

Probabilistic 4D Flight Planning in Structured Airspaces through Parallelized Simulation on GPUs

Daniel González Arribas^{*}, Eduardo Andrés-Enderiz[†], Manuel Soler[‡], Aniel Jardines[§], Javier García-Heras[¶]

Department of Bioengineering and Aerospace Engineering
Universidad Carlos III de Madrid

Leganés, Madrid, Spain

Email: ^{*}dangonza@ing.uc3m.es, [†]eandres@ing.uc3m.es, [‡]masolera@ing.uc3m.es, [§]ajardine@ing.uc3m.es, [¶]gcarrete@ing.uc3m.es

Abstract—The Air Traffic Management system is evolving to deal with efficiency, capacity, safety and environmental challenges. Progress along these fronts requires the development of trajectory planning and prediction tools that can deal with a complex and uncertain meteorological and operational context and go beyond the deterministic planning paradigm that underlies the technologies currently in place in ATM. In this work, we introduce a novel flight planning methodology to generate weather-optimal 3D flight plans in structured airspaces. By leveraging general-purpose computing on graphics processing units, we can simulate and evaluate multiple trajectory options under multiple scenarios in parallel, allowing us to provide quick iterations to a stochastic optimization algorithm. Our computational experiments show that our implementation can provide efficient solutions in seconds, as required in practical settings, while allowing for simple integration of future extensions thanks to its simulation-based nature.

Keywords—Aircraft trajectories, trajectory optimization, air traffic control, parallel programming, forecast uncertainty

I. INTRODUCTION

Strong growth in air traffic over the last years is leading to higher congestion and capacity shortfalls in some of the busiest airspaces in the world. As an illustration, during the busy third quarter that comprises the summer years, European airspace has seen 10-year record levels of delay in 2018 and 2019 [1], with 17.9 and 15.9 minutes per flight respectively. As the total amount of movements is expected to keep growing over the next decades [2], airport and airspace congestion will worsen unless ongoing and future mitigation efforts prove successful. In addition, the aviation sector faces increasing pressure to reduce its environmental footprint and to improve cost-efficiency and safety. By improving the efficiency of aircraft operations, the Air Traffic Management (ATM) system can make a critical contribution towards these goals, and multiple R&D initiatives (such as SESAR in Europe or NextGen in the US) have consequently been put in place to address the aforementioned challenges by developing and implementing new technologies, regulations, and operational procedures.

One critical channel through which these innovation initiatives can improve ATM performance in multiple areas is efficient management of uncertainty and predictability. The ATM system is a complex, interconnected socio-technical system where uncertainty is a fundamental feature, originating from multiple sources and propagating to every component of

the system. Successfully preventing, adapting to or recovering from disruption caused by uncertainty are, therefore, critical capabilities of a thriving ATM system.

At the center of the uncertainty diffusion process lies the aircraft trajectory. The execution of any given flight will differ from the planned trajectory and schedule due to uncertainty in the outcomes of airport processes (takeoff time and mass), weather, performance and pilot decisions. Then, as trajectories share a common airspace, they impact each other indirectly through deconfliction by air traffic controllers and flow management regulations. Finally, accumulated uncertainty is spread to arrival aerodromes and other airspace sectors, as well as through reactive operational modifications by airlines.

Thus, while individual trajectory analyses are by no means the only relevant topic in ATM-related uncertainty, trajectory prediction and optimization tools that take uncertainty into account will grow in importance in the next generation ATM systems. A healthy amount of effort is now devoted to study of uncertainty within the context of trajectory prediction (see, for example, [3], [4], [5], [6] for recent works on modeling various uncertainty sources); however, trajectory optimization lags behind prediction in the context of considering uncertainty.

Nonetheless, research interest is growing in this topic: in [7], it is proposed to employ weather uncertainty information, in the form of ensemble forecasts, as an input to the trajectory optimization process. Uncertainty quantification ideas are combined with a trajectory optimization approach in [8] in the context of conflict resolution. In previous works [9], we employed ensemble forecasts to model meteorological uncertainty at the flight planning stage within the context of a future, fully free-route airspace (FRA). However, the proposed methods, based on optimal control and direct collocation, are not applicable to the structured airspaces of today. A FRA is also assumed in [10], where ensemble forecasts are combined with a trajectory distance metric and a clustering algorithm to produce robust flight plans; here, robustness is defined as “optimality within the largest set of scenarios”. In [11], trajectory shapes are modified by a metaheuristic optimization algorithm in order to strategically plan traffic under uncertainty. Again, these works rely on continuous decision spaces. In contrast, the work of [12] deals with uncertain flight plan optimization on structured airspaces by

employing mixed-integer linear programming (MILP) tools; however, the computational cost (measured in minutes) is sizable and the assumptions are restrictive: fuel burn and its associated nonlinearities are disregarded and a single flight level is assumed.

Furthermore, even if one disregards uncertainty entirely, merely modeling these nonlinearities due to four-dimensional weather, nonlinear performance functions and other potential factors (such as convection or airspace congestion) represents a major algorithmic challenge. The traditional approach, widely applied in industrial practice, is often based on simplifying assumptions that reduce the problem to a two-stage process, starting with a shortest-path problem on the airspace graph to determine the route and finishing with the optimization of the vertical profile (altitude and airspeed) [13]. While this approach has been successful and provides good solutions most of the time, it is hard to extend it to incorporate additional features in efficient fashion and to consider the interplay between the vertical and the lateral profiles. Recent works, such as [14] and [15], show great progress in developing effective heuristics to path planning algorithms such as A* in order to tackle the 3D flight planning problem in a deterministic fashion. Nevertheless, it is not clear whether these heuristics can be extended to incorporate the additional features that will soon be demanded by incoming technologies and paradigms.

In summary, there is a clear need for flight planning algorithms in structured airspaces that 1) tackle the vertical and lateral profile in an integrated fashion, 2) consider the nonlinearities associated with performance and 4-dimensional weather and 3) are able to incorporate uncertainty and 4) are flexible enough to incorporate more complex features as extensions.

One potential piece of the puzzle was introduced in [16]. By leveraging the computational power of general-purpose computing on GPUs, it is possible to quickly simulate multiple flight options under multiple scenarios to feed an optimization algorithm. Furthermore, by using a simulation-based approach, incorporating new features is a natural extension. However, the mentioned work dealt again with free-routing airspaces only.

