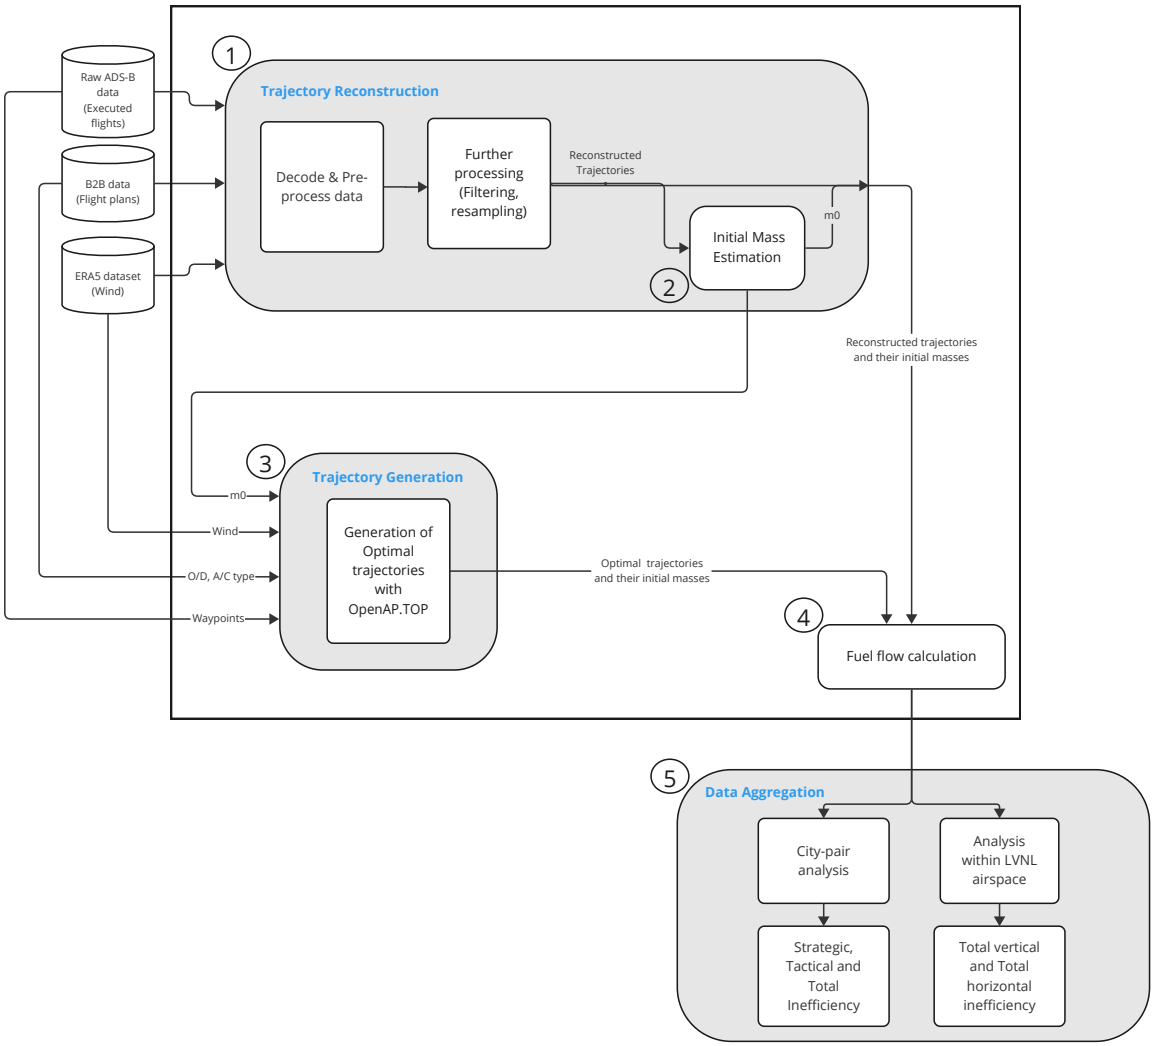

Fig. 2: Framework for inefficiency quantification.

The approach follows the framework outlined in Figure 3. The process begins with data collection and preprocessing. ADS-B data is gathered and EHS messages are decoded using the pyOpenSky, pyModeS and traffic libraries. DDR2 data is converted into time-series data. The data is then processed once again, and irregularities are corrected or removed where necessary, and flights are resampled using traffic. Flights that deviate from their planned destination or operate outside EUROCONTROL member states are excluded. The scope is further narrowed to city pairs that start or end in EHAM, focusing only on commercial flights operated by aircraft supported by OpenAP. Data is enriched to meet the requirements of OpenAP utilities. For example, flight plan data may be supplemented with speed and vertical rate information, or true airspeed (TAS) may be calculated for executed flights in which aircraft do not transmit EHS messages.

