

**In-Flight Cost Index Optimisation Upon Weather Forecast
Updates**

Prats, X., de la Torre, D. and Delgado, L.

This is a copy of the author's accepted version of a paper subsequently published in the proceedings of the 41st Digital Avionics Systems Conference (DASC). Portsmouth, VA, USA 18 - 22 Sep 2022.

The final published version will be available online at:

<https://doi.org/10.1109/DASC55683.2022.9925796>

© 2022 IEEE . Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

The WestminsterResearch online digital archive at the University of Westminster aims to make the research output of the University available to a wider audience. Copyright and Moral Rights remain with the authors and/or copyright owners.

In-flight cost index optimisation upon weather forecast updates

Xavier Prats¹, David de la Torre¹, and Luis Delgado²

¹Department of Physics-Aerospace Division, Technical University of Catalonia (UPC), Castelldefels, Spain

²School of Architecture and Cities, University of Westminster, London, United Kingdom

Abstract—This paper presents an optimisation framework to compute the altitude and speed profiles of a trajectory in the execution phase of the flight, such that the expected total cost (ETC) of the operation is minimised (*i.e.* modelling the expected cost of delay and fuel – including arrival uncertainties – at the arrival gate). This is achieved with a two-stage optimisation strategy: a trajectory optimiser that minimises a generalised direct operating cost function, for a given cost index; and an upper-level optimiser, which obtains the best cost index that minimises the ETC. Several case studies are presented for different departure delays, while considering the impact of two different weather forecast updates too: a region with relative high head-winds appearing half way across the flight; and a cold atmosphere scenario, with a tropopause altitude lower than standard conditions. ETC savings with respect to following the operational flight plan increase with departure delay, as expected. Due to the non-linearity of the cost function, however, the benefits of considering the weather update depend on the actual value of the departure delay, showing the convenience of integrating the proposed approach into a crew decision support tool in order to avoid sub-optimal decisions.

Index Terms—Trajectory optimisation; cost index; cost of delay; weather update

I. INTRODUCTION

In the execution of a flight, when a change in operational conditions arises (*e.g.* significant temporal deviations with respect to the flight plan are observed, or a new weather forecast is available), the crew, and/or the dispatcher monitoring the flight, might consider to modify the trajectory with the aim to minimise the expected flight costs (and potential disruptions) for the airline.

Current on-board flight management systems (FMS), can optimise the aircraft trajectory using a generalised direct operating cost (DOC) function that links fuel and time by the cost index (CI), the weighting parameter that relates cost of time and cost of fuel ($\text{DOC} = \text{Fuel} + \text{CI} \cdot \text{Time}$). Without support from dedicated tools, however, aircraft crew is limited to explore alternatives by trial and error by manually selecting different CIs.

Cutting-edge pilot decision support tools, such as the Flight Profile Optimiser (FPO) developed by PACE [1] or ClearPath developed by AVTECH [2], are gradually being deployed in commercial aviation with the aim to compute tactical trajectory updates. These systems offer a better user interface, with a significantly improved look and feel (running in tablets or electronic flight bags); along with a better

connectivity with the airline operating center and third parties, such as weather providers. In fact, the capability to use much more accurate weather data than the FMS (which is heavily limited by hardware and software constraints), together with improved trajectory optimisation algorithms, make these systems to clearly outperform FMS capabilities.

Aircraft crew (and/or ground dispatchers) can usually introduce an estimation on the expected delay at arrival and then, they might be able to explore different alternative trajectories. Yet, these systems use the generalised DOC as defined prior-departure with the CI of the operational flight plan (OFP). This approach conditions the final optimisation on crew/dispatchers experience on a given route and focuses on delay rather than the expected cost that will materialise as a function of the arrival time at the gate.

The actual cost of delay depends on complex factors, such as passengers itineraries (including their connections), maintenance and crew costs, reactionary delay, etc. This renders a non-linear and, usually, step-wise cost function that is difficult to integrate on trajectory optimisers. Moreover, the arrival time at the gate is affected by (uncertain) taxi-in time and possible delays in the terminal airspace.

The use of a simple DOC estimation in the optimisation means that the crew/dispatcher is required to manually transform the trade-offs between fuel and time obtained by the optimiser into the expected cost for the airline and adjust the outcome of the optimiser accordingly. This can lead to a manual exploration of alternatives and/or the adoption of sub-optimal decisions if downstream uncertainty is not properly considered – such as to recover some delay, at a high fuel expense, with no significant benefit at the end.

Pilot3, a Clean Sky 2 Research and Innovation action, aimed at overcoming some of these limitations [3], [4]. Firstly, by estimating the expected cost of delay for a given flight considering the most up-to-date information and modelling not only the costs directly associated with the flight; but also the reactionary effects for the flight (modelling of downstream rotations and breaching of curfews); and passengers (with their possible missed connections and estimation of re-booking, including possible compensations due to EC Regulation 261 [5]).

Secondly, by predicting time and fuel uncertainties at arrival using heuristics and/or machine learning estimators and integrating these in the expected cost of delay function.




Fig. 1. Trajectory optimisation scope for Pilot3

In this context, potential holdings and path stretching for sequencing and merging purposes in terminal airspace, along with and taxi-in, were considered.

Finally, by optimising the trajectory from the point Pilot3 is triggered (in flight) to the (expected) landing runway considering the expected total cost (ETC) for the flight. This is achieved with an optimisation framework composed of two stages: a low-level trajectory optimiser that minimises DOC for a given CI (considering weather updates, or operational constraints that might arise in the execution of the flight, if any) and generates trajectories which trade fuel and time; and an upper-level optimiser, which aims at obtaining the best CI such that the ETC at the arrival gate is minimised.

This paper presents this optimisation framework and how it is able to find this optimal CI, along with the associated trajectory, in order to minimise the ETC when departure delay is experienced, while also considering weather forecast updates. The expected costs (and trajectories) obtained by the Pilot3 optimiser are compared with two baseline trajectory plans: keep flying the OFP and optimising the trajectory without considering the new weather data (*i.e.* using Pilot3 to only react to departure delay).

