

Robust Aircraft Trajectory Optimization under Meteorological Uncertainty

Daniel González Arribas

in partial fulfillment of the requirements for the degree of Doctor in
Programa Interuniversitario en Mecánica de Fluidos

Universidad Carlos III de Madrid

Advisor 1:
Manuel Soler Arnedo

Advisor 2:
Manuel Sanjurjo Rivo

Tutor:
Manuel Soler Arnedo

Leganés, June 2019

Some rights reserved. This thesis is distributed under a
Creative Commons Attribution-ShareAlike 4.0 License. To
view a copy of this license, visit

<https://creativecommons.org/licenses/by-sa/4.0/>

Acknowledgements

These works have been financially supported by Universidad Carlos III de Madrid through a PIF scholarship; by Eurocontrol, through the HALA! Research Network grant 10-220210-C2; by the Spanish Ministry of Economy and Competitiveness (MINECO)’s R&D program, through the OptMet project (TRA2014-58413-C2-2-R); and by the European Commission’s SESAR Horizon 2020 program, through the TBO-Met project (grant number 699294).

Many people have knowingly or unknowingly or unwillingly done something that has helped me carry out the research described in this thesis. Inevitably, these lines can only acknowledge some of them. To the rest of them, I can only offer you an apology and some words: “whoops I can’t believe I forgot about you”, or “see you in court”, or “at least your absence in these pages will make everyone else feel more special by some weird supervillainy logic so thank you for your selfless sacrifice”. Besides, there is not *that* much blame to share.

In first place, I would like to express my most sincere gratitude towards my advisors, Prof. Manuel Soler and Prof. Manuel Sanjurjo. Their guidance, commitment, encouragement and understanding have been essential for the completion of the thesis. From the outset, Manuel Soler’s vision of the thesis project provided a clear roadmap that allowed us to focus our efforts on the core problems in a fruitful manner. During these years, I always felt that my work, views and proposals were valued and commended, and enhanced by their feedback. As advisors and colleagues, their contributions to this thesis are numerous. Plus, they are genuinely great guys.

Next, I want to thank Prof. Javier García-Heras, with whom I have worked closely the last years. He has provided very valuable feedback, and he has put a lot of hard work into several research projects that we were involved in. As a true team player, he has brought a lot to the research group, and I definitely owe him more than one.

I also offer my special thanks to Prof. Daniel Delahaye and the people at ENAC for their warm welcome to Toulouse, where I was able to discuss insightful ideas on ATM algorithms as I was starting my thesis. I would also like to express my gratitude towards Prof. Maryam Kamgarpour, who hosted my stay at ETH; in Zürich, I was fortunate to spend a few months among leading researchers on control theory and applications. In our subsequent collaborations, Prof. Kamgarpour has always offered excellent input and advice.

I would like to thank Dr. Javier López-Leonés and Dr. Enrique Casado, who supervised my stay with Dr. Miguel Vilaplana’s group at Boeing Research and Technology Europe, as well as Pablo, Luis, and Phillip. My time there was valuable and enjoyable,

and it improved my knowledge and thinking on trajectory planning and execution.

The implementation of the algorithms described in this thesis has been largely performed with the employment of open-source software and libraries. While the amount of people who have contributed to all these software packages is too large to acknowledge all of them, I would like to single out Joel Andersson and Joris Gillis for `CasADi`, Jeffrey S. Whitaker for `pygrib`, Andreas Wächter for `IPOPT`, and all other contributors on these projects and other scientific and engineering libraries. Thanks also to the very kind people at online Q & A sites for making computer pain more bearable.

