Trans. Japan Soc. Aero. Space Sci.
Vol. 67, No. 3, pp. 109–118, 2024
DOI: 10.2322/tjsass.67.109

Considering Expected TMA Holding into In-flight Trajectory Optimization*

Luis DELGADO,[1)] David DE LA TORRE,[2)] Jovana KULJANIN,[2)] and Xavier PRATS[2)†]

[1)]*Department of Architecture and Cities, University of Westminster (UoW), London NW15LS, United Kingdom*
[2)]*Department of Physics-Aerospace Division, Technical University of Catalonia (UPC)-BarcelonaTech, Castelldefels 08860, Spain*

Aircraft crew are aware of the delay they have experienced at departure. However, uncertainties ahead, and in particular holdings at arrival, can have an impact on the final performance of their operations. When optimizing a trajectory, the expected cost at the arrival gate should be considered. Consequently, taking into account potential congestion and extra delay at the arrival airspace is paramount to avoid making sub-optimal decisions during the early stages of a flight. This paper presents a framework to optimize trajectories in the execution phase of the flight considering expected delays at arrival. A flight from Athens (LGAV) to London Heathrow (EGLL) is used as an illustrative example, systematically exploring a range of departure delays and expected holdings at arrival.

Key Words: Trajectory Optimization, Flight Execution Phase, Uncertainty, Airline Costs

1. Introduction

The continuous growth of traffic demand is leading to the situation in which terminal maneuvering areas (TMA) could become the bottleneck of the entire air transportation system. In some concepts of operations, aircraft arriving at busy TMAs might be required to hold until sequenced to the approach – an event that will increase the flying time and fuel usage, and result in an arrival punctuality. Other TMAs try to avoid holding patterns, but at the expense of more intensive (and unpredictable for the aircraft crew) radar vectoring. This includes implementing sequencing and merging concepts such as tromboning or point merge;[1,2)] and/or tactically limiting the speeds of the aircraft upon their entrance into the TMA (i.e., linear holding), as performed with an Extended Arrival Manager (E-AMAN).[3)]

The actual holding time at the TMA could substantially vary from the expected (average) value due to uncertainties inherent in this process. For instance, an aircraft flying under good weather conditions and average traffic in the TMA will most likely have an average holding time. On the other hand, a flight approaching an airport with reduced capacity (e.g., due to bad weather) is more likely to present higher uncertainty.

Cutting-edge pilot decision support tools, such as Pacelab Flight Profile Optimizer (FPO) developed by PACE[4)] or ClearPath developed by AVTECH,[5)] are gradually being deployed in commercial aviation. These tools aim at computing tactical trajectory updates in order to improve the execution of the flight when uncertainty has already been materialized. Examples include a new weather forecast has been uplinked, a significant route shortcut has been granted by air traffic control, or an en-route flight level becomes available. Yet, these tools mainly optimize for direct operating costs (i.e., trip fuel and trip time) and do not consider downstream uncertainty on the operations (e.g., on the remaining flight, such as arrival delays). Generally, already accrued delay, such as delay experienced at departure, is assumed to be the same as that expected at arrival, and flight vertical and speed profiles are adjusted solely with these considerations.

Crew can usually introduce an estimation for expected delay at arrival, but this bases the optimization on the crew experience for a given route and focuses on delay rather than expected cost, which materializes as a function of the arrival time at the gate.

This might lead to sub-optimal decisions such as to recover some delay at a high fuel expense, with no significant benefit; or conversely, passenger missed connections which could have been averted by speeding up and trading some fuel due to the long delay is expected at the destination TMA, even if the initial departing delay is low. Therefore, as it is explored in this article, considering only the accrued (departing) delay is not sufficient, and expected disruptions, particularly holding at arrival, can have an important role in the outcome of trajectory optimization.

Note, that if an E-AMAN system is in place, the delay required per flight will be coordinated by this centralized system. However, flights will not have an arrival slot (and delay) assigned to them until they are closer to the airport. Therefore, airlines still have the possibility to try to recover delay if congestion is expected prior to having an arrival slot assigned. The prototype presented in this paper could also be used to assess the trade-off between delay and costs in a future concept of operations when arrival slots can be *negotiated*, as envisioned in SESAR for instance.

*Presented at 2022 International Workshop on ATM/CNS, 25–27 October 2022, Tokyo, Japan.
Received 30 November 2022; final revision received 22 March 2023; accepted for publication 22 June 2023.
†Corresponding author, xavier.prats@upc.edu

Pilot3, an innovative action funded under the Clean Sky 2 programme, aims at developing a software prototype for supporting crew decisions for civil aircraft in the execution phase of flight. By triggering the tool, the software provides at least two trajectory options, along with information on different indicators to aid the crew to select the most suitable one. This selection considers the multi-criteria business objectives of the airline, including the impact those decisions have on the airline's network. In particular, Pilot3 provides an optimization framework for trajectories considering the expected total cost and operational uncertainties.

In this paper, delays in the TMA (i.e., holdings in particular) are integrated in the cost function to optimize the trajectory. Our objective is to analyze how these holdings can affect the optimal trajectory, and hence, the decision performed by the crew. It is out of the scope of this paper to explain how this holding, and the details on the components of the cost function, are estimated.

A systematic range of departure delays and holding times are analyzed for a flight with destination to London Heathrow (EGLL), one of the most congested airports in Europe and one of the most challenging TMAs, since it serves several airports in the London area. London TMA typically operates with four holding stacks, entailing that aircraft are progressively being taken out of the bottom of the stack and vectored to the final approach, allowing aircraft holding at higher levels to descend to lower ones.[6] An analysis of traffic data from Sep. 2018 shows that around half of flights arriving to EGLL have some holding, which in extreme cases can reach up to 35 min.[7] The average holding times, however, are much shorter: about 8.5 min at the beginning of 2014 and falling to 7.5 min in 2016.[8]

2. Methodology

In a general case, Pilot3 could be triggered at any point of the flight (i.e., from the departure procedure to the initial descent). Pilot3 optimizes the aircraft trajectory from the current aircraft state (i.e., the moment Pilot3 is triggered) down to FL100 at the proximity of the destination airport (see Fig. 1). Below this altitude, the actions of the aircraft are significantly limited and highly standardized. Moreover, the aircraft trajectory is likely to be modified several times by tactical ATC intervention, thus forcing the pilot to no longer follow an optimized trajectory plan. In this regard, Pilot3 will compute the remaining trajectory plan assuming standard operations (i.e., a fixed sequence of aircraft intents) in a similar way it is currently done using on-board flight management systems.[9]

Operational uncertainties considered in the optimization framework include: holding time; distance to be flown during the final approach, sequencing and merging phase, which is understood as the distance from FL100 to the runway; and taxi-in time. It is assumed that all these uncertainties are experienced after FL100; even if the holding could be before, it is just a temporal displacement. Pilot3 will consider not only the average expected value of these sources of uncertainty, but their distribution when computing the expected cost of delay, as presented in Section 2.1.