II. OPTIMISATION FRAMEWORK

This section describes the scope of Pilot3 in-flight optimisation; then, the proposed two-stage optimisation is presented; and finally, some discussion regarding the benefits and limitations of the proposed approach is given.

A. Trajectory optimisation scope

Pilot3 could be triggered at any point of the flight and only the vertical and speed profiles are optimised. The lateral optimisation is out of the scope of current implementation, since it is considered that the flight has obtained a route clearance that might be difficult to update in flight within current air traffic management (ATM) paradigm.

As depicted in Figure 1, the trajectory is optimised from the current aircraft state (*i.e.* the moment Pilot3 is triggered) down to reaching FL100 (*i.e.* 10,000 ft) at the proximity of the destination. The trajectory plan from FL100 to the runway is computed assuming standard operations and using a fixed sequence of aircraft intents that depend on the aircraft type and potential altitude/speed constraints

in the arrival and approach procedures. The reason of not optimising the trajectory below FL100 is twofold:

- 1) the optimisation control space is significantly reduced since the aircraft is flying near the limits of the flight envelope (*i.e.* min/max speeds), the standard operating procedures constraint significantly the trajectory (*e.g.* approach speeds, glide path on the instrument landing system - ILS) and ATM strategic constraints might also be in place (*e.g.* speed and altitude limitations for certain legs);
- 2) the aircraft trajectory is likely to be modified several times by tactical ATC intervention, thus forcing the pilot to no longer follow the Pilot3 plan.

Note, however, that this does not mean that a fixed amount of time (and fuel) are considered for this final portion of the flight, but that the aircraft intents (*i.e.* how the flight is operated) are only fixed. Operational uncertainties at arrival (*i.e.* holdings, extra distance in path stretching for sequencing and merging purposes, and taxi-in times) are considered in order to compute the expected total cost (ETC) at the arrival gate. Uncertainties up to FL100 either by route changes and in the weather forecast are out of scope of the current implementation of Pilot3.

B. Optimisation approach

Figure 2 presents the optimisation framework developed in Pilot3, which is composed by several modules. This paper focuses on the trajectory optimiser and, to some extent, on the Objective Function Estimator, a key element needed to compute the ETC needed by the optimiser. As shown in the Figure, the generation of the cost function relies on the outputs of other modules of the Pilot3 tool, which are out of the scope of this paper (details can be found in [3], [4]).

1) *Objective Function Estimator*: The cost of delay, as expected at the arrival gate, is estimated by the Performance Indicators Estimator. This cost is composed of two components which are estimated as explicitly as possible [6], as shown in Figure 3:

- IROPs (irregular operations) cost, generated by passengers' (connecting and non-connecting) disruptions. This includes soft costs, associated to dissatisfaction for late arrival; and hard costs, due to compensation

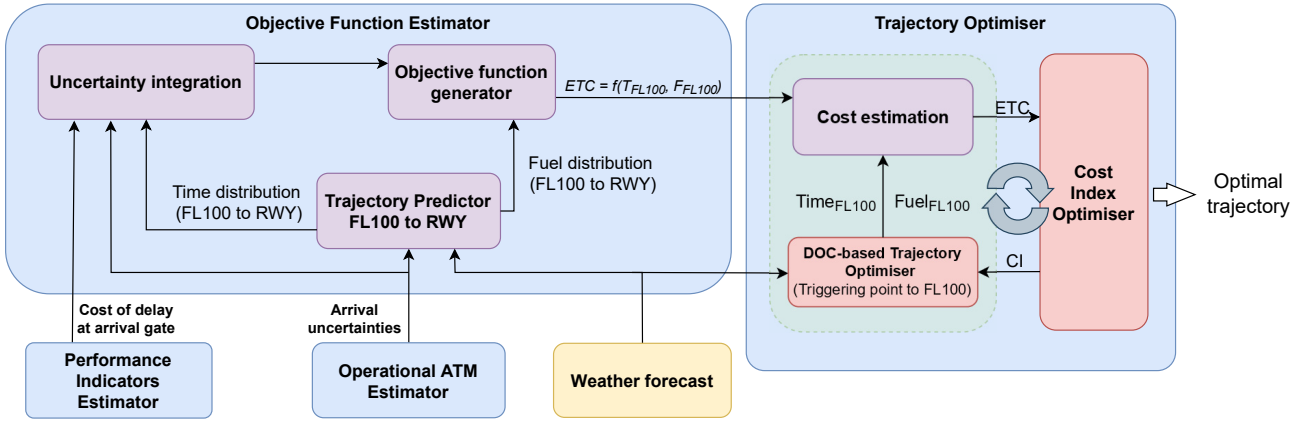


Fig. 2. Trajectory optimisation framework

as in Regulation 261 [5] and transfer and duty of care costs for passengers missing their connections. IROPs costs are non-linear and tend to be step-wise, as costs are triggered by events such as missing connections, as observed in Figure 3.

- Other costs which are not related to passengers management: crew and maintenance costs (as reported in [6]) and reactionary delay costs [7] from the propagation of delay in subsequent flights, the potential breach of curfew, and expected pre-tactical fleet management actions, such as aircraft swapping or cancellations [4].

Even with a certainty on the arrival time at the gate, the estimation of these costs can be complex and uncertain: for example, for the same arrival time passengers might or might not miss a connection depending on the status of the rest of fleet of the airline. The Performance Indicators Estimator uses the most up to date available information to model these uncertainties [3].

As presented in Section II-A, Pilot3 considers the uncertainties associated with the operations at arrival: holding time, sequencing and merging distance (*i.e.* distance from FL100 to the runway), and taxi-in time. These parameters

are estimated by the Operational ATM Estimator (see Fig. 2), which provides not only the expected (or average) time of these processes, but their probability distribution [8].