In the current work, we propose a computationally efficient and extensible flight planning solution in structured airspaces. By developing novel modelling concepts, based on the optimization over probability distributions of flight plans, we show that this parallel-simulation-with-optimization approach can be extended to structured airspaces with the same advantages. In addition, we implement a randomized optimization algorithm, inspired by popular techniques in stochastic optimization, that leverages this parallel simulation framework to produce near-optimal flight plans in a timescale of seconds.

The paper is organized as follows. After the introduction (Section I), we overview the mathematical implications of flight planning concepts under uncertainty in Section II. Section III introduces the modelling and simulation strategy, while Section IV delves into the optimization methodology. We show preliminary results from our prototype implementation in Section V and finish up with some conclusions and future

work in Section VI.

II. OPERATIONAL FRAMEWORKS UNDER UNCERTAINTY

In a deterministic operational setting, where it is assumed that the relevant outcomes can be computed a priori for any action sequence, a lot of practical considerations regarding the implementation of these actions (such as when or how frequently can flight plans be computed or modified, or when is relevant information available) do not impact the formulation of the corresponding mathematical optimization problem. However, if one allows for the existence of uncertainty or stochasticity in some of the involved variables, this conclusion no longer holds. In that case, the resulting mathematical optimization problem depends not only on the shape, magnitude or structure of this uncertainty, but also on the operational framework that determines when and which actions can or must be taken in response to unexpected conditions. Making different assumptions about the operational context leads to different kinds of dynamic optimization problems. In the scope of this work, we will employ the following taxonomy:

- In a “naïve” or **deterministic** flight planning framework, uncertainty is deemed to be negligible or, alternatively, it might be assumed that the optimal decision is close even if the outcome is uncertain. Any unexpected development is managed afterwards, possibly by amending the flight plan during execution with another deterministic calculation. This approach is the simplest and often close to optimal, but it may result in inefficient flight plans and heavily degraded ATM performance in certain instances, such as days of intense convective activity featuring numerous reroutes and conflicts.
- The **robust** flight planning concept assumes that the computed flight plan will be followed closely, but the performance in terms of variables such as fuel burn or arrival time will be impacted by uncertainty. Therefore, the aim is to design a flight plan that is not merely optimal for a baseline scenario, proving instead good performance across a number of potential realizations of the uncertainty. Thus, this concept represents a more realistic representation of actual aircraft operations than the deterministic one.
- The **dynamic** flight planning framework allows the flight plan to be continuously modified as the uncertainty is realized; this is the approach taken under models based on Markov Decision Processes (MDPs). By taking advantage of all the information available until the implementation of a control decision, as well as the possibility of taking multiple choices in the future, employing this framework could let a single aircraft attain the best theoretical efficiency. However, current technologies and operational procedures do not support constant flight plan updates because of the associated CNS and controller workload requirements. Therefore, this concept corresponds to a potential future ATM system rather than the present one.
- Finally, a **hybrid** setting combines elements of both the robust and dynamic paradigms, by assuming that

flight plans can be amended while being executed but opportunities to do so are limited or costly. While this approach leads to challenging mathematical modelling, it is the closest to current and near-term operations.

In the current work, we will work within a *robust* framework, as in the works mentioned in Section I; however, we will also introduce some ideas connected to the *hybrid* concept when discussing probabilistically-executed flight plans in Section IV-A.

III. FLIGHT PLAN SIMULATION

We consider the problem of determining the optimal flight plan in a structured airspace under uncertainty. For the purposes of exposition, we restrict ourselves in the current work to consider only the cruise phase of the flight, as it illustrates the main principles of our proposed approach.

A. Modelling

The airspace structure is modeled as a directed acyclic graph $G = (V, E)$, where V represents a set of navigation waypoints that are connected by airway edges $e \in E$. We define as $\text{Out}(n)$ the set of airway segments outgoing from node n . The cruise is assumed to start at the first waypoint after the top of climb; we denote it as the origin node $o \in V$; similarly, we assume that the trajectory will end at the last node before the top of descent, the destination node $d \in V$. We assume that the graph has been preprocessed to remove any node from which there is no path to the destination node. The aircraft is allowed to change its airspeed over the cruise phase within its flight envelope, and it is allowed to perform step climbs and/or descents between discrete flight levels when entering a new airway segment (though the frequency of those changes will be restricted, as described in Section IV-A).

In this setting, a flight plan \mathcal{F} is described by a tuple $(\mathcal{R}, \overline{\text{FL}}, \overline{\text{M}})$, where:

- \mathcal{R} represents the route or lateral path; it is represented by a valid non-cyclic path between the origin and destination nodes, i.e., a sequence of waypoints $\mathcal{R} := (r_0, r_1, \dots, r_{n_r})$ such that $r_0 = o$, $r_{n_r} = d$, $r_k \in V \forall k \in \{1, \dots, n_r\}$, $\{r_k, r_{k+1}\} \in E \forall k \in \{1, \dots, n_r - 1\}$ and $r_i \neq r_j$ when $i \neq j$.
- A vertical profile $\overline{\text{FL}}$, composed by an ordered sequence of tuples of the form (r_k, FL_k) indicating that the aircraft will switch to the flight level FL_k when reaching the waypoint r_k .
- An Mach schedule $\overline{\text{M}} := (\text{M}_0, \dots, \text{M}_{n_r})$ indicating the target Mach number M_k at waypoint r_k .