After processing and filtering, the next step involves estimating each flight's initial mass. The model developed by Tassanbi is used for this estimation, and the results are necessary inputs for the trajectory optimization step. Trajectory optimization involves generating two reference trajectories for each flight. The first is a theoretically optimal trajectory without constraints, while the second is a trajectory fixing the ERBT route. Fuel consumption is then calculated for each flight. This is done iteratively for every timestep.

Once fuel consumption calculations for every flight is complete, the results are aggregated. Data is grouped by city pair, airline, and flight phase to identify total, strategic and tactical inefficiencies using Equations (1), (2) and (3) respectively. These PIs were chosen as they provide an appropriate level of granularity for a high-level analysis. An additional case study is then conducted to analyze horizontal and vertical inefficiencies within the airspace controlled by LVNL, using Equations (4) and (5) respectively, as these indicators are of particular interest to ANSPs in assessing their performance. This analysis focuses on flights operating in the Amsterdam FIR up to FL245. The segment of the flights operating in this airspace are closely examined, and inefficiencies are calculated using the second reference trajectory as a baseline.

To validate the framework, an initial pilot analysis is performed on flights between Amsterdam and Paris. Amsterdam and Paris are two major transit hubs in Europe, closely linked by the KLM-Air France network. Their strong connectivity leads to high traffic volume and a larger sample size, making this route a suitable choice for the start of the analysis. A subset of data is processed for this route to reconstruct trajectories, estimate initial masses, generate optimal trajectories, calculate fuel consumption, and measure inefficiencies. This analysis is first conducted for a single day and then expanded to cover the selected time frame. Once the pilot is completed successfully, the process is extended to the rest of the city pairs.

IV. EXPECTED RESULTS & LIMITATIONS

The anticipated outcomes of this study include insights into flight inefficiency, while acknowledging inherent limitations. Firstly, inclusion of wind and more representative initial mass will have a non-negligible influence on the results, with the exact extent of this impact being an area of particular interest. With regards to the city pair analysis, city pairs passing through similar airspaces could have comparable inefficiency results as they are subject to the same restrictions, weather patterns, and ANSP capacity constraints. Airline operational policies are also expected to contribute to inefficiency differences for the same city-pairs as carriers have differing strategies for fuel tankering, speed management, and route selection. Airport congestion may be another influencing factor, particularly for city pairs involving high-traffic hubs, where increased holding times and suboptimal arrival procedures will contribute to inefficiency. Lastly, the contribution of flight phases to the overall inefficiency will be observable through this analysis.

This study is subject to several limitations. Firstly, the analysis is limited to aircraft types supported by OpenAP, with approximately one-third of these types supported for fuel flow calculations. The approach to dealing with this limitation is still unclear at this phase, but must be addressed during the research. The OpenAP model has inherent limitations, which may affect the accuracy of the results. It is also important to note that these results are indicative, as the full benefit pool cannot be explicitly identified due to the interplay of sources of inefficiency.

The LVNL case study is particularly limited by the exclusion of intercontinental city pairs and en-route aircraft. This omission excludes a significant percentage of flights in LVNL airspace from the analysis, which could influence results. This limitation primarily stems from resource constraints, particularly time, associated with conducting this thesis. Nevertheless, the results provided will be indicative, and the developed methodology can be extended in this study if time permits, or in future studies, to include all flights traversing Dutch airspace.

Lastly, another limitation in this research is the high computational load of quantifying flight inefficiencies. Processing ADS-B data, generating optimal trajectories, and fuel flow calculations require high processing time and resource allocation, and this is generally one of the limitations in the implementation of fuel-based metrics for inefficiency analysis.