A trajectory predictor is then used to compute the fuel and time distributions for these processes by numerically integrating the trajectory, in backwards, from the runway threshold to FL100. This trajectory predictor uses the weather forecast available; aircraft performance models and data from the EUROCONTROL's base of aircraft data (BADA) [9]; and a predefined sequence of aircraft intents that are aircraft- and procedure-dependent.

The convolution of the time distributions for holding, sequencing and merging and taxi-in¹ yields to the distribution of the time required to reach the arrival gate from FL100. Therefore, the expected cost of delay can be computed for a given arrival time at FL100, calculating in this way the distribution of arrival times at the gate and using the cost of delay at the gate computed by the Performance Indicator Estimator [3], [4]. The use of distributions of arrival time, instead of only average values, is important due to the non-linearity of the cost of delay function.

As shown in Figure 3, this mathematical process applied for all possible arrival times at FL100, yields to a temporal shift of the cost of delay function and a smoothing of the expected cost of delay function.

The Objective Function Estimator adds the cost of the expected fuel used from FL100 to reaching the arrival gate and the expected cost of delay to provide to the Trajectory Optimiser with a function that translates the arrival time at FL100 and the fuel required from the triggering point of Pilot3 to FL100 into the ETC. The ETC will therefore include the expected cost of the fuel from triggering point to the gate and the expected cost of delay.

2) *Trajectory optimisation:* Different CI lead to different trajectories that minimise a general direct operating costs (DOC) cost function ($DOC = Fuel + CI \cdot Time$); and therefore, to different ways to trade flight time and fuel consumption.




Fig. 3. Cost of delay as a function of arrival time at gate or FL100.

¹Note that in some cases some of these processes might not be present, *e.g.* an airport not using holdings at arrival




Fig. 4. CI-based optimisation

Thus, a CI optimiser is implemented in the Trajectory Optimiser module (see Fig. 2) aiming at finding the best CI that produces a trajectory for which the ETC is minimised.

A simple, but robust, a recursive grid search method has been implemented to find the optimal CI: the cost function is evaluated in regularly and coarsely spaced CIs between arbitrary bounds, finding a first optimum. Then, the process is recursively repeated by setting a finer distance between CIs and as new bounds the CIs of the previous grid adjacent to the optimum, until an integer CI precision is reached.

Figure 4 provides a visual example of this optimisation approach. The blue circles represent the ETC corresponding to the DOC-optimal trajectories for a given CI, while the red cross shows the ETC optimal solution found at each iteration of the algorithm. The figure shows how the algorithm converges to a CI that minimises the ETC by refining the search around the optimum after each iteration.

For each CI explored in this process, a trajectory optimisation algorithm is launched aiming at minimising the conventional DOC function given above. Dynamo is used for this purpose, which consists on a high-resolution trajectory optimiser based on an aircraft point-mass model (*i.e.* 3 degree of freedom) that is able to use realistic weather data and accurate aircraft performance data [10]. In the current version of Pilot3, only Dynamo's altitude and speed trajectory optimisation capability is used, which is based on an advanced and fast grid-search method.

Although Dynamo is flexible enough to specify different kinds of vertical profiles, the implementation embedded in Pilot3 assumes a standard CAS/Mach climb (and a standard Mach/CAS descent) profile, along with potential speed and/or altitude constraints depending on the specific departure, arrival or approach procedure being flown.

The value of the optimal CAS/Mach pair for climb (and descent) are taken from pre-computed tables, where these optimal speeds are given as a function of the aircraft type, CI, aircraft mass and atmospheric conditions. These speeds are typically known as ECON speeds (from *economic*), since they optimise a conventional DOC function and can be computed in advance for a wide set of input combinations. Then, when running the optimiser, only a table look-up is needed, making this process very efficient from a

computational point of view.

Regarding the cruise phase, Dynamo determines the best altitude profile using a grid-search methodology: at arbitrarily-defined decision points, a feasible collection of discrete cruise flight levels is explored and the one with the lowest cost is selected. The optimal cruise speed for each potential flight level is also selected from pre-computed tables. Is it worth noting that this cruise speed (*i.e.* ECON Cruise Mach) is regularly updated along the cruise (even if the cruise flight level is kept constant) in order to account for changes in aircraft mass and/or atmospheric conditions. A set of inhibit distances after each resulting step-climb or step-descent (and also after the top of climb and before the top of descent) ensure the trajectory plan is operationally sound. In this line, all potential flight levels where the aircraft would not be able to maintain a minimum rate of climb are discarded. For this paper, this minimum rate of climb available is set to 500 ft/min, which corresponds to the typical constraint set in most airspace.

For each candidate trajectory resulting from this grid-search, a numerical integration is performed using the aircraft equations of motion (*i.e.* the point-mass model), the weather, data and the aircraft performance models (and data). As a result of the integration the flight time and fuel consumption for each option is determined and, using the CI, the DOC for each option is computed. Finally, the algorithm selects the option with lower DOC.

The algorithm builds the trajectory in a modular and iterative manner, ensuring continuity in mass, time, speed and altitude throughout the flight due to the fact the equations of motion are integrated. The result of the optimisation is a four-dimensional trajectory (lateral and vertical flight profiles along time) from the triggering point until the arrival airport. Nevertheless, other operational or flight variables are also available, such as the fuel on board, wind components along the route, operational speeds, aerodynamic lift and drag, etc.

Note that the value of CI is given as an input parameter for the optimisation and is therefore constant for the whole flight. Thus, for a given flight (*i.e.* for a given aircraft type, initial conditions, weather conditions, etc.), the CI determines the choice of altitudes and speeds (the whole vertical profile in fact) and, by extension, the resulting fuel consumption and flight time.

Wrapping up, the main outputs of Dynamo within the Pilot3 framework are, for a given CI, the fuel consumption from the triggering point until FL100; and the arrival time at FL100. As explained above (recall Fig. 2), these values will be explicitly used by the CI optimiser in order to obtain the trajectory (*i.e.* the CI) that minimises the ETC.