To obtain the performance associated to a flight plan, we must compute the corresponding trajectory with a dynamical model of the aircraft. As is commonly done in ATM studies, the state of the aircraft is assumed to evolve according to the point-mass model, where the aerodynamic and propulsive performance of the aircraft is given by the BADA model [17], [18]. Neglecting turn dynamics and flight path angle dynamics (which take place at relatively fast timescales for the flight planning problem), the equations of motion are the following:

Equations of motion

State variables: ϕ, λ, h, v, m

Control variables: C_T, γ, χ

Differential equations:

$$\begin{bmatrix} \dot{\phi} \\ \dot{\lambda} \\ \dot{h} \\ \dot{v} \\ \dot{m} \end{bmatrix} = \begin{bmatrix} \frac{v \cos \gamma \cos \chi + w_y}{(R_M(\phi) + h)} \\ \frac{v \cos \gamma \sin \chi + w_x}{(R_N(\phi) + h) \cos \phi} \\ v \sin \gamma \\ \frac{T(C_T) - D(C_L)}{m} - g \sin \gamma \\ -f_c(C_T) \end{bmatrix}$$

$$\text{where } C_L(\gamma) = \frac{2mg \cos \gamma}{\rho v^2 S}$$

Here, the state is composed by the latitude ϕ , the longitude λ , the altitude h , the true airspeed (TAS) v and the mass m , while the control vector is composed by a thrust coefficient C_T , the flight path angle γ and the true heading is χ . In addition, (w_x, w_y) are the components of the wind, R_M and R_N are the Earth's ellipsoid radii of curvature in the meridian and the prime vertical respectively, T and D are the magnitude of the thrust and drag forces, g is the Earth's gravity, f_c is the fuel burn rate and S is the wetted surface of the aircraft. For clarity, the functional dependences on the remaining environmental variables (the air density ρ , the temperature T and the pressure P) have been omitted; the Mach number M can be computed with the usual formula $M = v/a(T)$, with a being the speed of sound on dry air at temperature T .

As the weather variables at a given pressure level (ρ, T, w_x, w_y) are taken from a forecast, there is some degree of uncertainty in their values, due to incomplete knowledge of the state of the atmosphere, model error in physical parametrizations, computational limitations and nonlinear, sometimes chaotic, dynamics. We will characterize this uncertainty by employing Ensemble Prediction System (EPS) forecasts, a standard tool in modern meteorology for dealing with uncertainty [19]. An EPS is composed by N_{EPS} individual forecasted scenarios (with EPS usually ranging from 10 to 50), called "ensemble members", each one run with strategically perturbed initial conditions or model parameters, thus providing a probabilistic estimation of the future state of the atmosphere. In this work, we denote the set of all weather-related variables as a random variable \mathcal{W} that takes discrete values in the set $\{\mathcal{W}_1, \dots, \mathcal{W}_{N_{\text{EPS}}}\}$ with probability $\mathbb{P}(\mathcal{W} = \mathcal{W}_j) = N_{\text{EPS}}^{-1} \forall j$.

Note that the route specified by the flight plan fixes the Mach number (which, for a given temperature, determine the

true airspeed v) as well as the course ψ at any point in the trajectory. Therefore, as the velocity composition on the horizontal plane shows (see Figure 1), any uncertainty in the magnitude and direction of the wind will generate uncertainty in the heading and, most, importantly, the groundspeed, which determines the time at which the aircraft will overfly each waypoint in the route, as well as how much time the aircraft will spend at each leg (thus influencing fuel burn). Uncertainty in temperature, which determines the propulsive and aerodynamic performance of the aircraft as well as the airspeed, will also contribute to uncertainty in fuel burn.

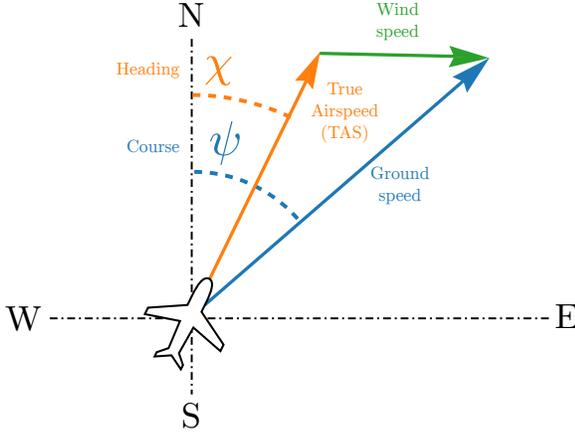


Fig. 1: Wind triangle when $\gamma \approx 0$

Finally, the time at which the aircraft crosses the origin node is modelled as a Gaussian variable $t_0 \sim \mathcal{N}(\bar{t}_0, \sigma_{t_0})$; analogously, the mass of the aircraft at the origin node is $m_0 \sim \mathcal{N}(\bar{m}_0, \sigma_{m_0})$. Both quantities may be correlated, with a correlation coefficient ρ_{mt} ; they are assumed to be independent of the weather realization, though again this assumption can be relaxed. As these distributions will be accessed by the algorithm by sampling, other distributions can be employed by the algorithm.

B. Trajectory Integration

For a given realization of \mathcal{W} , t_0 and m_0 , we can compute in a deterministic fashion the trajectory associated to a flight plan \mathcal{F} and obtain the performance figures of interest; in this work, we will restrict ourselves to the consideration of the final mass m_f and the arrival time t_f . The corresponding calculation, which we will now proceed to explain briefly, is denoted by:

$$[t_f, m_f] = \text{TI}(\mathcal{F}, \mathcal{W}, t_0, m_0) \quad (1)$$

Where TI stands for ‘‘trajectory integration’’ In order to perform this integration, each airway segment in the route \mathcal{R} is discretized into an equispaced grid of geographical points, with a resolution that allows a step climb to be performed in a single interval (though finer grids can be employed by slightly modifying the calculation method). Then, the grid conformed

by the union of the grids of every segment in the route is employed to integrate the following two differential equations:

$$\frac{dt}{ds} = v_G^{-1} \quad (2)$$

$$\frac{dm}{dt} = -f_c(C_{T,req}) \quad (3)$$

where s is the ground distance flown, v_G is the groundspeed (computed from h , v , γ , ψ and (w_x, w_y) through geometrical relationships) and $C_{T,req}$ is the thrust parameter required to match the altitude and speed profile. The integration scheme employed is Heun’s method, which approximates the solution of an ODE system of the form $\dot{y}(t) = f(t, y(t))$ by:

$$\begin{aligned} \tilde{y}_{i+1} &= y_i + \Delta t f(t_i, y_i) \text{ (predictor)} \\ y_{i+1} &= y_i + \frac{\Delta t}{2} [f(t_i, y_i) + f(t_{i+1}, \tilde{y}_{i+1})] \text{ (corrector)} \end{aligned} \quad (4)$$

where the step size Δt is determined by the spacing between grid nodes. While both equations are integrated in lockstep, the time evolution (independent of the mass profile) is integrated first, increasing the accuracy of the mass calculations.