Sin duda, estos años no hubieran sido lo mismo sin la inmejorable compañía e incondicional ayuda y apoyo de mis compañeros de despacho y departamento. A David, Sara, Dani y Mick, cuyos cacharros espero ver brillando muy muy tenuamente en las esferas celestas, gracias por todos los grandes momentos; a mis compañeros de exilio en el Inframundo, Álex, Bin y Marco (ahora Señor del ídem); a Carlos, mi reincidente vecino de despacho; a Manu, por los ratos en los que compartimos nuestro irracionalmente grande entusiasmo sobre pequeños pero absolutamente críticos detalles sobre gráficas o programación. A Luca y Rocco: I, for one, welcome our new Italian overlords. A Gonzalo, quien espero que algún día me perdone por dedicarle la referencia a Los Simpson a otras personas, y a Güemes, que ha pasado a una vida mejor (remunerada). A Toni, por su influyente investigación sobre los límites del humor gracias a la cual me puedo librar del honor de soltar los peores chistes jamás escuchados en el edificio Juan Benet. A Daniele, Massoud, Valentin, y todos los miembros pasados de nuestro grupo de investigación, y a Eduardo, Nacho y los miembros futuros. A todos los demás estudiosos del Plasma, de quienes he aprendido mucha terminología sobre propulsión eléctrica cuyo significado no sé si quiero descubrir algún día.

Dado que aguantarme es demasiado trabajo para unas pocas personas, mucha gente ha contribuido durante estos años, y es justo agradecérselo. A Álvaro, Aarón, Carlos y Sara, gracias por todas las tardes compartidas. A Marco, Carmen, Javi, Isa, Juli, Celia, Pablo, Arturo, Vir, Elena y Thomas: la mayoría ya sabíais en la carrera lo que había, no se admiten devoluciones. A Miriam, María, Pablo, Ana, Laura, Rebeca y Alberto, gracias por los maravillosos saraos. A Juan, David, Juampe, Maritxu, María, Mario, Miguel, Xiana y Rosana, por tantas conversaciones sobre lo que mola y lo que no. A Fer y todos los compañeros de la raqueta, gracias por los numerosos puntazos.

Finalmente, quiero agradecer a mi familia su aprecio y comprensión durante estos años y, en especial, a Silvia y Laura, a Charco, Trasto, Truco y Coco, y a quienes nunca podré agradecer lo suficiente todo su apoyo y cariño: gracias, mamá y papá.

Published and Submitted Content

The following publications are partially included in the thesis.

Journal papers

1. 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, pp. 673–688, mar 2018 (<https://doi.org/10.2514/1.G002928>). Material from this publication is included mainly in Chapter 4 and Chapter 5 and to a smaller extent in Chapter 1, Chapter 2 and Chapter 3. Author roles: Daniel González Arribas (DGA) carried out the development and implementation of the method and the analysis of the results; the paper was written by DGA with feedback from Manuel Soler (MS) and Manuel Sanjurjo-Rivo (MSR).
2. D. González-Arribas, M. Soler, M. Sanjurjo-Rivo, M. Kamgarpour, and J. Simarro, “Robust aircraft trajectory planning under uncertain convective environments with optimal control and rapidly developing thunderstorms,” *Aerospace Science and Technology*, vol. 89, pp. 445–459, jun 2019 (<https://doi.org/10.1016/j.ast.2019.03.051>). Material from this publication is included mainly in Chapter 7 and Chapter 8 and to a smaller extent in Chapter 1, Chapter 2 and Chapter 3. Author roles: DGA carried out the development and implementation of the algorithm and the analysis of the results; the data was provided by JS and the paper was written by DGA with input and feedback from JS, MS, MSR and Maryam Kamgarpour (MK).

Book chapters

1. D. González Arribas, M. Soler, M. Sanjurjo Rivo, J. García-Heras Carretero, D. Sacher, U. Gelhardt, J. Lang, T. Hauf, and J. Simarro, “Robust optimal trajectory planning under uncertain winds and convective risk,” in *Air Traffic Management and Systems III*, Springer Nature Singapore Pte Ltd., 2019 (forthcoming). Material from this publication is included mainly in Chapter 6 and to a smaller extent in Chapter 1 and Chapter 3. Author roles: DGA carried out the development and implementation of the optimization algorithm, the simulations and the analysis of the results; the convective indicator was developed by Daniel Sacher (DS), Ulrike Gelhardt (UG), and Juergen Lang (JL), and the chapter was written by DGA with feedback from the remaining coauthors.