C. Approach benefits and drawbacks

The approach described presents a set of benefits. First, the outcome is an easier to communicate trajectory (as it is based on a single optimal CI) that could be computed with any CI-based (*i.e.* DOC-based) trajectory optimiser, but

TABLE I
MAIN CHARACTERISTICS TO COMPUTE THE OPERATIONAL FLIGHT PLAN (OFP).

Flight schedule	Flight dispatch	Other operational information
Airline: Lufthansa	Cost Index (CI): 10 kg/min	Estimated taxi-out: 10'
Aircraft type: A321-111	Cost of fuel: 0.5 Eur/kg	OFP trip time: 138'
Stage: LEMD - EDDF	Payload: 171 PAX [†] + 1,000 kg Cargo	Taxi-in and padding at arrival: 7'
SOBT: 06h35 UTC	EMPAX 1C arrival + ILS RWY07C approach	65 connecting PAX
SIBT: 09h10 UTC	Weather forecast issue/applicability: 2016-07-28 00h UTC	

[†]According to the EU-OPS 1.620 [11] flights “within the European region” shall account 84kg (+ 13kg of luggage) per adult passenger.

SOBT/SIBT: Scheduled off/in-block time – UTC: Coordinated universal time – PAX: passengers – ILS: instrumental landing system – RWY: runway

which minimises the ETC of the flight. Secondly, there is only one control variable used in this ETC optimisation: the CI. Thus, the CI is used as proxy to control the trade-off between time and fuel of the trajectory existing in the ETC function. This simplifies the optimisation. Third, the framework is modular and flexible. The cost function is computed independently of the Trajectory Optimiser and if new information is available it can be updated independently of the optimisation approach. Similarly, new improvements can be provided for the different components independently, *e.g.* a different DOC-based optimiser or Cost Index optimiser could be used.

The main drawback of the presented approach is that the space of search for the generation of the trajectories (vertical and speed profiles explored) are limited by the capabilities of the CI based trajectory optimiser. In some cases the same arrival time might be achievable with lower fuel if the trajectory is modified independently of the restrictions related to use a constant CI relationship for the whole trajectory as the current implementation of the trajectory optimiser relies on the use of pre-optimised speed tables as a function of CI, airplane mass and atmospheric conditions. In other words, a better trajectory could be found with different segments flown at different CI.

III. SCENARIO AND CASE STUDIES

To assess the optimisation framework presented in this paper, a flight from Madrid (LEMD) to Frankfurt-Main (EDDF) is modelled as described in Table I.

Individual passenger itineraries (with their connections) are modelled based on historical data from IATA's PaxIS and Global distribution Systems datasets, as in previous research projects [12]. Figure 5 shows the different passenger groups with connections at EDDF for the flight under study. 65 passengers (38%) have an onward connection. The number of passengers of each group is indicated along the time where their connecting flight is scheduled to depart.

A significant number of passengers (greater than 40) have a flight departing earlier than 12h. A first increment on passenger related costs is then expected around 10h20, which would correspond to the group of 20 passengers missing their connection (considering the minimum connecting time, if the flight arrives after 10h20 they won't be able to reach their connecting flight [12]). This can be observed in Figure 3. The fact that passenger miss connections do

not necessarily increase significantly the expected cost of delay if, for example, they can be re-accommodated and arrive to their final destination before being entitled to compensation due to Regulation 261.




Fig. 5. Passenger groups connecting at EDDF for the flight subject to study

A. Operational flight plan (OFP)

The OFP has been generated as follows: the route (*i.e.* sequence of waypoints) is obtained from EUROCONTROL's Demand Data Repository 2 (DDR2) [13]; then, the altitude and speed trajectory profiles are optimised with Dynamo (see section II-B2), using aircraft performance data and models from BADA v4.2 and weather forecast from the European Centre for Medium-Range Weather Forecasts (ECMWF) ERA5 models ². The optimisation criterion for this optimisation is the generalised DOC described above, assuming that the cruise Mach is kept constant for a given cruise flight level.

For flight and fuel planning purposes, EMPAX 1C arrival is considered, as it is the longest arrival in Frankfurt. Since this airport operates a tromboning procedure at approach, the German AIP, as in the AIRAC 2013 (issued December 2020) [14], requests the operators to consider 83 NM from SPESA to the landing runway as average flight distance for fuel planning purposes.

Figure 6 shows the OFP trajectory with the climb, cruise and descent phases represented, respectively, by green, blue and red segments. Figure 7 shows the resulting vertical and speed profiles of the OFP trajectory with the along-track and cross-wind components at different altitudes (coloured

²<https://www.ecmwf.int/>

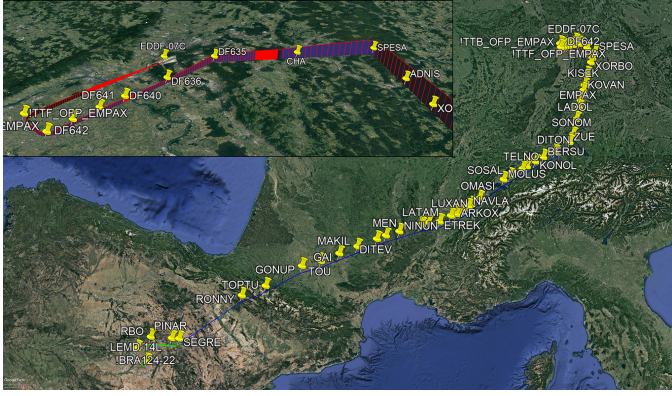

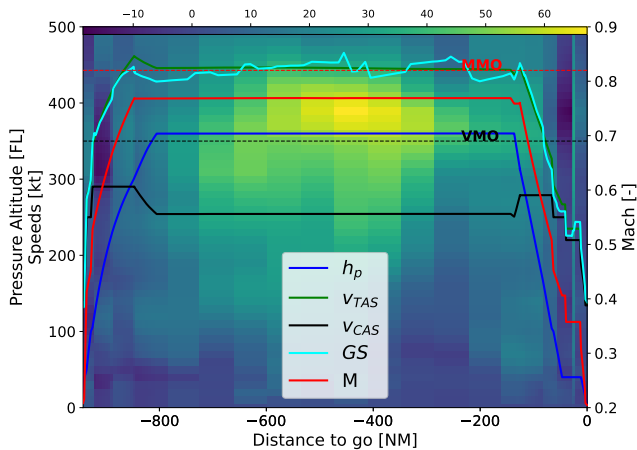


Fig. 6. OFP route Horizontal trajectory profile (Detail of the descent trajectory in the upper-left side).