We can now define the expected final mass and time as:

$$[\mathbb{E}[t_f], \mathbb{E}[m_f]] = \mathbb{E}[\text{TI}(\mathcal{F}, \mathcal{W}, t_0, m_0)] \quad (5)$$

We can approximate the expectation operator by an unweighted average between all ensemble members $\{\mathcal{W}_j\}$:

$$[\mathbb{E}[t_f], \mathbb{E}[m_f]] \approx \frac{1}{N_{\text{EPS}}} \sum_j \text{TI}(\mathcal{F}, \mathcal{W}_j, t_0^j, m_0^j) \quad (6)$$

where we sample $t_0^j, m_0^j \sim t_0, m_0$ independently for each member in Monte Carlo-like fashion. As each calculations performs the same (or very similar) operations on different data, this computation can be easily parallelized in order to compute the trajectory in each scenario quickly; then, the sum can be done using parallel reduction techniques. We do so by employing CUDA[20], [21], a framework for general-purpose computing on Graphics Processing Unit (GPGPU) on devices manufactured by NVIDIA. Additionally, this choice allows us to employ the texture units on a GPU to perform fast interpolation of the weather variables.

IV. PROBABILISTIC FLIGHT PLAN GENERATION

We will now proceed to employ the scheme described in Section III-B to generate optimized flight plans. To do so, we must start by defining the objective function to be optimized. Again, while we note that our approach could accommodate a number of extensions, we choose a common objective in ATM:

$$J(\mathcal{F}) = \underbrace{\bar{m}_0 - \mathbb{E}[m_f]}_{\text{expected fuel burn}} + \text{CI} \cdot \underbrace{(\mathbb{E}[t_f] - \bar{t}_0)}_{\text{expected flight time}} \quad (7)$$

where CI denotes the ‘‘cost index’’, i.e., the relative cost of time with respect to fuel, in the view of the user. Our goal is, then, to solve the following optimization problem:

$$\min_{\mathcal{F}} J(\mathcal{F}) \quad (8)$$

As discussed in Section I, optimizing directly on the flight plan decision space (which is composed by both discrete and continuous decision variables that may interact in nonconvex fashion) is challenging and prevents the use of gradient-based and gradient-inspired methods, the main workhorse of applied optimization. Instead, we will borrow a popular technique from some reinforcement learning algorithms [22]; we will replace the problem defined by Equation 8, where optimization is performed on the space of possible flight plans, by the problem defined by Equation 9, where optimization is performed on the space of probability distributions defined over flight plans.

$$\min_{p(\mathcal{F})} \mathbb{E}_{p(\mathcal{F})}[J(\mathcal{F})] \quad (9)$$

If \mathcal{F}^* is the optimal solution of Problem 8, with optimal cost J^* , then a probability distribution that has $\mathbb{P}(\mathcal{F} = \mathcal{F}^*) = 1$ produces the same value of J in Problem 9; it is also clear that $\mathbb{E}_p[J(\mathcal{F})] \geq J^*$ if p is a valid probability distribution. Therefore, both formulations are equivalent, but the second one may be easier to search; we can do so by parameterizing the distribution p with a parameter vector $\theta \in \mathbb{R}^\Theta$ and defining a third variation of the problem:

$$\min_{\theta} \mathbb{E}_{p(\mathcal{F};\theta)}[J(\mathcal{F})] \quad (10)$$

This problem is no longer equivalent in general to the first problem, as it depends on whether the parameterization can represent a distribution where $\mathbb{P}(\mathcal{F} = \mathcal{F}^*) = 1$. If it cannot, then the optimal solution of Equation 10 may be worse than the optimal solution of Equation 9 by a finite amount. We will present our chosen parameterization in Section IV-A, while describing how it can accommodate solutions that are arbitrarily close to the optimal non-probabilistic flight plan.

A. Flight Plan Encoding

We introduce the Probabilistic-execution Flight Plan ($\mathcal{P}\mathcal{F}$) as a distribution over the space of possible flight plans. In a $\mathcal{P}\mathcal{F}$, the aircraft may take any path between the origin and the destination nodes: when arriving at a node, it can then choose any of the outgoing airway segments. The probability of choosing each airway is determined by a vector of parameters $\Phi \in \mathbb{R}^{n_{cr}}$ in a manner that will be described afterwards. In a regular flight plan, the flight level changes and the Mach profile are defined only along the chosen route; instead, in a $\mathcal{P}\mathcal{F}$, the target Mach numbers $\hat{M} \in \mathbb{R}^{|\mathcal{E}|}$ and flight levels $\hat{FL} \in \mathbb{R}^{|\mathcal{E}|}$ (note that $|\cdot|$ denotes the cardinality of a set) will be defined in the entire graph, regardless of whether they lie within the sampled route or not. Thus, we define the parameter vector of a $\mathcal{P}\mathcal{F}$ as:

$$\theta = \begin{bmatrix} \Phi \\ \hat{M} \\ \hat{FL} \end{bmatrix} \quad (11)$$

To sample the route of a $\mathcal{P}\mathcal{F}$, a vector of i.i.d. uniform random variables $\xi_{cr} \sim U(0, 1) \in \mathbb{R}^{n_{cr}}$ is generated. Then, we compute a vector of bits $b \in \mathbb{R}^{n_{cr}}$ with the following rule:

$$b_i = \begin{cases} 1 & \text{if } S(\Phi_i) \leq \xi_i \\ 0 & \text{if } S(\Phi_i) > \xi_i \end{cases} \quad (12)$$

where S is a sigmoid function $S(x) = \frac{1}{2} \left(1 + \frac{x}{\sqrt{1+x^2}} \right)$. Each waypoint with more than one outgoing airway segment is assigned one or more of these bits; the resulting binary number determines the chosen airway. Thus, the parameters Φ influence the likelihood that any airway is chosen. Clearly, as the parameters Φ grow towards ∞ or $-\infty$, the route choice converges to a deterministic lateral path, thus converging towards the corresponding deterministic performance.