As all uncertainties are limited to the FL100 to gate phase, and the optimization of the trajectory finishes when reaching FL100 in the descent, the optimizer developed can be considered deterministic. This optimizer will minimize the expected total cost computed as expected cost of fuel and expected cost of delay as a function of arrival time at FL100, as presented below.

2.1. Cost function modeling

The total cost that a flight will experience is composed of two components: cost of fuel and cost of delay.

Cost of fuel considers the amount of fuel used by the optimized trajectory from triggering point to FL100, and the expected fuel used for the processes from FL100 to gate (i.e., holding, sequencing and merging, and taxi-in).

The cost of delay depends on the arrival time at the gate, as this will be translated into reactionary delay, passenger satisfaction, compensations and missed connections, crew and maintenance costs, etc. This cost function is highly non-linear and discontinuous. It can be seen as a step-wise function, as increments are produced and linked to events; for example, reaching the threshold for having to compensate passengers due to Regulation 261[11] if they are entitled, or breaching a curfew at the end of the day due to reactionary delay.[12]

Most of the events that generate the cost of delay have some degree of uncertainty associated: the uncertainty of when the cost will actually materialize (e.g., how many passengers will actually miss their connection for a given arrival time at the gate), and the uncertainty of the cost value itself

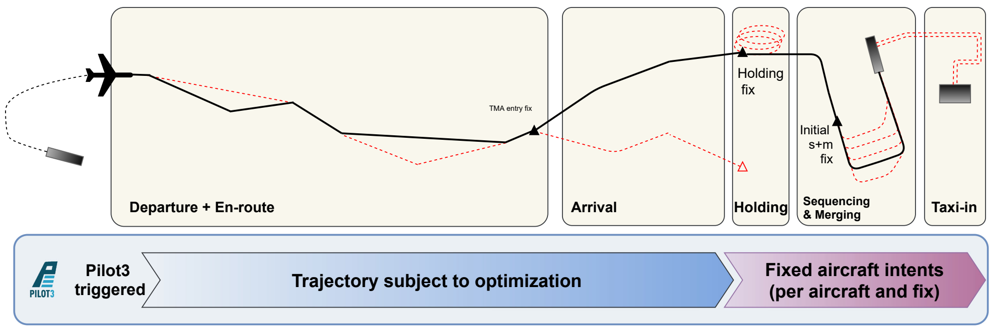

Fig. 1. Trajectory optimization concept for Pilot3.

Table 1. Main characteristics to compute the operational flight plan (OFP).

Flight schedule	Flight dispatch	Other operational information
Airline: British Airways	Cost Index: 10 kg/min	Estimated taxi-out: 10 min
Aircraft type: A320-231	Cost of fuel: 0.5 Eur/kg	OFP trip time: 216 min
Stage: LGAV–EGLL	Payload: 144 passengers* + 1,000 kg Cargo	Buffer at arrival (taxi-in and padding): 9 min
SOBT: 05h15 UTC	LOGAN 2H arrival + ILS approach to runway 09R	Planned holding point: LAM
SIBT: 09h10 UTC	Weather forecast issue/applicability: 2016-07-28 5h UTC	124 connecting passengers

*According to the EU-OPS 1.620[10] flights "within the European area" shall account per adult passenger 97 kg (luggage included).
SOBT/SIBT: Scheduled off/in-block time, UTC: Coordinated universal time, ILS: instrumental landing system.

(e.g., how many passengers will claim compensation even if entitled). With these considerations, the optimization framework estimates the expected cost of delay as a function of arrival time at the gate.[9] Note that the details on how these computations are performed are out of the scope of this article, but they are estimated based on the European Airline Delay Cost Reference Values.[12]

Then, given an arrival time at FL100, the actual time of arrival at the gate, and hence, the associated cost of delay is determined by the convolution of the stochastic processes of potential holding, final approach/sequencing and merging from FL100 to runway and taxi-in times.[13] The distribution of times for these processes are considered instead of just their average time. This is required to compute the expected cost of delay function because the estimated cost of delay at the gate is non-linear, as previously indicated. As explained in Section 4, the consideration of this uncertainty will have a significant impact on the shape of the expected cost function, and hence, on the outcome of the optimization and the crew and flight behavior.

By combining the expected cost of fuel and delay, Pilot3 computes the expected total costs (ETC) function that is used for the optimization. This approach substantially differs from the most widely used methods in flight planning, in which the direct operating costs (DOC) are computed as a weighted sum of cost of fuel, (nominal) cost of trip time, and route charges (i.e., air navigation fees). The so-called cost index (CI) is the weighting parameter that relates the cost of time versus the cost of fuel in this kind of approach.[14]

2.2. Trajectory optimization: CI as a proxy

The trajectory optimizer software Dynamo,[15] which uses a point-mass representation of the aircraft and high-fidelity aircraft performance and meteorological data, is able to optimize the vertical profile of a flight for a given CI. This is done by selecting the set of flight altitudes via grid search and speeds via pre-optimized tables as function of CI that minimize a DOC-type cost function.

Therefore, for a given CI, a trajectory can be generated with Dynamo. The expected cost of delay function when reaching FL100, as explained in Section 2.1, is then used to compute the expected total cost of the trajectory: cost of fuel obtained using Dynamo from triggering point to FL100; average cost of fuel for final phases of the flight (i.e., holding, sequencing and merging, and taxi-in); and expected cost of delay.[9]

With this framework, the optimization performed consists of obtaining the best CI such that the ETC is minimized. This is done using a binary-search algorithm.[16] Note that even if CI is used to generate the trajectories, the final optimization minimizes the total cost function (i.e., ETC) and not DOC.

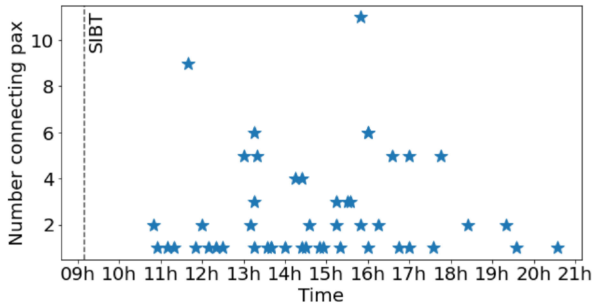

Fig. 2. Passenger groups connecting at EGLL into follow-up flights on the LGAV–EGLL flight.

3. Scenario and Case Studies

This section specifies the flight and case studies used to illustrate the methodology proposed in this paper, and explains the impact of having estimations of holdings at arrival with different degrees of certainty. For this purpose, a flight from Athens (LGAV) to London Heathrow (EGLL) was selected taking into account the flight schedule, dispatch, and operational considerations summarized in Table 1.