Appendix

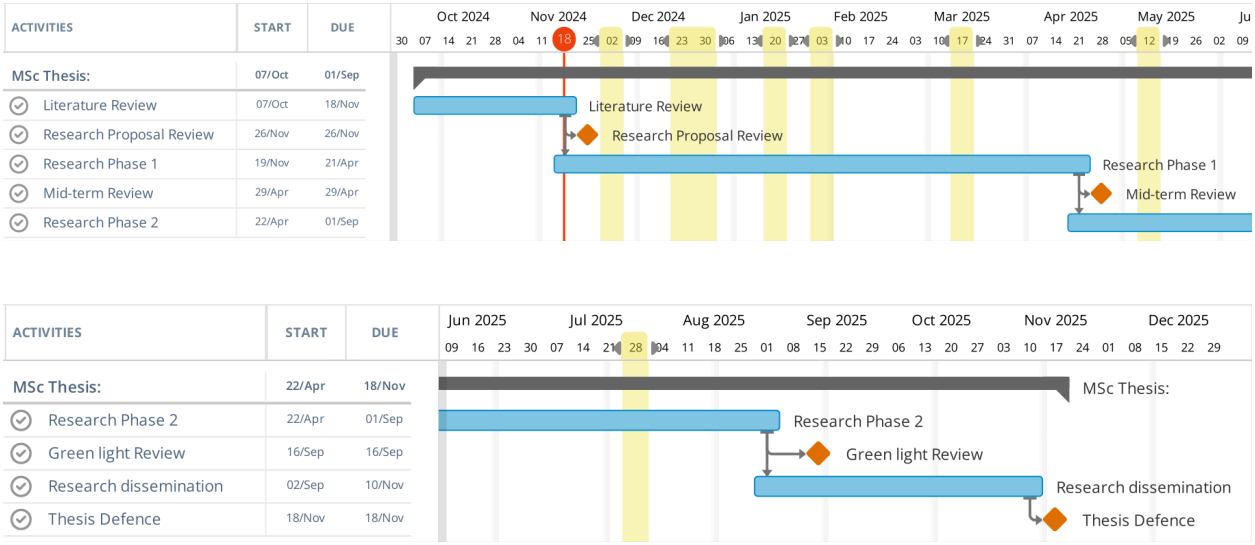

Fig. 3: Gantt Chart including thesis milestones. Yellow bars depict holidays.

Task	Start Date	Due Date
Literature Review	07/Oct	18/Nov
Research Proposal Review	26/Nov	26/Nov
Research Phase 1	19/Nov	21/Apr
Mid-term Review	29/Apr	29/Apr
Research Phase 2	22/Apr	01/Sep
Green Light Review	16/Sep	16/Sep
Research Dissemination	02/Sep	10/Nov
Thesis Defence	18/Nov	18/Nov

TABLE I: Thesis Phases and Deadlines (for additional clarity).

References

[1] European Commission, "Reducing emissions from aviation," [Online]. Available: https://climate.ec.europa.eu/eu-action/transport/reducing-emissions-aviation_en, 2024, accessed: Nov. 14, 2024.

[2] D. S. Lee, D. W. Fahey, A. Skowron, M. R. Allen, U. Burkhardt, Q. Chen, S. J. Doherty, S. Freeman, P. M. Forster, J. Fuglestvedt, A. Gettelman, R. R. De León, L. L. Lim, M. T. Lund, R. J. Millar, B. Owen, J. E. Penner, G. Pitari, M. J. Prather, R. Sausen, and L. J. Wilcox, "The contribution of global aviation to anthropogenic climate forcing for 2000 to 2018," Atmospheric Environment, vol. 244, p. 117834, 2021.

[3] EUROCONTROL, "European aviation EUROCONTROL forecast update 2024-2030," [Online]. Available: https://www.eurocontrol.int/sites/default/files/2024-02/eurocontrol-seven-year-forecast-2024-2030-february-2024.pdf, Feb. 2024, accessed: Nov. 14, 2024.

[4] T. G. Reynolds, "Air traffic management performance assessment using flight inefficiency metrics," Transport Policy, vol. 34, pp. 63–74, 2014.

[5] EUROCONTROL Performance Review Unit, "Performance insight 9: Gate-to-gate co2 emissions in europe - a holistic approach," [Online]. Available: https://www.eurocontrol.int/sites/default/files/2024-06/eurocontrol-prc-performance-insight-9.pdf, EUROCONTROL, Tech. Rep., 2024, accessed: Nov. 15, 2024.

[6] "SES environmental performance and targets," [Online]. Available: https://www.easa.europa.eu/eco/eaer/topics/air-traffic-management-and-operations/ses-environmental-performance-targets, 2024, accessed: Nov. 14, 2024.

[7] M. S. Ryerson, M. Hansen, and J. Bonn, "Time to burn: Flight delay, terminal efficiency, and fuel consumption in the national airspace system," Transportation Research Part A: Policy and Practice, vol. 69, pp. 286–298, 2014.

[8] E. Calvo, J. M. Cordero, L. D'Alto, J. López-Leonés, and M. Vilaplana, "A new method to validate the route extension metric against fuel efficiency," 2024.

[9] X. Prats, R. Dalmau-Codina, and C. Barrado, "Identifying the sources of flight inefficiency from historical aircraft trajectories: A set of distance-and fuel-based performance indicators for post-operational analysis," Jun. 2019.

[10] M. Schäfer, M. Strohmeier, V. Lenders, I. Martinovic, and M. Wilhelm, "Bringing up OpenSky: A large-scale ADS-B sensor network for research," in ACM/IEEE International Conference on Information Processing in Sensor Networks, Apr. 2014.

[11] J. Sun, The 1090 Megahertz Riddle: A Guide to Decoding Mode S and ADS-B Signals, 2nd ed. Delft: TU Delft OPEN Publishing, 5 2021.

[12] J. Bronsvoort, P. Zissermann, S. Barry, and G. McDonald, "A framework for assessing and managing the impact of ANSP actions on flight efficiency," Air Traffic Control Quarterly, vol. 23, no. 1, pp. 29–53, 2015.

[13] J. Sun, J. M. Hoekstra, and J. Ellerbroek, "OpenAP: An open-source aircraft performance model for air transportation studies and simulations," Aerospace, vol. 7, no. 8, p. 104, 2020.