To complete the sampled flight path, it is necessary to generate the Mach schedule and the flight level switching profile. We will do so by introducing a probabilistic Mach and flight level switching feature that is not required by the probabilistic transformation, but is instead important for the operational acceptability of the solution (though it can be disabled with certain choices of the corresponding parameters). In first place, we the Mach and flight level (FL) “guidance” profiles as smoothed versions of the profile determined by the FL and Mach values of the segments. In second place, the actual Mach number and flight level switch to the discrete value that is closest¹ to the guidance profile in a probabilistic fashion, with the switching probability being small until a certain distance has been flown since the last change. In this way, we prevent continuous Mach adjustments that may generate undesired wear in the engines and we limit the frequency of the flight level changes, in order to prevent very frequent climbs and descents that would generate extra workload for the air traffic controllers.

Algorithm 1 formalizes the described process for sampling the $\mathcal{P}\mathcal{F}$. We define the following parameters and notation: $||l||$ denotes the length of leg l , ℓ_M and ℓ_{FL} denote respectively the Mach and FL guidance smoothing factors, D_M and D_{FL} denote the average minimum distance between switches, and σ_M and σ_{FL} denote the execution variability of these changes. \hat{M}_l and \hat{FL}_l denote the target Mach and flight level for the segment l (from the parameter vector θ).

The rationale for the probabilistic formulation of the route choice is the one exposed at the beginning of Section IV, i.e., for convenience of optimization; however, the probabilistic formulation of the profile adjustment can also be justified by the existence of uncertainty in the time at which the aircraft obtains ATC clearance to move to an adjacent flight level.

B. Optimization

As the $\mathcal{P}\mathcal{F}$ corresponding to a vector of parameters θ is now stochastic, evaluating $\mathbb{E}_\theta[J(\mathcal{F})] := \mathbb{E}_{p(\mathcal{F};\theta)}[J(\mathcal{F})]$ requires sampling multiple flight plans and executing them. However, we can take advantage of the fact that we are already sampling on the environmental uncertainty to sample simultaneously on the space of exogenous sources of uncertainty and on the space of potential flight plans associated to θ :

¹For the flight level changes, available flight levels are separated by jumps of 2000 feet; for the Mach profile, available settings are multiples of 0.01

Algorithm 1: Sampling a flight plan from a \mathcal{PF}

```
1 Initialize  $\mathcal{R} \leftarrow (o)$ ,  $\bar{\mathbf{M}} \leftarrow (\mathbf{M}_0)$ ,  $\bar{\mathbf{FL}} \leftarrow (\mathbf{FL}_0)$ 
2 Initialize  $\mathbf{WP} \leftarrow o$ ,  $\mathbf{M} \leftarrow \bar{\mathbf{M}} \leftarrow \mathbf{M}_0$ ,  $\mathbf{FL} \leftarrow \bar{\mathbf{FL}} \leftarrow \mathbf{FL}_0$ 
3 Initialize  $L_{\mathbf{FL}} \leftarrow 0$  # Distance since the last FL change
4 Initialize  $L_{\mathbf{M}} \leftarrow 0$  # Distance since the last Mach
  change
5 Sample  $\xi_{cr} \sim U(0, 1) \in \mathbb{R}^{n_{cr}}$ 
6 Sample  $\xi_{\mathbf{FL}} \sim U(0, 1) \in \mathbb{R}^{|\mathcal{E}|}$ 
7 Sample  $\xi_{\mathbf{M}} \sim U(0, 1) \in \mathbb{R}^{|\mathcal{E}|}$ 
8 while  $\mathbf{WP} \neq d$  do
9    $\mathbf{WP}_+ \leftarrow \text{next\_waypoint}(\mathbf{WP}, \xi_{cr}, \Phi)$ 
10   $l \leftarrow \{\mathbf{WP}, \mathbf{WP}_+\}$ 
11   $\delta_{\mathbf{M}} \leftarrow \|l\|/\ell_{\mathbf{M}}$ 
12   $\delta_{\mathbf{FL}} \leftarrow \|l\|/\ell_{\mathbf{FL}}$ 
13   $\tilde{\mathbf{M}} \leftarrow (1 - \delta_{\mathbf{M}})\bar{\mathbf{M}} + \delta_{\mathbf{M}}\hat{\mathbf{M}}_l$ 
14   $\tilde{\mathbf{FL}} \leftarrow (1 - \delta_{\mathbf{FL}})\bar{\mathbf{FL}} + \delta_{\mathbf{FL}}\hat{\mathbf{FL}}_l$ 
15  if  $S((L_{\mathbf{FL}} - D_{\mathbf{FL}})/\sigma_{\mathbf{FL}}) \geq \xi_{\mathbf{FL}}^l$  then
16    if  $\tilde{\mathbf{FL}} \geq \mathbf{FL} + 10$  then
17       $\mathbf{FL} \leftarrow \mathbf{FL} + 20$ 
18       $L_{\mathbf{FL}} \leftarrow 0$ 
19      Append  $(\mathbf{WP}_+, \mathbf{FL})$  to  $\bar{\mathbf{FL}}$ 
20    else if  $\tilde{\mathbf{FL}} \leq \mathbf{FL} - 10$  then
21       $\mathbf{FL} \leftarrow \mathbf{FL} - 20$ 
22       $L_{\mathbf{FL}} \leftarrow 0$ 
23      Append  $(\mathbf{WP}_+, \mathbf{FL})$  to  $\bar{\mathbf{FL}}$ 
24  if  $S((L_{\mathbf{M}} - D_{\mathbf{M}})/\sigma_{\mathbf{M}}) \geq \xi_{\mathbf{M}}^l$  then
25     $\mathbf{M} \leftarrow 0.01 \cdot \text{round2int}(100 \cdot \tilde{\mathbf{M}})$ 
26     $L_{\mathbf{M}} \leftarrow 0$ 
27   $L_{\mathbf{FL}} \leftarrow L_{\mathbf{FL}} + \|l\|$ 
28   $L_{\mathbf{M}} \leftarrow L_{\mathbf{M}} + \|l\|$ 
29  Append  $\mathbf{M}$  to  $\bar{\mathbf{M}}$ 
30  Append  $\mathbf{WP}_+$  to  $\mathcal{R}$ 
31   $\mathbf{WP} \leftarrow \mathbf{WP}_+$ 
```