(a) Background colour: along-track wind.



(b) Background colour: cross-track wind.

Fig. 7. OFP vertical and speed trajectory profiles.

backgrounds). In these plots, pressure altitude (hp) for the whole trajectory is depicted together with Mach number (M), calibrated airspeed (CAS), true airspeed (TAS) and ground speed (GS). It is worth mentioning that the sudden changes in wind components in these figures are due to

track changes in the lateral route, which change the relative wind direction with respect to the route. Finally, these plots also depict the maximum operational speeds for that aircraft type: MMO (maximum Mach in operation) and VMO (maximum CAS in operation). As observed in the Figure, the OFP for this scenario consists on a cruise at FL360 and Mach 0.77.

B. Definition of case studies

First, for all case studies in this paper it is considered that all passengers are entitled to Regulation 261 compensation if delay thresholds at their final destination are met [5]. This leads to the cost of delay function presented in Figure 3.

The case studies consider that Pilot3 is triggered shortly after 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.

1) *No weather update*: A first set of case studies explore the solutions given by Pilot3 for different departure delays with a nominal weather forecast. A flight might depart late for a combination of factors: leaving the gate with deviation with respect to schedule; taxi-out time different than planned; route shortcuts (or path stretching) in the departure phase of the flight; experiencing different weather conditions than those used when planning the flight, etc. These time differences translate into differences, with respect to the OFP, in the time of arrival at the TOC.

For these cases, the weather forecast used by Pilot3 to re-plan the trajectory at the TOC is the same that was used to compute the OFP (*i.e.* no weather forecast update is considered at the TOC). Thus, these case studies aim to illustrate the trade-off between delay recovery and extra fuel usage, and how this results in a new trajectory plan optimised by Pilot3. Time differences at the TOC ranging from -10 to 60 minutes at 5 minute intervals are simulated.

2) *Weather updates:* A second and third set of case studies consider that a new weather forecast has been up-linked to the aircraft and it is considered by Pilot3 when re-planning the trajectory at the TOC.

In particular, two different forecasts are considered: a case where a region of relatively high headwinds (ranging between 70 to 40 kt) is observed, approximately, between FL290 and FL360, for the second half of the flight; and a case with a cold atmosphere, implying a (low) tropopause altitude around FL250³.

The ECMWF ERA5 weather forecast corresponding to 2018-04-17 with issue/applicability time of 00h UTC was used to simulate the first case (headwind). The forecast of 2018-02-07 with issue/applicability time of 00h UTC was used for the second case (cold atmosphere).

³The tropopause is the atmospheric boundary between the troposphere and the stratosphere where an abrupt change in the temperature lapse rate occurs: from a positive (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






Fig. 8. Pilot3 solution vs. keep flying the OFP



(a) Delay at TOC: 25'



(b) Delay at TOC: 45'

Fig. 9. Examples of vertical and speed profiles for the Pilot3 optimisation without weather forecast update

IV. RESULTS

This section presents the results for all case studies presented in the previous section. In order to validate our approach, the Pilot3 optimised trajectory plan is compared with two baseline trajectory plans: keep flying the OFP and optimising the trajectory without considering the new weather data (*i.e.* only reacting to departure delay). The main metric selected for comparison is the expected total costs (ETC) of both trajectory plans.

A. No weather update

Figure 8 shows the results for the first set of case studies, where different delays at TOC are explored and the weather

forecast used by Pilot3 is the same used when computing the OFP. Recall that time, fuel and expected savings are obtained by comparing the Pilot3 solution (that reacts to delay at TOC) with a baseline solution consisting in keep flying the OFP. As observed in the figure, as delay at TOC increases, the CI selected by the optimiser (*i.e.* the one minimising the ETC) increases too, from 0 to 279 kg/min.

For the case where the aircraft arrives 5' earlier than expected at the TOC, the resulting optimised trajectory suggests to slightly slow down (arriving 1' later than the OFP), in order to save some fuel. This is achieved by flying at CI 0. Yet, the cost savings compared with the baseline trajectory are negligible. This is due to the relatively low CI used to compute the OFP (10 kg/min).

Then, for positive delays at the TOC, the Pilot3 solution progressively recovers delay – at the expense of burning more fuel – yielding in all cases to some cost savings (if compared with the costs incurred if the aircraft keeps flying the OFP). The amount of delay recovered therefore varies from extra 1 minutes of delay accrue to just over 11 minutes of saving. As observed in the figure, cost savings are higher for higher delays at the TOC, since the cost of delay grows quickly some minutes after the SIBT and therefore, recovering some delay yields high expected savings (see Figure 3). The benefits obtained range from being negligible, for very low delays at TOC, to close to 2,100 €.

The results obtained are aligned with the expected behaviour used by airlines, *i.e.* the highest the initial delay the more delay is tried to be recovered. The optimisation provided by Pilot3, however, ensures that the amount of delay recovered (trading fuel) is optimal considering the total expected costs, preventing the crew to try to recover *too much* delay without a clear benefit.

Finally, it is worth noticing how beyond 45' of delay at TOC the optimal CI does not increase anymore. This is due to the fact that the maximum delay that can be recovered has been reached (for such a relative short flight), since (ground) speed cannot be further increased. Hence, fuel and time differences for both trajectories remain constant beyond this point. Yet, it is worth noting how the expected cost savings do increase for delays higher than 45' even if the delay recovered is maintained at around 11'. This is due to the non-linearity of the cost of delay, as shown in Fig. 3.