Conference presentations

1. D. González Arribas, M. Soler, M. Sanjurjo Rivo, J. García-Heras Carretero, D. Sacher, U. Gelhardt, J. Lang, T. Hauf, and J. Simarro, “Robust Optimal Trajectory Planning under Uncertain Winds and Convective Risk,” in *EIWAC 2017 - The 5th ENRI International Workshop on ATM/CNS*, oct 2017 (https://e-archivo.uc3m.es/bitstream/handle/10016/26962/robust_EIWAC_2017.pdf). This publication serves as the basis of the homonymous book chapter; thus, material from this publication is included in Chapter 6 as well as in Chapter 1 and Chapter 3. Author roles: DGA carried out the development of the optimization algorithm and some of the analysis of the results; the convective indicator was developed by DS, UG, JL; MS contributed to the simulations and plots and the paper was written by DGA and MS with feedback from the remaining coauthors.

The material from these sources included in this thesis is not singled out with typographic means and references.

Other Research Merits

The following publications describe additional research performed during the thesis that is related to various degrees. Nevertheless, their content is not included in this thesis.

Journal papers

1. V. Courchelle, M. Soler, D. González-Arribas, and D. Delahaye, “A simulated annealing approach to 3d strategic aircraft deconfliction based on en-route speed changes under wind and temperature uncertainties,” *Transportation Research Part C: Emerging Technologies*, vol. 103, pp. 194–210, jun 2019
2. D. Hentzen, M. Kamgarpour, M. Soler, and D. González-Arribas, “On maximizing safety in stochastic aircraft trajectory planning with uncertain thunderstorm development,” *Aerospace Science and Technology*, vol. 79, pp. 543–553, Aug. 2018
3. D. González-Arribas, M. Soler, J. López-Leonés, E. Casado, and M. Sanjurjo-Rivo, “Automated optimal flight planning based on the aircraft intent description language,” *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, vol. 233, pp. 928–948, jan 2018

Book chapters

1. 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*, pp. 453–466, Springer International Publishing, Sept. 2018

Conference presentations

1. V. Courchelle, D. Delahaye, D. González-Arribas, and M. Soler, “Simulated annealing for strategic traffic deconfliction by subliminal speed control under wind uncertainties,” in *SIDs 2017 - Proceedings of the Seventh SESAR Innovation Days*, SESAR, nov 2017
2. D. González-Arribas, D. Hentzen, M. Sanjurjo-Rivo, M. Soler, and M. Kamgarpour, “Optimal aircraft trajectory planning in the presence of stochastic convective weather cells,” in *17th AIAA Aviation Technology, Integration, and Operations Conference*, p. 3431, American Institute of Aeronautics and Astronautics, June 2017

3. P. Machuca, D. González-Arribas, D. Morante-González, M. Sanjurjo-Rivo, and M. Soler, “Robust optimization of descent trajectories on irregular-shaped bodies in the presence of uncertainty,” in *Proceedings of the AAS/AIAA Astrodynamics Specialist Conference*, SESAR, aug 2017
4. D. González Arribas, M. Soler, and M. Sanjurjo Rivo, “Wind-based robust trajectory optimization using meteorological ensemble probabilistic forecasts,” in *SIDs 2016 - Proceedings of the Sixth SESAR Innovation Days*, SESAR, dec 2016
5. D. Gonzalez Arribas, M. Soler, and M. S. Rivo, “Wind-optimal cruise trajectories using pseudospectral methods and ensemble probabilistic forecasts,” in *Proceedings of the 5th International Conference on Application and Theory of Automation in Command and Control Systems - ATACCS '15*, pp. 160–167, ACM, ACM Press, 2015
6. D. Gonzalez-Arribas, M. Sanjurjo-Rivo, and M. Soler, “Optimization of path-constrained systems using pseudospectral methods applied to aircraft trajectory planning,” *IFAC-PapersOnLine*, vol. 48, no. 9, pp. 192–197, 2015

Awards

- SESAR Young Scientist Award 2017, second place.