Individual passenger itineraries, with their connections, are modeled based on historical data from IATA's PaxIS and Global distribution systems datasets, as in previous research projects.[17] Figure 2 shows the different passenger groups with connections at EGLL for the flight chosen for this study. The number of passengers of each group is indicated together with the time where their connecting flight is scheduled to depart. Note how some passengers have a long waiting time at EGLL before their subsequent flight (i.e., their connection will not be missed even if some arrival delay is experienced). However, the first passenger group with a connection is for a flight scheduled at 10h50 (i.e., the SIBT of the LGAV–EGLL flight is 9h10). Considering a standard minimum connecting time of 84 min at EGLL,[17] some passenger groups will start missing connections with an arrival delay greater than just 16 min. As will be explained in Section 4, the fact that passengers miss connections does not necessarily result in a significant increase in the expected cost of delay if, for example, they can be re-accommodated and arrive at their final destination before being entitled to compensation due to Regulation 261, if entitled to this.

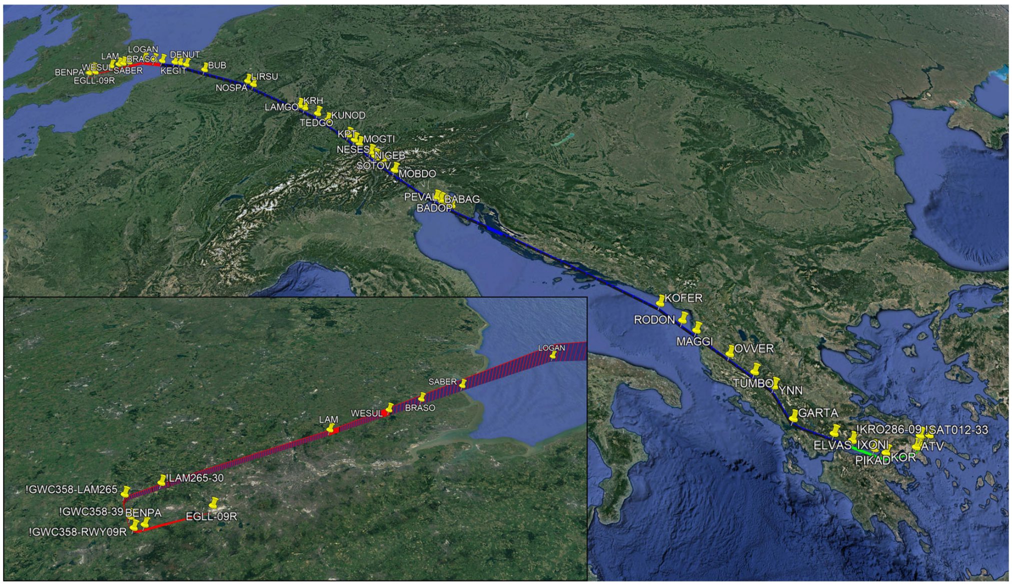

Fig. 3. OFP route: Horizontal trajectory profile. Detail of the descent trajectory in the lower-left side.

3.1. Operational flight plan (OFP)

Considering the scenario depicted in Table 1, an OFP is generated as follows: The route (i.e., sequence of waypoints) is obtained from EUROCONTROL's Demand Data Repository 2 (DDR2);[18] then, the vertical (and speed) trajectory profile is optimized with Dynamo[15] using aircraft performance data from EUROCONTROL's BADA v4.2[19] and weather forecast information from the European Centre for Medium-Range Weather Forecasts (ECMWF) ERA5.[‡] The optimization criterion for this optimization is the standard DOC function, assuming that the cruise Mach is kept constant for a given cruise flight level.

Figure 3 presents the OFP trajectory with the climb, cruise, and descent phases represented, respectively, by green, blue and red segments. Figure 4 shows the resulting vertical and speed profiles of the OFP trajectory with the along-track and cross-wind components at different altitudes (colored backgrounds). In these plots, the pressure altitude (hp) for the whole trajectory is depicted together with the Mach number (M), calibrated airspeed (CAS), true airspeed (TAS), and ground speed (GS). It is worth mentioning that the (apparently) sudden changes in ground speed of these figures, such as that observed at around 1,300 NM from the destination airport, are due to track changes in the lateral route, which change the relative wind direction along and cross-track, and therefore the resulting ground speed. These plots also depict the maximum operational speeds for that aircraft type: maximum Mach in operation (MMO) and maximum CAS in operation (VMO). It is worth noting that pressure and temperature data given in the ERA5 weather forecast are also considered in the optimization process. The impact of temperature, for instance, on the optimization is analyzed in more detail in another publication.[16]

As shown in Fig. 4, the OFP for this scenario consists of an initial cruise at FL360 followed by a step-climb to FL380 approximately 850 NM from the destination airport. The optimal cruise speed resulting for this OFP is M.770. The first half of the cruise is mainly affected by a relatively strong crosswind component of approximately 60 kt, while a relatively mild headwind and crosswind components dominate the remaining cruise.

Arrival procedures at EGLL are obtained from the UK AIP,[20] AIRAC 2111 (i.e., issued on Nov. 4 2021). The arrival procedure LOGAN 2H, which ends at the Lambourne fix (LAM) is used. This is the fix where the holding pattern is located. For flight and fuel planning purposes (i.e., to compute the OFP), the approach to runway 09R is chosen since it is the longest possible.

3.2. Definition of case studies

This paper explores systematically different departure delays and expected holdings at the London TMA. All case studies consider that Pilot3 is triggered when reaching the top of climb (TOC). At that moment, it is assumed that the aircraft crew evaluates the status of the flight with respect to time adherence.

A flight might depart late for a combination of factors: leaving the gate with a deviation in the planned schedule, a taxi-out time different than originally planned, route shortcuts or path stretching in the departure phase, etc. These factors eventually materialize into deviations in the time of arrival at the TOC, if compared with the OFP. In order to cover a range of operational departures, delays from −10 to 180 min at the TOC in 5-min intervals are simulated.

Regarding holding at arrival, a range between 0 and 25 min at 5-min intervals is explored; these being the common holding times at EGLL.[§] This delay may not be known by the crew since it is experienced at arrival. The impact of this de-

[‡]https://www.ecmwf.int/

[§]Note that 0 min of holding implies that the flight would fly the arrival procedure as planned without any holding at the holding point.