$$[\hat{m}_f, \hat{t}_f] = \frac{1}{N_{\text{EPS}}} \sum_j \text{TI}(\mathcal{F}_j, \mathcal{W}_j, t_0^j, m_0^j) \quad (13)$$

$$\mathbb{E}_{\theta}[J(\mathcal{F})] \approx \hat{J} := \bar{m}_0 - \hat{m}_f + \text{CI} \cdot (t_f - \bar{t}_0) \quad (14)$$

where \mathcal{F}_j is sampled from θ independently. In this way, we can reduce the problem to the stochastic optimization problem:

$$\min_{\theta} \mathbb{E}[\hat{J}] \quad (15)$$

We employ Algorithm 2 to optimize this objective function. The described method is based on the V1 version of the Augmented Random Search (ARS) algorithm from [23]. From a starting point θ_0 , this scheme produces progressively better iterates by generating n random search directions $\omega \in \mathbb{R}^{\Theta}$ and estimating $\hat{J}^+ := \hat{J}(\theta + \mathbf{S}\omega)$ and $\hat{J}^- := \hat{J}(\theta - \mathbf{S}\omega)$, where $\mathbf{S} \in \mathbb{R}^{\Theta \times \Theta}$ is a diagonal scaling matrix. Then, the decision variables θ are advanced along every search direction in proportion to the differential $\hat{J}^+ - \hat{J}^-$. Thus, ARS can

be thought of as a kind of estimated gradient descent, where the estimated gradient is computed through centered finite differences on a projection along random directions.

The algorithm allows us to take full advantage of the high degree of parallelism offered by modern GPUs, as we are now sampling along a new dimension; i.e. we are evaluating $\text{TI}(\cdot)$ $2n \times N_{\text{EPS}}$ trajectories simultaneously each iteration. In order to increase convergence speed, we add one of the ‘‘momentum with lookahead’’ heuristics studied in [24] to the algorithm; we denote the update velocity as $\nu \in \mathbb{R}^{\Theta}$. Note that, since we are not solving a problem of the form studied in [23], there are some minor differences in the definition of the decision vector, which in our case is closer to what these authors name ‘‘Basic Random Search’’.

Algorithm 2: Augmented Random Search with Nesterov updates

```
1 Hyperparameters: step-size  $\alpha$ , Nesterov momentum
  factor  $\beta$ , number of directions sampled per iteration
   $n$ , diagonal scaling matrix for the exploration noise  $\mathbf{S}$ 
2 Initialize:  $\nu \leftarrow 0$ ,  $\theta \leftarrow \theta_0$ 
3 while stopping condition not met do
4   Sample  $\omega_i \sim \mathcal{N}(0, 1) \in \mathbb{R}^{\Theta}$ ,  $i \in \{1, \dots, n\}$ 
5    $r_i^+ = \hat{J}(\theta + \beta\nu + \mathbf{S}\omega_i)$ ,  $i \in \{1, \dots, n\}$ 
6    $r_i^- = \hat{J}(\theta + \beta\nu - \mathbf{S}\omega_i)$ ,  $i \in \{1, \dots, n\}$ 
7    $\bar{r} \leftarrow \frac{1}{2n} \sum_i (r_i^+ + r_i^-)$ 
8    $\sigma_R \leftarrow \sqrt{\frac{1}{2n} \sum_i ((r_i^+ - \bar{r})^2 + (r_i^- - \bar{r})^2)}$ 
9    $\nu \leftarrow \beta\nu - \frac{\alpha}{n\sigma_R} \sum_i (r_i^+ - r_i^-)\omega_i$ 
10   $\theta \leftarrow \text{clamp}(\theta + \nu)$ 
```

Here, $\text{clamp}(\cdot)$ is a function that projects the updated decision variables inside the flight envelope, to prevent the proposal of infeasible actions to the simulation engine.

V. RESULTS

We study the flight of a narrow-body twinjet taking place the 7th of June, 2018, from a waypoint in the neighbourhood of Madrid to a waypoint in the neighbourhood of Berlin. For the route graph G , we take the full airspace graph from that day and generate the subgraph induced by the paths with length of at most 102.5% of the length of the minimum ground distance path; the integration resolution is set to 10 nmi. The weather forecast is taken from the ECMWF EPS[25].

We launch the simulations on a consumer-grade NVIDIA 1060 GTX graphics card, a device featuring 1280 CUDA cores at a clock speed of 1.6 GHz; we perform the calculations on single-precision. Table I describes the setup of the main computational hyperparameters; only basic tuning has been performed in them, suggesting that systematic tuning or automatic scheduling at different optimization epochs may improve the computational performance.

Parameter	Value	Parameter	Value
ℓ_M	200 nmi	ℓ_{FL}	100 nmi
D_M	250 nmi	D_{FL}	300 nmi
σ_M	10 nmi	σ_{FL}	10 nmi
α	0.02	β	3/4
n	25	S (for Ψ entries)	1
N_{EPS}	32	S (for M entries)	0.15
		S (for FL entries)	15

TABLE I: Parameter values

Figure 2 illustrates the performance profile of the optimization process for $CI = 0.5$ kg/s, i.e., the comparison of the performance of the algorithm for 100 different runs. The two inner bands represent the inner 66 and 90 percentiles, while the lighter outer band represents the full range of the profiles. It can be observed that the optimization process converges consistently to the optimum performance, and it does so quickly, with almost all of the runs converging close to the optimum in around 600 iterations. In our machine, that takes around 2 seconds, which is compatible with desired computational speeds for operational flight planning tools.

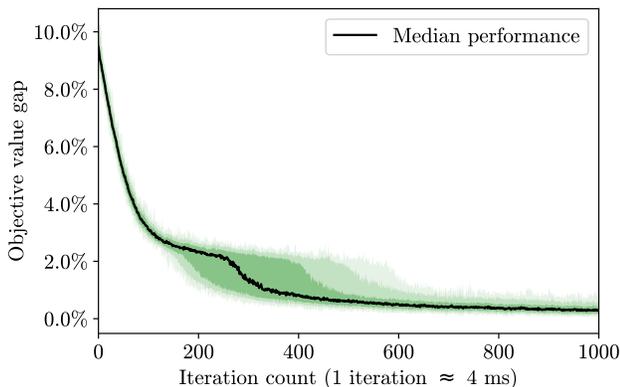


Fig. 2: Performance profiles.