Abstract

The Air Traffic Management (ATM) system in the busiest airspaces in the world is currently being overhauled to deal with multiple capacity, socioeconomic, and environmental challenges. One major pillar of this process is the shift towards a concept of operations centered on aircraft trajectories (called Trajectory-Based Operations or TBO in Europe) instead of rigid airspace structures. However, its successful implementation (and, thus, the realization of the associated improvements in ATM performance) rests on appropriate understanding and management of uncertainty. Due to its complex socio-technical structure, the design and operations of the ATM system are heavily impacted by uncertainty, proceeding from multiple sources and propagating through the interconnections between its subsystems.

One major source of ATM uncertainty is weather. Due to its nonlinear and chaotic nature, a number of meteorological phenomena of interest cannot be forecasted with complete accuracy at arbitrary lead times, which leads to uncertainty or disruption in individual air and ground operations that propagates to all ATM processes. Therefore, in order to achieve the goals of SESAR and similar programs, it is necessary to deal with meteorological uncertainty at multiple scales, from the trajectory prediction and planning processes to flow and traffic management operations.

This thesis addresses the problem of single-aircraft flight planning considering two important sources of meteorological uncertainty: wind prediction error and convective activity. As the actual wind field deviates from its forecast, the actual trajectory will diverge in time from the planned trajectory, generating uncertainty in arrival times, sector entry and exit times, and fuel burn. Convective activity also impacts trajectory predictability, as it leads pilots to deviate from their planned route, creating challenging situations for controllers. In this work, we aim to develop algorithms and methods for aircraft trajectory optimization that are able to integrate information about the uncertainty in these meteorological phenomena into the flight planning process at both pre-tactical (before departure) and tactical horizons (while the aircraft is airborne), in order to generate more efficient and predictable trajectories.

To that end, we frame flight planning as an optimal control problem, modeling the motion of the aircraft with a point-mass model and the BADA performance model. Optimal control methods represent a flexible and general approach that has a long history of success in the aerospace field. As a numerical scheme, we use direct methods, which can deal with nonlinear systems of moderate and high-dimensional state spaces in a computationally manageable way. Nevertheless, while this framework is well-developed in the context of deterministic problems, the techniques for the solution of practical

optimal control problems under uncertainty are not as mature, and the methods proposed in the literature are not applicable to the flight planning problem as it is now understood.

The first contribution of this thesis addresses this challenge by introducing a framework for the solution of general nonlinear optimal control problems under parametric uncertainty. It is based on an ensemble trajectory scheme, where the trajectories of the system under multiple scenarios are considered simultaneously within the same dynamical system and the uncertain optimal control problem is turned into a large conventional optimal control problem that can be then solved by standard, well-studied direct methods in optimal control. We then employ this approach to solve the robust flight plan optimization problem at the planning horizon. In order to model uncertainty in the wind and estimating the probability of convective conditions, we employ Ensemble Prediction System (EPS) forecasts, which are composed by multiple predictions instead of a single deterministic one. The resulting method can be used to optimize flight plans for maximum expected efficiency according to the cost structure of the airline; additionally, predictability and exposure to convection can be incorporated as additional objectives. The inherent tradeoffs between these objectives can be assessed with this methodology.

The second part of this thesis presents a solution for the rerouting of aircraft in uncertain convective weather scenarios at the tactical horizon. The uncertain motion of convective weather cells is represented with a stochastic model that has been developed from the output of a deterministic satellite-based nowcast product, Rapidly Developing Thunderstorms (RDT). A numerical optimal control framework, based on the point-mass model with the addition of turn dynamics, is employed for optimizing efficiency and predictability of the proposed trajectories in the presence of uncertainty about the future evolution of the storm. Finally, the optimization process is initialized by a randomized heuristic procedure that generates multiple starting points. The combined framework is able to explore and as exploit the space of solution trajectories in order to provide the pilot or the air traffic controller with a set of different suggested avoidance trajectories, as well as information about their expected cost and risk.