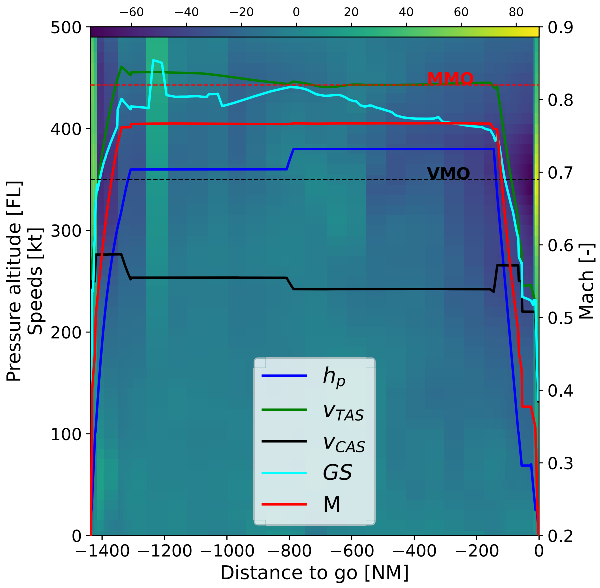

(a) Background color: along-track wind.

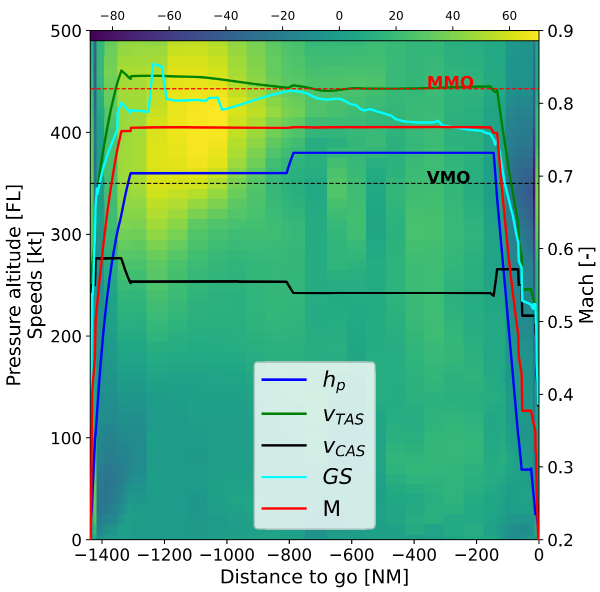

(b) Background color: cross-track wind.

Fig. 4. OFP vertical and speed trajectory profiles.

lay if known by the system is analyzed (i.e., using an estimator able to predict it).

For all case studies, it is considered that all passengers are entitled to Regulation 261 compensation if delay thresholds at their final destination are met.[11] Note that passengers are only entitled to this compensation if the airline is deemed responsible for their delay.

4. Results

This section presents the results for the scenario and case studies presented in the previous section.

The optimization performed using Pilot3 is compared in terms of expected total cost (EUR), fuel (kg), and time deviation (min) with respect to the default alternative of maintaining the operational flight plan (OFP).

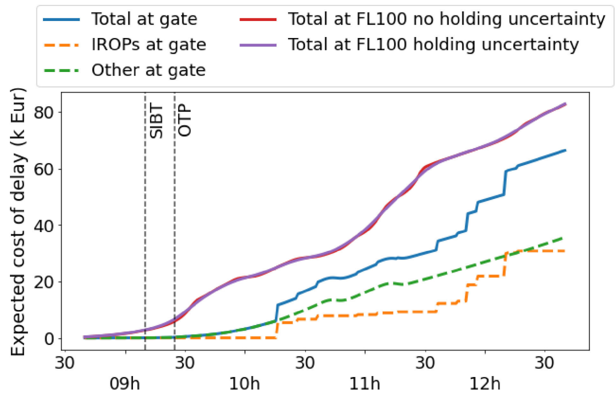

Fig. 5. Expected total cost of delay (and components) as a function of arrival time at gate and at FL100.

4.1. Cost function

Figure 5 shows the breakdown for the expected cost of delay as a function of the arrival time at the gate (i.e., in blue). As observed, the total cost of delay consists of IROPs costs (e.g., passenger compensation costs, assuming that passengers are entitled to compensation due to Regulation 261 if delay thresholds are passed) and other costs (i.e., reactionary delay, crew, and maintenance related costs). These costs are estimated as explained in Section 2.1, and as previously mentioned, based on the European Airline Delay Cost Reference Values.[12]

Some costs dominate other costs, with reactionary costs propagated in subsequent rotations being overall the main component. For this flight, the time allowed for rotations is relatively tight: the aircraft has 40 min for the rotation at EGLL before departing to LIRF. Therefore, if the arrival to EGLL is delayed, the probability of delay being propagated to LIRF, and to subsequent flights, is high.

Recall from Section 3 that this flight has a relatively tight buffer at arrival with only 9 min for taxi-in and padding (i.e., the difference between the estimated landing time and the SIBT), and that some passenger groups will start missing connections with an arrival delay greater than just 16 min. However, as observed in the cost function, Fig. 5, cost associated to these missed connections will not have a significant impact until around 1 h of delay when, as shown in Fig. 2, a group with nine passengers will miss their connection. This is due to the fact that, in some cases, if passengers can be re-accommodated into subsequent flights arriving to their destination before the Regulation 261 threshold, the cost for the airline for these missed connections can be rather low even if the delay is attributed to the airline. However, once compensations, and duty of care (e.g., waiting until next day for a flight) are due, cost increases in sharp steps.

Finally, and as explained in Section 2.1, considering the arrival processes with their uncertainties (i.e., holding, sequencing and merging, and taxi-in), the expected cost of delay can be expressed as a function of the arrival time at FL100. These arrival processes will be translated into an average shift of the cost function to be used by the Pilot3 optimizer. This shift includes the addition of the expected times of these arrival processes. For the particular example of Fig. 5, 46.3 min are obtained, resulting from holding (i.e., 20 min), sequencing and merging (17.7 min), and taxi-in (8.6 min).

In Fig. 5, two costs are presented as a function of arrival time at FL100, both considering an expected holding of 20 min, but in one case no uncertainty is considered (i.e., red) while in the other a normal distribution with a sigma of 6 min is used (i.e., purple).

When considering uncertainty, the time from FL100 to the runway is a probability distribution. Therefore, the expected cost at FL100 is computed considering estimation of the arrival time; hence, considering this probability density function over the cost of expected total cost of delay at gate. Note how uncertainty *smooths* the expected cost function. Yet, the two cost curves overlap as, in this example, the expected time of the holding does not change. For the results presented in the next section, no holding uncertainty will be considered to facilitate the discussion.

4.2. Analysis of the Pilot3 optimized solution

Figure 6 presents the results of the optimization of the trajectories using Pilot3 for the range of departing delays and holding times explained in Section 3.2. Three different results are shown: expected costs savings of the optimized trajectory with respect to maintaining the OFP, Fig. 6(a); expected delay to be recovered by the optimized trajectory, Fig. 6(b); and finally, the variation in fuel consumption, Fig. 6(c).