The obtained lateral route is represented in Figure 3, along with the route network considered and the shortest path in the graph. It illustrates the principle that the shortest path in geodesic terms is not necessarily the optimal path in the presence of weather, as winds and temperature play a significant role in route efficiency.

We now proceed to solve the problem for different values of the Cost Index parameter. The route is the same in all of the three cases, but the vertical profile (Figure 4) shows some differences. Looking at the Mach profile, the solutions exhibit expected behaviour: the optimal speed profile increases with higher CI settings, as the optimizer chooses to arrive faster at the cost of burning more fuel. The true airspeed shifts along the route due not only to Mach changes, but also to temperature variations with location and altitude. We can observe that the optimal flight level is higher for the low CI case. Finally, Figure 5 compares the groundspeed profiles (as well as their uncertainty) between different values of the cost index, as well as the additional cumulative fuel burn and flight time with

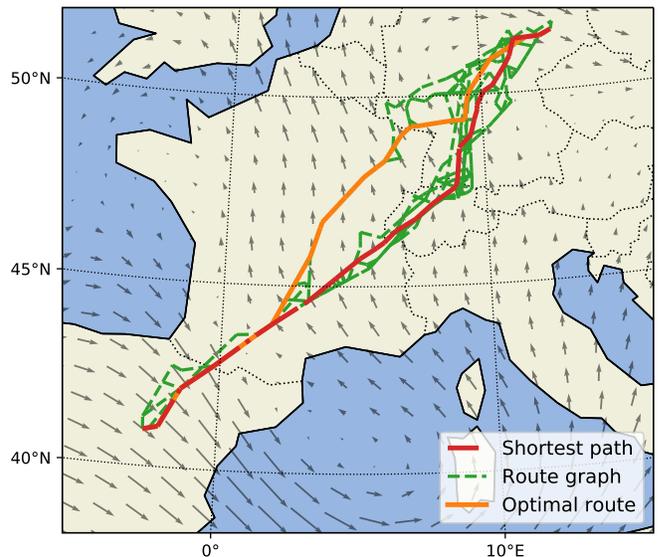


Fig. 3: Lateral route and FL 300 wind forecast for t_0

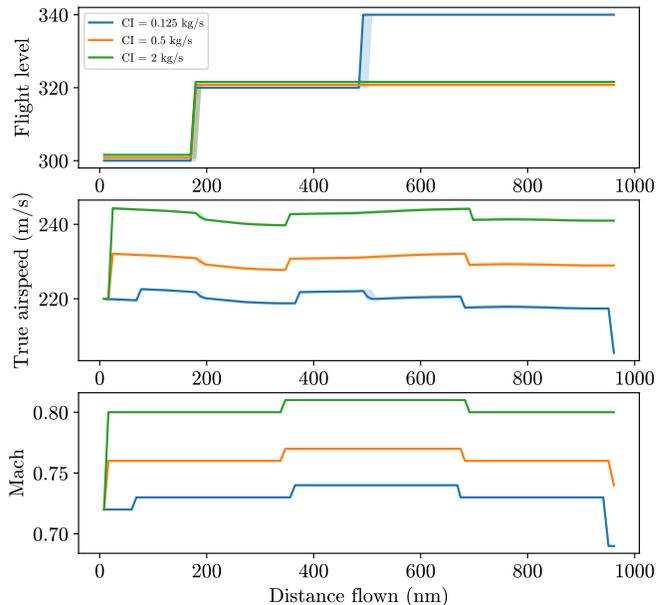


Fig. 4: Vertical profiles.

respect to the minimum-fuel case. We can verify that, indeed, the higher CI solutions end up burning more fuel in order to shave a few minutes off the total flight time.

VI. CONCLUSIONS

In this work, we have introduced a flight planning architecture able to handle 4D flight planning problem under uncertainty in an extensible and efficient manner. We have done so by introducing novel modelling concepts that allow us to harness the highly parallel computational power available in modern GPUs in order to compute trajectories under multiple

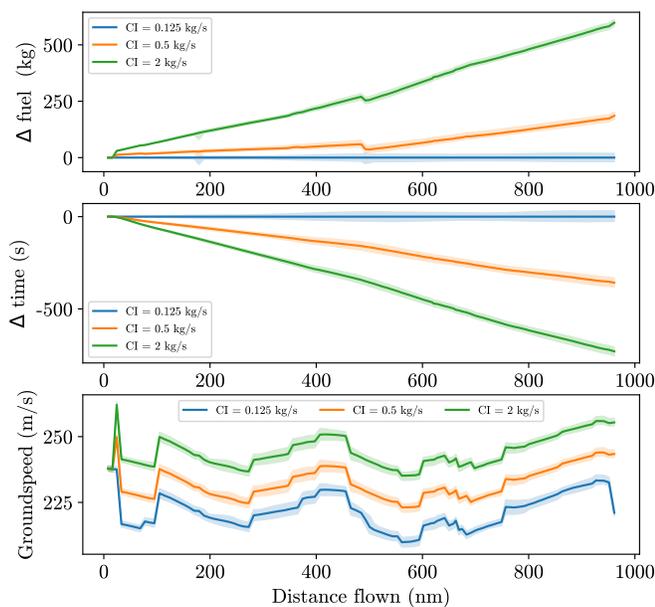


Fig. 5: Groundspeed, mass profiles. Shaded regions depict the uncertainty range.

scenarios and multiple flight plan options simultaneously. By exploring the decision space in a probabilistic fashion, we can make use of effective stochastic optimization algorithms to quickly obtain a solution, and thus meet the computational demands of the future ATM toolchains.

This work opens up some potential avenues for future work that we intend to explore. The first path involves tuning and improving the algorithm to enhance its convergence properties, as well as designing test campaigns to characterize the performance of the algorithm and ability to find the best optima under multiple conditions and flight types.

The second one is to leverage the simulation-based nature of the proposed solution to incorporate additional factors such as climb and descent profiles, overflight costs, convective cell developments or airspace congestion, thus proving its ability to deal with the complex problems for which it was introduced.