In order to recover delay, as the lateral trajectory is maintained, the flight needs to increase its ground speed (GS). This can be achieved by either increasing the cruise Mach and descent speeds; and/or by modifying the vertical profile. In general, lower flight levels would represent an increase in true airspeed (TAS) for a given Mach, since the speed of sound (that directly depends on air's temperature) is typically higher at lower altitudes.

For delays at the TOC of 5', 10', 15' and 20', the increase in GS is achieved by increasing the Mach number in cruise and the Mach/CAS pair of the descent, thus slightly pushing forward the top of descent. When the delay at the TOC is 25', Pilot3 also proposes a similar increase in cruise

Mach and descent Mach/CAS, but also to descend to FL340 approximately 50 NM after the TOC (see Figure 9(a)). This plot compares the speed and vertical profiles for optimised and baseline trajectories: solid lines correspond to the Pilot3 solution, while pale lines represent the OFP trajectory.

For 30' and 35' of delay at the TOC, the resulting trajectory is essentially the same as the previous one, but the descent to FL340 is suggested right after the TOC. For these cases the Mach number is already the maximum allowed⁴. Then, for 40' of delay at the TOC, Pilot3 suggests to immediately descent to FL300. Since the Mach number in cruise cannot longer be increased, the only way to increase the GS is to fly at a lower altitude. Hence, for 45' of delay at the TOC, the proposed new cruise altitude is lowered to FL260, reaching the maximum delay that can be recovered in this flight (see Figure 9(b)), since altitudes below FL260 have not been allowed for these experiments assuming they are not operationally valid.

The results obtained are as expected on these type of delay recovery strategies, *i.e.* as flying time is reduced, cruising altitude decreases while speed increases.

B. Headwind weather update

For the second set of case studies it is assumed that a new weather forecast update is made available at the TOC, presenting a significant headwind zone in the second half of the flight. Figures 10(a) and 10(b) show the vertical profiles of the Pilot3 solutions when optimising the trajectory at the TOC with the new weather profile (white trajectories) and compares them with Pilot3 solutions obtained when the nominal weather forecast is used (*i.e.* the trajectories presented and discussed in section IV-A and shown in red in the Figure). Recall that when optimising with the nominal weather forecast, a gradual usage of lower FLs is observed in order to increase the GS.

As observed in the figure, the solutions that take into account the new weather conditions behave differently. For low values of TOC delay (up to 20' as shown in Fig. 10(a)), Pilot3 suggests to perform a step climb halfway to destination approximately (the exact location depends on the particular TOC delay) in order to avoid the headwind area by flying above it. The only way to further increase the GS and recover more delay is to fly at lower altitudes (as discussed in Section IV-A). With the new weather update the optimal trajectory for delays higher than 20' consists on a descent to FL260 in order to avoid cruising inside the headwind region (see Fig. 10(b)). As seen in the figure, for the case of 25' of TOC delay, this abrupt step descent is done at approximately 250NM before the destination, while for 30' of TOC delay it is suggested to be initiated 50 NM after the TOC, while for higher delays Pilot3 suggests to immediately descent to FL260 and recover as much delay as possible.

⁴In this paper it has been chosen as the Maximum Mach in Operations (MMO) minus an operational buffer of M0.02

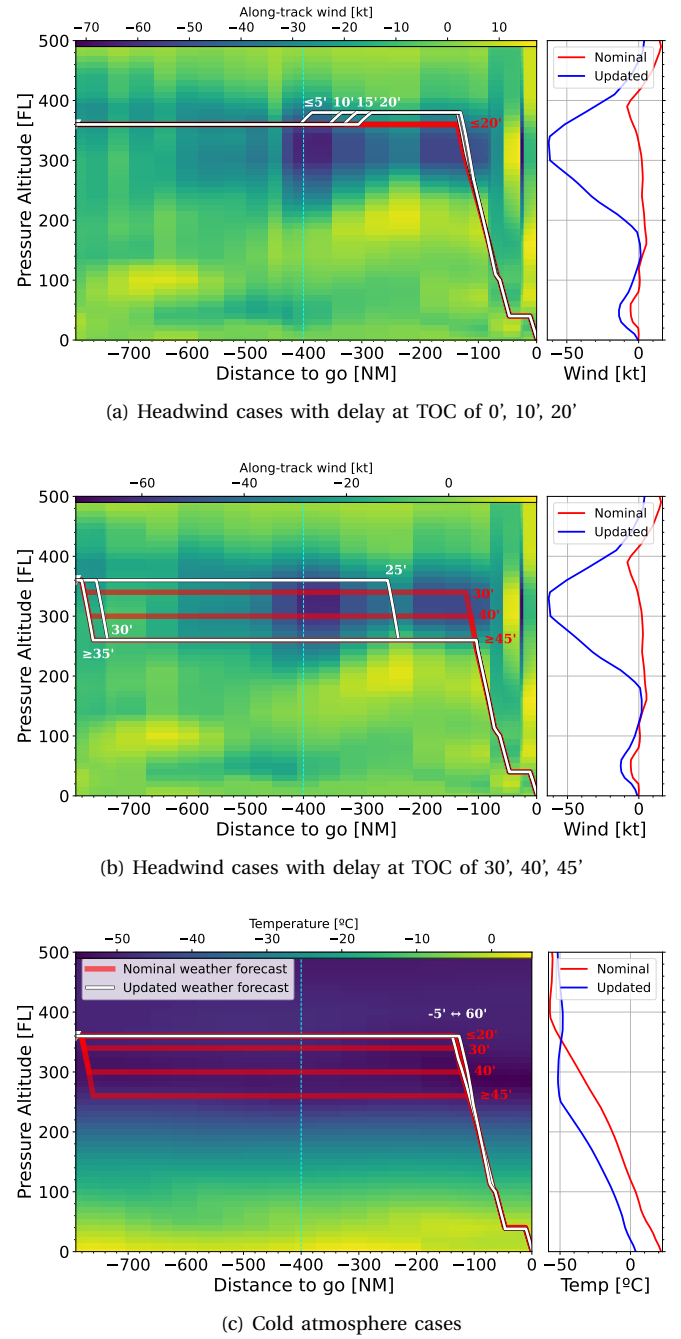


Fig. 10. Along-track wind or temperature profiles along the route corresponding to the updated weather forecast (coloured background). Vertical section (at -400 NM) comparing the two forecasts. Altitude profiles of the different P3 solutions (labeled with the corresponding TOC delay)