The first aspect to notice is that in both trajectories under comparison (optimized and OFP), the final arrival processes are modeled in the same manner. This means that the fuel and time required to reach the gate from FL100 will be the same for both cases, as both model the same holding, final arrival, and taxi-in processes.

In this study, it is assumed that the estimation of the holding time is deterministic, and this is translated into a shift in the cost function to be used by the optimizer as previously shown, the result obtained by the optimization depends only on the total delay expected (i.e., departure plus holding), and not on how this delay is shared among them. This can be observed in Fig. 6, where the expected cost savings, delay recovered and extra fuel usage are presented as a function of expected holding time and departure delay. The outcome of the optimization depends on the aggregated expected delay (i.e., departing and holding delay combined).

Nevertheless, the division between both sources of delay are kept independent in this figure, as this facilitates a more precise operational analysis. For example, if departing on time, the crew not usually considers the recovery of any delay. Yet, if the expected holding delay is 15 min, a recovery of around 5 min represents savings close to 100 EUR. In current operations, the crew relies on their expertise of previous operations to decide if the delay should be recovered without having a clear view on the impact of these decisions in terms of the expected costs for the airline.

Figure 6(b) shows how there is a maximum optimal amount of delay recovered for this flight (i.e., 22.5 min), at

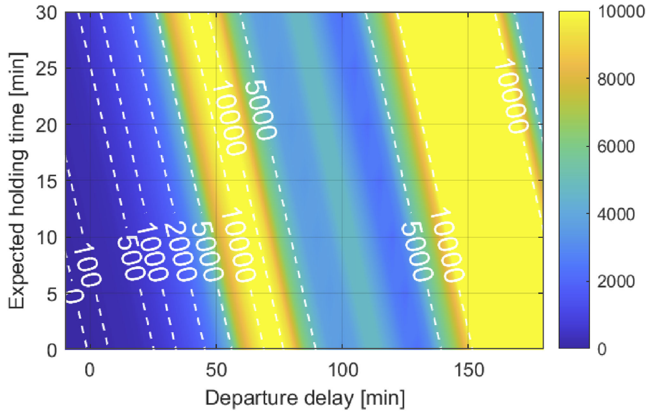

(a) Expected cost ssavings (EUR)

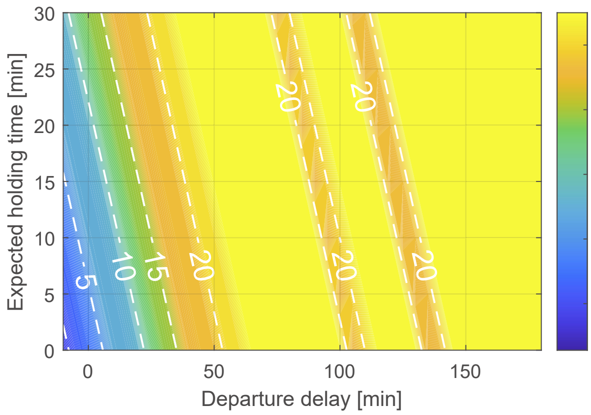

(b) Delay recovered (min)

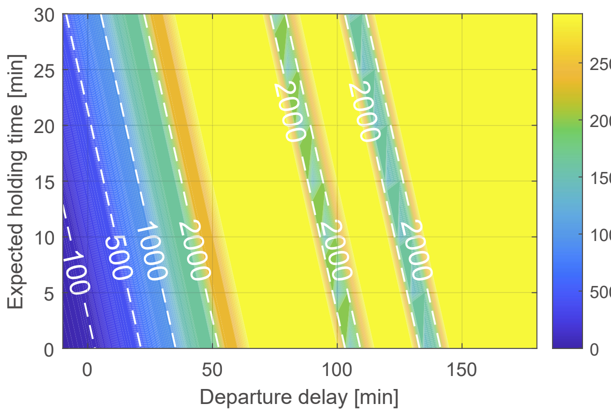

(c) Extra fuel usage (kg)

Fig. 6. Results of Pilot3 optimization for the different departure delays and expected holding times.

the expense of burning 2,936 kg of extra fuel, if compared with the OFP (see Fig. 6(c)). As expected, as the departure delay and expected holding increases, the amount of delay to be recovered also increases.

It is worth noticing, however, contrary to simple rules of thumb, there are regions where recovering less than the maximum possible delay is more suitable, even in the case of high initial delays. For instance, if the expected arrival delay is 115 min considering that departing and holding delay are translated into arrival delay (i.e., no correction is done to the flight due to the late departure and the flight is operated as planned until Pilot3 optimization is triggered), the delay re-

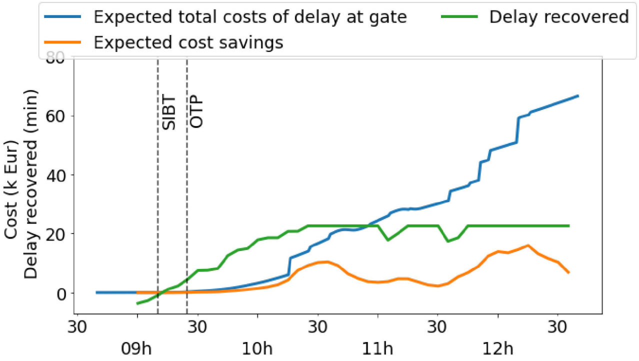

Fig. 7. Expected total cost of delay as a function of the arrival time at the gate if flying the OFP. For each of these arrival times, the delay recovered and expected cost savings of the Pilot3 solution are also given.

covered by the optimized trajectory is 17.8 min, instead of the maximum of 22.5 min. In this case, the extra fuel consumption is reduced from 2,936 kg to 1,446 kg. This behavior is also observed for expected arrival delays of 120, 145 and 150 min.|| For those cases, the amount of recovered delay is 20.0, 17.4 and 18.6 min, respectively; and the corresponding extra fuel consumption, with respect to the OFP, is 2,096, 1,368 and 1,637 kg, respectively. This can be easily observed by the *color bands* appearing in Fig. 6(b) and 6(c) for these amounts of expected delays.

In a similar manner, the savings obtained do not evolve in a monotonous way as a function of the total expected delay at arrival. As observed from the cost function (see Fig. 5), the cost of delay is non-linear and presents regions where higher benefits by recovering delay can be achieved than other possibilities, as seen in Fig. 6(a). If the event that triggers a given cost is non-recoverable (e.g., even with the maximum possible recovery passengers will miss their connection), it might be more economical to recover less delay and save fuel.

These outcomes are therefore a reflection of the shape of the expected cost function and will vary for different flights under different conditions. The framework described in this article allows these regions to emerge as a function of the departing and expected holding delay.