REFERENCES

- [1] "All-causes delay and cancellations to air transport europe for q3 2019," Tech. Rep., dec 2019.
- [2] Eurocontrol, "European aviation in 2040 - challenges of growth," Eurocontrol, 2018. [Online]. Available: <https://www.eurocontrol.int/articles/challenges-growth>
- [3] E. Casado, M. L. Civita, M. Vilaplana, and E. W. McGookin, "Quantification of aircraft trajectory prediction uncertainty using polynomial chaos expansions," in *2017 IEEE/AIAA 36th Digital Avionics Systems Conference (DASC)*, 2017, pp. 1–11.
- [4] G. Zhu, C. Matthews, P. Wei, M. Lorch, and S. Chakravarty, "En route flight time prediction under convective weather events," in *2018 Aviation Technology, Integration, and Operations Conference*. American Institute of Aeronautics and Astronautics, jun 2018. [Online]. Available: <https://doi.org/10.2514/6.2018-3670>
- [5] R. Vazquez and D. Rivas, "Propagation of initial mass uncertainty in aircraft cruise flight," *Journal of Guidance, Control, and Dynamics*, vol. 36, no. 2, pp. 415–429, 2013. [Online]. Available: <https://doi.org/10.2514/1.57675>
- [6] J. Cheung, J.-L. Brenguier, J. Heijstek, A. Marsman, and H. Wells, "Sensitivity of flight durations to uncertainties in numerical weather prediction," in *SIDs 2014 - Proceedings of the SESAR Innovation Days*, 2014.
- [7] J. Cheung, A. Hally, J. Heijstek, A. Marsman, and J.-L. Brenguier, "Recommendations on trajectory selection in flight planning based on weather uncertainty," in *SIDs 2015 - Proceedings of the SESAR Innovation Days*, 2015.
- [8] Y. Matsuno, T. Tsuchiya, J. Wei, I. Hwang, and N. Matayoshi, "Stochastic optimal control for aircraft conflict resolution under wind uncertainty," *Aerosp. Sci. Technol.*, vol. 43, pp. 77–88, Jun. 2015. [Online]. Available: <https://doi.org/10.1016/j.ast.2015.02.018>
- [9] D. González-Arribas, M. Soler, and M. Sanjurjo-Rivo, "Robust aircraft trajectory planning under wind uncertainty using optimal control," *Journal of Guidance, Control, and Dynamics*, vol. 41, no. 3, pp. 673–688, mar 2018. [Online]. Available: <https://doi.org/10.2514/1.g002928>
- [10] K. Legrand, S. Puechmorel, D. Delahaye, and Y. Zhu, "Robust aircraft optimal trajectory in the presence of wind," *IEEE Aerospace and Electronic Systems Magazine*, vol. 33, no. 11, pp. 30–38, 2018.
- [11] W. Dai, J. Zhang, D. Delahaye, and X. Sun, "A heuristic algorithm for aircraft 4d trajectory optimization based on bezier curve," 2019.
- [12] A. Franco, D. Rivas, and A. Valenzuela, "Optimal aircraft path planning in a structured airspace using ensemble weather forecasts," in *SIDs 2018 - Proceedings of the Eighth SESAR Innovation Days*. SESAR, dec 2018.
- [13] S. S. Altus, "Flight planning – the forgotten field in airline operations," Jeppesen, Tech. Rep.
- [14] M. Blanco, R. Borndörfer, N.-D. Hoang, A. Kaier, A. Schienle, T. Schlechte, and S. Schlobach, "Solving time dependent shortest path problems on airway networks using super-optimal wind," in *16th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2016)*, ser. OpenAccess Series in Informatics (OASIS), M. Goerigk and R. Werneck, Eds., vol. 54. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2016, pp. 12:1–12:15. [Online]. Available: <http://drops.dagstuhl.de/opus/volltexte/2016/6536>
- [15] A. N. Knudsen, M. Chiarandini, and K. S. Larsen, "Heuristic variants of a* search for 3d flight planning," in *Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, W.-J. van Hoeve, Ed. Cham: Springer International Publishing, 2018, pp. 361–376.
- [16] D. González-Arribas, M. Sanjurjo-Rivo, and M. Soler, "Multiobjective optimisation of aircraft trajectories under wind uncertainty using GPU parallelism and genetic algorithms," in *Computational Methods in Applied Sciences*. Springer International Publishing, Sep. 2018, pp. 453–466. [Online]. Available: https://doi.org/10.1007/978-3-319-89890-2_29
- [17] A. Nuic, "User manual for the base of aircraft data (BADA) rev 3.11," Eurocontrol Experimental Centre, Tech. Rep. 13/04/16-01.
- [18] E. Gallo, F. Navarro, A. Nuic, and M. Iagaru, "Advanced aircraft performance modeling for ATM: Bada 4.0 results," in *2006 IEEE/AIAA 25TH Digital Avionics Systems Conference*. IEEE, Oct. 2006, pp. 1–12. [Online]. Available: <https://doi.org/10.1109/dasc.2006.313660>
- [19] P. Bauer, A. Thorpe, and G. Brunet, "The quiet revolution of numerical weather prediction," *Nature*, vol. 525, no. 7567, pp. 47–55, Sep. 2015. [Online]. Available: <https://doi.org/10.1038/nature14956>
- [20] "Cuda programming guide," NVIDIA Corporation, 2010.
- [21] A. Klöckner, N. Pinto, Y. Lee, B. Catanzaro, P. Ivanov, and A. Fasih, "Pycuda and pyopencl: A scripting-based approach to gpu run-time code generation," *Parallel Computing*, vol. 38, no. 3, pp. 157–174, 2012.
- [22] B. Recht, "A tour of reinforcement learning: The view from continuous control," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 2, no. 1, Dec. 2018. [Online]. Available: <https://doi.org/10.1146/annurev-control-053018-023825>
- [23] H. Mania, A. Guy, and B. Recht, "Simple random search of static linear policies is competitive for reinforcement learning," in *Advances in Neural Information Processing Systems*, 2018, pp. 1800–1809.
- [24] Y. Nesterov and V. Spokoiny, "Random gradient-free minimization of convex functions," *Foundations of Computational Mathematics*, vol. 17, no. 2, pp. 527–566, nov 2015. [Online]. Available: <https://doi.org/10.1007/s10208-015-9296-2>
- [25] R. Buizza, "The TIGGE global, medium-range ensembles," ECMWF, Tech. Rep., nov 2014.