Figure 11(a) shows the results of the different optimisations, as a function of the delay at the TOC, when comparing the Pilot3 optimised trajectory (reacting to delay at TOC and considering the new weather update) with the trajectory consisting on keep flying the OFP, but simulating it with the new weather conditions. In order to perform this simulation, the altitude and speed (*i.e.* Mach or CAS) profiles of the OFP are numerically integrated with Dy-

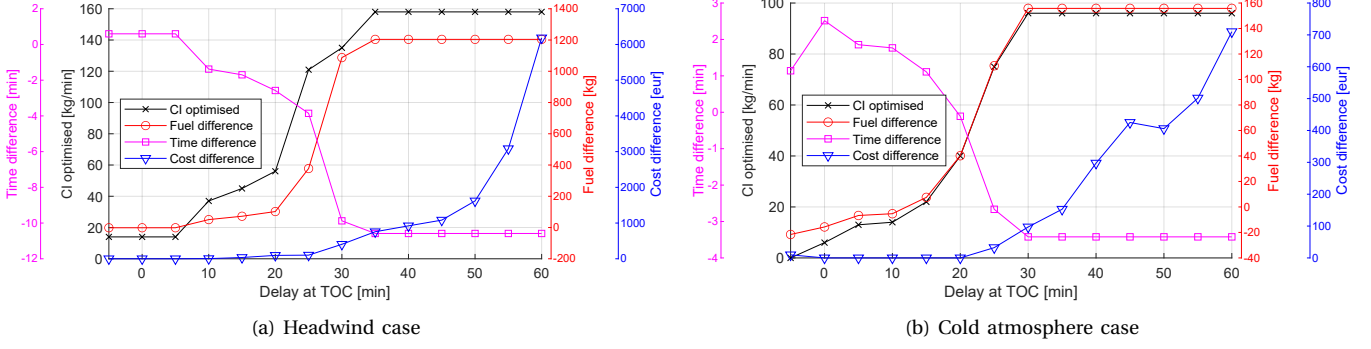


Fig. 11. Pilot3 solution using the weather update vs. keep flying the OFP (simulating the OFP with the atmospheric conditions of the weather update)

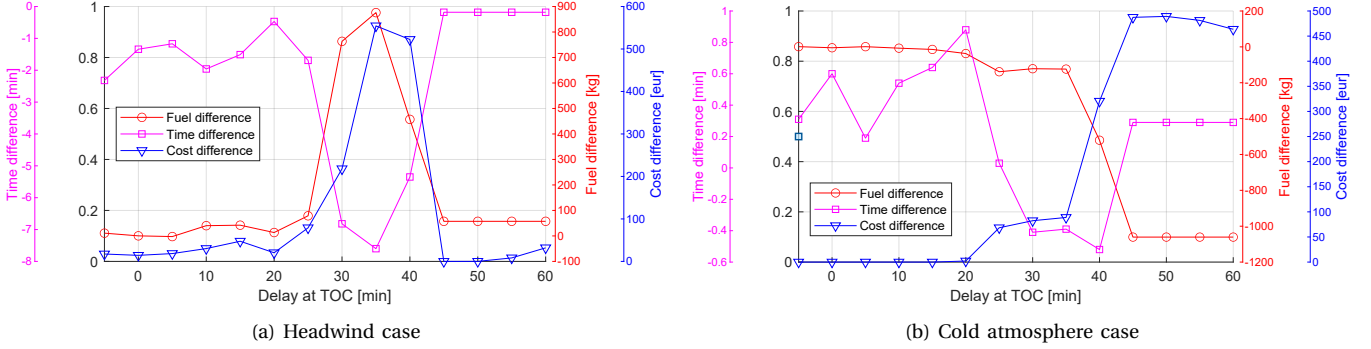


Fig. 12. Pilot3 solution using the weather update vs. Pilot3 solution using the nominal weather forecast (simulating the resulting trajectory with the atmospheric conditions of the weather update)

namo using the new weather conditions (hence yielding to a different GS profile if compared with the OFP and, consequently to different flight time and fuel consumed).

As in the case where the weather is not updated (Fig 8), for larger delays more time is recovered at the expense of burning more fuel. The ETC savings, however, are higher in this case reaching more than 6,000 € for 60' of delay at the TOC. Note, however, how the maximum CI used is lower (158 kg/min instead of 279 kg/min of the no weather update case), meaning that the maximum delay that can be recovered is achieved earlier for delays at TOC of 35'.

The abrupt descents to FL260 proposed by Pilot3 change significantly the amount of delay recovered. As shown in Figure 12(a), the differences in flight time, comparing the two Pilot3 solutions, can reach up to almost 8' for the case with TOC delay of 35' (still significant for such a short flight). This leads to fuel differences around 900 kg and a difference in ETC of 550 €.

It is worth noting that significant cost savings are obtained only in the interval ranging from 25' to 40' of TOC delay. This is because these are the cases where the Pilot3 solution that does not take into account the weather update plans a cruise within the strong headwind zone, thus leading to non-optimal trajectories. For TOC delays greater than 45' the two Pilot3 trajectories are essentially the same (*i.e.* immediate descent to FL260), while for TOC delays of 20' or lower the fact to consider the new weather

update bring marginal benefits (obtained with the step climbs described above).