To simplify the analysis of the impact of the cost function on the behavior of the optimizer, Fig. 7 presents the cost function at the gate as a function of the expected arrival time at the gate by maintaining the OFP (i.e., when actions to recover departure delay or expected arrival holdings are not taken). For each of these arrival times, the delay recovered and expected cost savings of the Pilot3 solution are also provided (i.e., with respect to the expected outcome if OFP is maintained).

First, if the expected arrival time at the destination gate is before the SIBT, Pilot3 will generate a trajectory that slows down the flight, as the extra delay generated will be compensated with fuel savings. Then, as the expected arrival time increases, due to the initial delay, the delay recovered tends to increase, but as previously mentioned, this does not increase monotonically (remember, from previous discussion that recovering the maximum amount of delay is not always the optimal decision). Finally, cost savings are closely related to local variations of the cost function as a function of the arrival time (i.e., when sharp increments are observed); for example, due to passengers missing connections.

4.3. Optimized vertical and speed profiles

This section discusses the speed and vertical profiles of the Pilot3 optimized trajectories corresponding to the case studies presented above. Figure 8 shows the expected cost of delay as a function of the arrival time at FL100; and the time that FL100 should be overflown in order to arrive at the gate at the SIBT. Six representative cases have been selected for further analysis in this section. Hence, for each of these cases, Fig. 8 also shows the expected cost of fuel, the expected total cost, the arrival time at FL100 if the crew keeps flying the OFP for the whole flight, and the arrival time at FL100 for the Pilot3 optimized solution. Figure 9, in turn, shows the vertical and speed profiles of the six selected trajectories: the OFP trajectory is depicted with pale lines, and the optimized Pilot3 trajectory in solid darker lines.

When the expected arrival time at the destination gate is before the SIBT, Pilot3 proposes to slow down and slightly increase the trip time, if compared with the OFP. This is observed in Fig. 8 for the case with −10 min of expected arrival delay (i.e., if flying OFP is kept). In Fig. 9(a), it can be observed that Pilot3, when triggered at the TOC, selects a slightly lower Mach number and slightly delays the planned step climb from FL360 to FL380. In doing so, some fuel can be saved without having an impact on the cost of delay (i.e., the Pilot3 solution still arrives before the SIBT), yielding to a better expected total cost when compared to maintaining the OFP (see Fig. 6).

For expected delays at FL100 when flying OFP ranging from 5 to 15 min, the Mach number of the Pilot3 solution progressively increases with this expected delay, while the vertical profile is very similar to the OFP profile. When the expected delay ranges from 20 to 30 min, the planned step climb at FL380 is not done anymore, and Pilot3 proposes to perform the whole cruise at FL360, while increasing the Mach number as well. Figure 9(b) shows the results for a particular case when the expected delay at FL100 (if flying the OFP) is 20 min. In this case, the cruise Mach number has increased to M.794, while the delay recovered using the Pilot3 solution is 7.5 min, as shown in Fig. 8; this resulting from the expense of burning extra fuel volume of 267 kg.

If the expected delay at FL100 when flying the OFP reaches 35 min, then Pilot3 proposes descending to FL320 at about halfway through the cruise, as observed in Fig. 9(c). In general, lower flight levels would represent an increase in true airspeed (TAS) for a given Mach number, since the speed of sound, which depends on air temperature directly, is typically higher at lower altitudes. In this case, the delay

||As stated before, note that these expected arrival delays will be reached by combining the departure delay and expected holding time, which add to these values (e.g., 120 min of expected arrival delay will be achieved with 100 min of departing delay and 20 min of holding time, or 110 min of departing delay and 10 min of holding time, etc.).

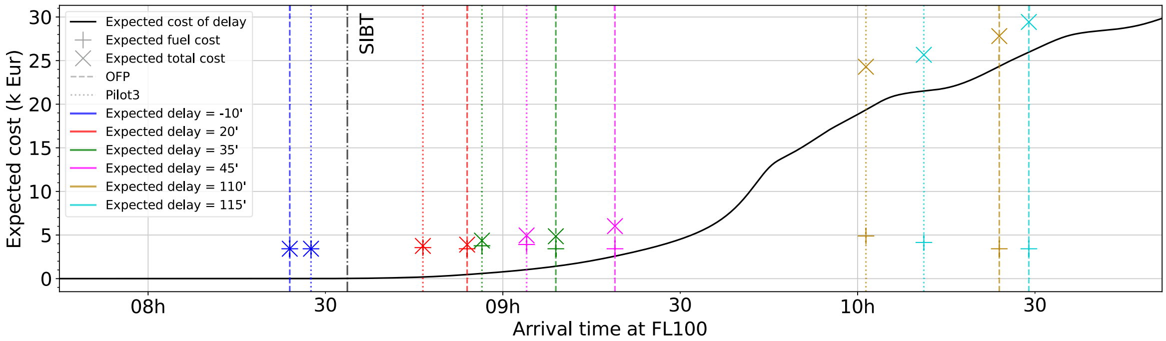

Fig. 8. Expected costs for six representative cases with different expected delays at FL100 when flying with OFP.

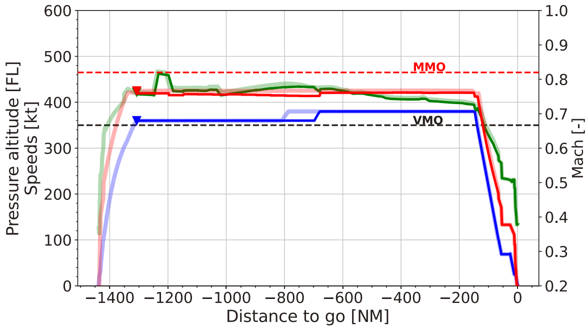

(a) Expected delay = –10 min

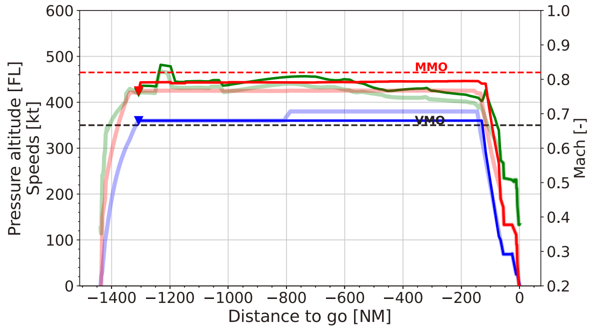

(b) Expected delay = 20 min

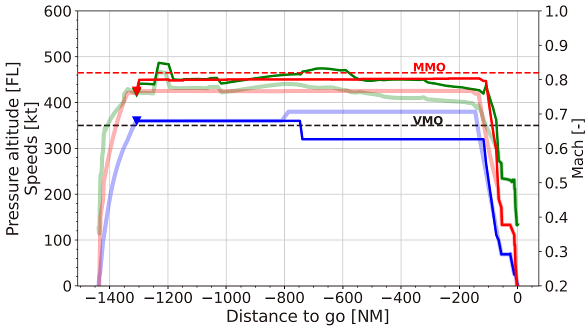

(c) Expected delay = 35 min

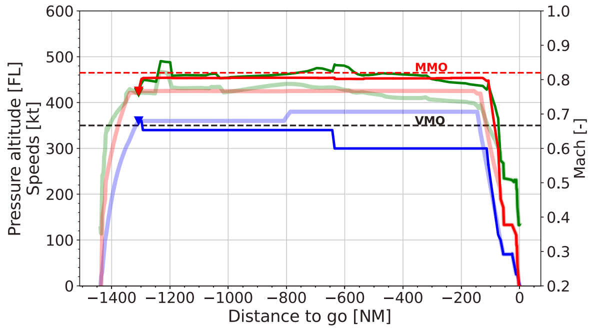

(d) Expected delay = 45 min

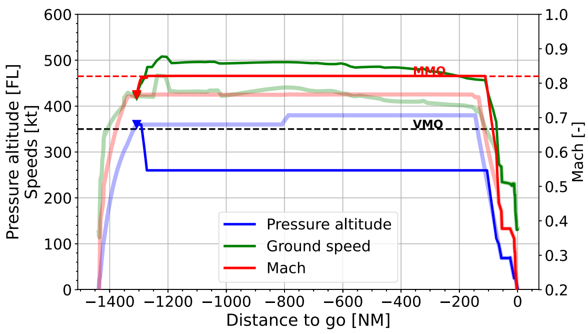

(e) Expected delay = 110 min

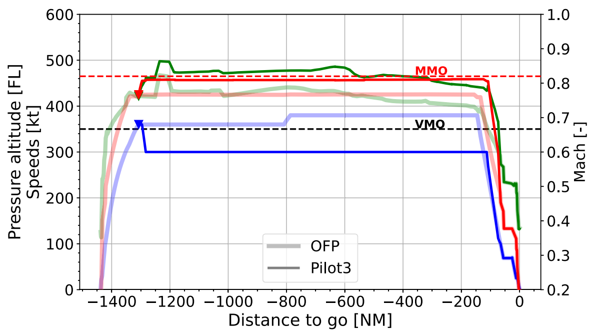

(f) Expected delay = 115 min

Fig. 9. Vertical and speed profiles for six representative cases with different expected delays at FL100 when flying with OFP.

recovered using the Pilot3 solution is 12.5 min, as shown in Fig. 8; representing an extra fuel burn of 686 kg. For cold atmospheres with relatively low tropopause altitudes,[¶] the speed of sound at lower cruise altitudes could barely increase (i.e., the air temperature was almost constant), meaning that lower cruise levels would not represent an increase in true airspeed. In such a situation, the amount of delay that can be effectively recovered is significantly reduced since the only operational option to recover delay is to keep the same altitude and increase the cruise Mach number, which is already close to its maximum. For an example of Pilot3 optimized

[¶]The tropopause is the atmospheric boundary between the troposphere and the stratosphere where an abrupt change in the temperature lapse rate occurs: from a negative (and mostly constant) rate in the troposphere (i.e., temperature linearly decreases with altitude) to an almost null rate in the lower layers of the stratosphere.

trajectories in cold atmospheric conditions, the reader is referred to.[16]

When the delay at FL100 while flying the OFP is expected to be 40–45 min, the Pilot3 solution proposes to immediately descend to FL340 at the triggering point (i.e., the TOC for these case studies) and perform a second descent at about half of the cruise to FL300. The case for an expected delay of 45 min is shown in Fig. 9(d), where the two descents through the cruise are observed, while the cruising Mach number is increased close to MMO. For this case, the delay recovered is 14.9 min (see Fig. 8) and the extra fuel burnt is 977 kg.

When the expected delay ranges 50–60 min, the Pilot3 solution proposes to immediately descend to FL300 at the triggering point, and to maintain this altitude until the end of the cruise phase. The proposed Mach number also increases for those cases with higher expected delays. When the expected delay ranges 65–70 min, a descent to FL280 and an increase in the Mach number to its maximum (MMO) is suggested.

For expected delays at FL100 while flying the OFP ranges 75–110 min, the suggested cruise flight level drops to FL260 and the whole cruise is also performed at MMO. This is the trajectory that can recover the maximum amount of delay (22.5 min) at the expense of burning 2,936 kg of extra fuel (see Fig. 6). Figure 8 shows the cost break down for the case the expected delay is 110 min, while Fig. 9(e) shows the vertical and speed profiles of the Pilot3 optimized trajectory for this case.

As discussed in the previous section, contrary to simple rules of thumb, there are cases where recovering less than the maximum possible delay is better in terms of expected total costs – even if the expected delay at FL100 while flying the OFP is high. Recall Fig. 7, where it was seen that the delay recovered was not always monotonically increasing. For example, if the expected delay at FL100 while flying the OFP is 115 min, the delay recovered (17.8 min) is lower than in the previous case when the expected delay was 110 min and the suggested recovered delay was 22.5 min, as is the case for the extra fuel consumption (i.e., 1,446 kg instead of 2,936 kg). Figures 8 and 9(f) show this particular case, where it can be observed that the suggested cruise flight level be reduced to FL300 instead of FL260, as in the previous case. Furthermore, the cruise Mach number is now slightly lower than MMO. For an expected delay of 120 min, the trajectory is similar, but the cruise descent is to FL280.

For expected delays at FL100 while flying the OFP ranges 125–140 min, the trajectory leading to the maximum delay recovery is proposed again by Pilot3 (see Fig. 9(e)). For 145–150 min of expected delay, the amount of delay recovered by the Pilot3 solution is again lower than the maximum, and the resulting trajectories are similar to the one shown in Fig. 9(f). Finally, for expected delays more than 155 min, the delay recovered is again the maximum possible and the proposed trajectories are like the one shown in Fig. 9(e). Recall Fig. 6, where different *colored bands* were easily observed, meaning that the amount of recovered delay and corresponding extra fuel consumption do not always increase with the expected delay. In fact, they directly depend on the actual shape of the cost function.

5. Conclusions

Operations in the terminal maneuvering area (TMA) may induce a negative effect on the performance of a flight, mainly due to extending the total flight duration and increasing fuel consumption. Even if a flight has departed on time or earlier than expected, the uncertainties ahead might require some apparently counter-intuitive actions, such as speeding up the flight. The departing delay is known by the crew, but considering uncertainties ahead, and holding in particular, is critical to avoid sub-optimal decisions. This paper has introduced the consideration of holding in TMA in the optimization process of aircraft trajectories during the execution phase of the flight.