C. Cold atmosphere weather update

For the third set of case studies it is assumed that a new weather forecast update is available at the TOC, showing a colder atmosphere along the flight (recall Section III-B2). A colder atmosphere also means a lower tropopause altitude. As observed in Fig. 10(c), the tropopause altitude for the nominal weather forecast is around FL380, while for the updated forecast it lies around FL250. This means that air temperature is more or less constant above this altitude and therefore, the speed of sound as well.

Since the speed of sound is approximately constant above FL250, altitude changes above this level will make almost no difference in terms of TAS for a given Mach number. Recall that Pilot3 quickly reaches the maximum Mach number allowed and for higher delays at TOC the only way to recover delay is by flying to lower altitudes (with, in principle, hotter air). This gradual usage of lower cruise altitudes was seen in section IV-A, for the nominal weather forecast results (depicted as red trajectories in Fig. 10(c)). With the updated weather forecast, however, these descents are not longer useful, because for a given Mach number the TAS does not increase. Since no big differences in winds aloft are observed in this layer of the atmosphere, the GS does not increase either. This is why the optimal

trajectory proposed by Pilot3 keeps cruising at the original OFP altitude (FL360), as depicted by the white trajectories in Fig. 10(c). Only slight differences are appreciated in the optimal position of the top of descent.

Like in previous section, Figure 11(b) shows the results of the different optimisations, as a function of the delay at the TOC, when comparing the Pilot3 optimised trajectory (reacting to delay at TOC and considering the new weather update) with the trajectory consisting on keep flying the OFP, but simulating it with the new weather conditions.

As in the case where the weather is not updated (Fig 8), for larger delays more time is recovered at the expense of burning more fuel. Yet, due to the impossibility to increase the GS significantly no much delay can be recovered (around 3.5' at the most, instead of the 11' obtained without the weather update), leading also to lower cost benefits when compared with keeping the OFP (approximately up to 700 €). The limitation to recover more delay for this particular weather scenario is also observed by looking at the selected CI, which quickly saturates to 96 kg/min.

If we compare the Pilot3 optimisations (using the new forecast vs. using the nominal one) the results shown in Fig. 12(b) are obtained. For delays at the TOC of 20' or lower the two trajectories are essentially the same (recall Fig. 10(c)), leading to the same cost figures. Cost savings become apparent for 40' of TOC delay or higher, since these are the cases where the Pilot3 optimal trajectory that considers the nominal forecast suggests to fly at lower altitudes, leading in reality to sub-optimal decisions due to the (actual) colder atmosphere.

V. CONCLUSIONS

The optimisation framework presented is able to generate a trajectory which minimises the expected total costs (ETC) for the flight considering uncertainties at arrival. This is achieved by implementing a two step optimisation process: for a given cost index (CI) an optimal trajectory is obtained by minimising a classical cost function trading time and fuel; and an upper-level optimiser selects the best CI such as the ETC are minimised.

In general terms, when a flight has to trade fuel for time this is achieved by increasing the speed (operating at a higher Mach, and/or by selecting a lower cruising altitude). These results, however, are severely impacted by the weather conditions. As shown in this paper particularly when non-nominal weather is encountered.

The optimisation framework discussed is flexible but relies on having an estimation of the cost function as updated as possible. This presents some challenges as the materialisation of cost can be uncertain and depend on factors external to the flight. Future research should address how these factors are considered.

Finally, the space of search is limited by the outcome of the CI based trajectory optimiser, *i.e.* only solutions which are feasible by modifying the CI can be explored. This approach is effective and produce results which can

be validated, *i.e.* a cost index which can be reviewed by the crew and ground operators. However, it could be possible that more complex trajectories (not feasible by optimising only the direct operating cost related with the CI which is maintained along the route) can produce even lower expected operating costs. To achieve this another optimisation framework should be used, *e.g.* integrating the expected cost of delay into a grid-search algorithm which optimises both speeds and flight levels. This should be further researched.

ACKNOWLEDGMENTS

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

REFERENCES

- [1] PACE. Aircraft analysis and evaluation. <http://pace.txtgroup.com/products/aircraft-analysis-evaluation/>, 2016.
- [2] AVTECH. Clearpath. <http://www.avtech.aero/clearpath/>, 2022.
- [3] Pilot3 Consortium. Crew assistant decision model description (final release). Technical report, Pilot3 Consortium, 2022. D4.3, Ed. 01.00.
- [4] Pilot3 Consortium. Verification and validation report. Technical report, Pilot3 Consortium, 2022. D5.2, Ed. 01.01.
- [5] European Commission. 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 repealing Regulation (EEC) No 295/91, 17/02/2004, 2004.
- [6] Andrew Cook and Graham Tanner. European airline delay cost reference values, updated and extended values. Technical report, University of Westminster, 2015. v4.1.
- [7] Paolino De Falco and Luis Delgado. Prediction of reactionary delay and cost using machine learning. In *Airline group of the International Federation of Operational Research Society (AGIFORS)*, 2021.
- [8] Luis Delgado, Gerald Gurtner, and Andrew Cook. Pilot3 crew multi-criteria decision support tool estimating performance indicators and uncertainty. In *Airline group of the International Federation of Operational Research Society (AGIFORS)*, 2020.
- [9] EUROCONTROL. BADA overview: Base of Aircraft Data (BADA) EUROCONTROL's aircraft performance model, 2015.
- [10] Ramon Dalmau, Marc Melgosa, Santi Vilardaga, and Xavier Prats. A Fast and Flexible Aircraft Trajectory Predictor and Optimiser for ATM Research Applications. In *8th Proc. of the International Conference on Research in Air Transportation*, Castelldefels, Spain, 2018.
- [11] 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(July), 2008.
- [12] Gerald Gurtner, Luis Delgado, and Damir Valput. An agent-based model for air transportation to capture network effects in assessing delay management mechanisms. *Transportation Research Part C: Emerging Technologies*, 133:103358, 2021.
- [13] EUROCONTROL. DDR2 Reference Manual v.1.0.1, 2015.
- [14] DFS. AIRAC 2013 (2020-12-03), 2020.