As presented, if holdings can be estimated with a high accuracy, considering it is equivalent to assuming a later departure. However, consideration of the expected cost of delay produces results that are more complex than rules of thumb. The amount of delay to be recovered not only depends on the departure delay, but also on the expected holding and the characteristics of the cost of delay for each particular flight. Adding uncertainties to the predictions modifies the shape of the cost function used by the optimizer making it *smoother* and providing solutions that are closer to average behaviors. Although these uncertainties might have a small impact on the results, the actual prediction of holding can be critical in situations where the expected arrival time is close to events that trigger additional costs, such as potential passenger missed connections.

The optimization presented in this paper could be incorporated in crew support tools such as the Pacelab Flight Profile Optimizer (FPO) developed by PACE.[4] Modern flight optimizers rely on cloud computing mitigating computational issues associated with limited capabilities of electronic flight bags (EFBs) or flight management systems (FMSs). As presented in this paper, there are two elements that play a key role on the results of the optimization: the cost function and the uncertainty on the operations.

In order to construct the cost function, as described in Section 2.1, Pilot3 would benefit from up-to-date operational data (e.g., information on airline fleet status to estimate reactionary delay) and passengers' itineraries with their characteristics (e.g., to estimate missed connections). Uncertainty might be present on these datasets, and estimators relying on either heuristics or machine learning could be used. Different estimators could be defined for each parameter, and the most relevant one used based on data availability and other operational limitations.[9]

The estimation of uncertainty, like holding as addressed in this paper, is also paramount to shape the cost function, and therefore produce solutions that minimize the expected cost. Advanced heuristics and machine-learning models could be used for this task[7] too.

How all of these factors are considered when building the cost function, and their impact on optimization can be complex will be subject of further research.[13] Once again, a cloud infrastructure with access to airline's real-time systems could be useful.

Finally, it is worth noting how the optimization approach described in this paper could be used in a distributed manner (i.e., each flight optimizing their own trajectory) or as part of a centralized system that considers several linked flights by the airline (e.g., flights with connecting passengers between them).

With all of these considerations, future work should focus on holdings predictions and the impact of their inaccuracy of these predictions on the performance obtained by the optimizer.

Acknowledgments

This work has been performed as part of Pilot3 innovation action which has received funding from the Clean Sky 2 Joint Undertaking (CS2JU) under grant agreement 863802. The CS2JU receives support from the European Union's Horizon 2020 research and innovation programme and the CS2JU members other than the Union. The opinions expressed herein reflect the authors' views only. Under no circumstances shall the CS2JU be responsible for any use that may be made of the information contained herein.

References

1) Ivanescu, D., Shaw, C., Tamvaclis, C., and Kettunen, T.: Models of Air Traffic Merging Techniques: Evaluating Performance of Point Merge, 9th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference, Hilton Head, South Carolina, September, 2009.
2) Boursier, L., Favennec, B., Hoffman, E., Trzmiel, A., Vergne, F., and Zeghal, K.: Merging Arrival Flows without Heading Instructions, USA/Europe Air Traffic Management Research and Development Seminar, July, 2007, pp. 1–8.
3) European Commission: Commission Implementing Regulation (EU) No 716/2014 of 27 June 2014 on the Establishment of the Pilot Common Project Supporting the Implementation of the European ATM Master Plan, 2014.
4) PACE: Aircraft Analysis and Evaluation, https://pace.txtgroup.com/products/aircraft-analysis-evaluation/, 2016 [Last accessed: Sep. 2022].
5) AVTECH: Clearpath, https://www.avtech.aero/clearpath/, 2022 [Last accessed: Sep. 2022].
6) Zammit-Mangion, D., Rydell, S., Sabatini, R., and Jia, H.: A Case Study of Arrival and Departure Managers Cooperation for Reducing Airborne Holding Times at Destination Airports, Proceedings of the 28th International Congress of the Aeronautical Sciences (ICAS 2012), International Council of the Aeronautical Science, 2012, pp. 4411–4420.
7) Dispatcher3 Consortium: Verification and Validation Report, Tech. Rep., December 2022, D5.2, Ed. 01.01.
8) NATS: Heathrow Holding Times on the Decline Thanks to New Technology, https://www.nats.aero/news/heathrow-holding-times-on-the-decline-thanks-to-new-technology/, 2016 [Last accessed: Sep. 2022].
9) Pilot3 Consortium: Crew Assistant Decision Model Description (Final Release), Tech. Rep., 2022, D4.3, Ed. 01.00.
10) European Commission: Commission Regulation (EC) No 859/2008 of 20 August 2008 Amending Council Regulation (EEC) No 3922/91 as Regards Common Technical Requirements and Administrative Procedures Applicable to Commercial Transportation by Aeroplane, *Official Journal of the European Union*, **2005** (2008), pp. 124–125.
11) Regulation (EC) No 261/2004 of the European Parliament and of the Council of 11 February 2004 Establishing Common Rules on Compensation and Assistance to Passengers in the Event of Denied Boarding and of Cancellation or Long Delay of Flights, and Repealin, *Official Journal of the European Union*, no. 295/91, 2004.
12) Cook, A. and Tanner, G.: European Airline Delay Cost Reference Values, Updated and Extended Values, University of Westminster, Tech. Rep., 2015, v4.1.
13) Pilot3 Consortium: Verification and Validation Report, Tech. Rep., March 2022, D5.2, Ed. 01.01.
14) Airbus: Getting to Grips with the Cost Index, Flight Operations Support and Line Assistance, Tech. Rep., 1998.
15) Dalmau, R., Melgosa, M., Vilardaga, S., and Prats, X.: A Fast and Flexible Aircraft Trajectory Predictor and Optimiser for ATM Research Applications, 8th Proceedings of the International Conference on Research in Air Transportation (ICRAT), Castelldefels, Spain, 2018.
16) Prats, X., de la Torre, D., and Delgado, L.: In-flight Cost Index Optimisation upon Weather Forecast Updates, 41st Proceedings of the AIAA/IEEE Digital Avitonics Systems Conference (DASC), 2022.
17) Gurtner, G., Delgado, L., and Valput, D.: An Agent-based Model for Air Transportation to Capture Network Effects in Assessing Delay Management Mechanisms, *Trans. Research Part C: Emerging Technologies*, **133** (2021), p. 103358.
18) Eurocontrol: DDR2 Reference Manual v.1.0.1, 2015.
19) BADA Overview: Base of Aircraft Data (BADA) EUROCONTROL's Aircraft Performance Model, 2015.
20) NATS: AIRAC 2111 (2021-11-04), 2021.

Takeshi Tsuchiya
Associate Editor