The proposed methods are tested on example scenarios based on real data, showing how different user priorities lead to different flight plans and what tradeoffs are then present. These examples demonstrate that the solutions described in this thesis are adequate for the problems that have been formulated. In this way, the flight planning process can be enhanced to increase the efficiency and predictability of individual aircraft trajectories, which would lead to higher predictability levels of the ATM system and thus improvements in multiple performance indicators.

Resumen

El sistema de gestión del tráfico aéreo (*Air Traffic Management*, ATM) en los espacios aéreos más congestionados del mundo está siendo reformado para lidiar con múltiples desafíos socioeconómicos, medioambientales y de capacidad. Un pilar de este proceso es el gradual reemplazo de las estructuras rígidas de navegación, basadas en aerovías y *way-points*, hacia las operaciones basadas en trayectorias. No obstante, la implementación exitosa de este concepto y la realización de las ganancias esperadas en rendimiento ATM requiere entender y gestionar apropiadamente la incertidumbre. Debido a su compleja estructura socio-técnica, el diseño y operaciones del sistema ATM se encuentran marcadamente influidos por la incertidumbre, que procede de múltiples fuentes y se propaga por las interacciones entre subsistemas y operadores humanos.

Uno de los principales focos de incertidumbre en ATM es la meteorología. Debido a su naturaleza no-lineal y caótica, muchos fenómenos de interés no pueden ser pronosticados con completa precisión en cualquier horizonte temporal, lo que crea disrupción en las operaciones en aire y tierra que se propaga a otros procesos de ATM. Por lo tanto, para lograr los objetivos de SESAR e iniciativas análogas, es imprescindible tener en cuenta la incertidumbre en múltiples escalas espaciotemporales, desde la predicción de trayectorias hasta la planificación de flujos y tráfico.

Esta tesis aborda el problema de la planificación de vuelo de aeronaves individuales considerando dos fuentes importantes de incertidumbre meteorológica: el error en la predicción del viento y la actividad convectiva. Conforme la realización del viento se desvía de su previsión, la trayectoria real se desviará temporalmente de la planificada, lo que implica incertidumbre en tiempos de llegada a sectores y aeropuertos y en consumo de combustible. La actividad convectiva también tiene un impacto en la predictibilidad de las trayectorias, puesto que obliga a los pilotos a desviarse de sus planes de vuelo para evitarla, cambiando así la situación de tráfico. En este trabajo, buscamos desarrollar métodos y algoritmos para la optimización de trayectorias que puedan integrar información sobre la incertidumbre en estos fenómenos meteorológicos en el proceso de diseño de planes de vuelo en horizontes de planificación (antes del despegue) y tácticos (durante el vuelo), con el objetivo de generar trayectorias más eficientes y predecibles.

Con este fin, formulamos la planificación de vuelo como un problema de control óptimo, modelando la dinámica del avión con un modelo de masa puntual y el modelo de rendimiento BADA. El control óptimo es un marco flexible y general con un largo historial de éxito en el campo de la ingeniería aeroespacial. Como método numérico, empleamos métodos directos, que son capaces de manejar sistemas dinámicos de alta dimensión con costes computacionales moderados. No obstante, si bien esta metodología

es madura en contextos deterministas, la solución de problemas prácticas de control óptimo bajo incertidumbre en la literatura no está tan desarrollada, y los métodos propuestos en la literatura no son aplicables al problema de interés.

La primera contribución de esta tesis hace frente a este reto mediante la introducción de un marco numérico para la resolución de problemas generales de control óptimo no-lineal bajo incertidumbre paramétrica. El núcleo de este método es un esquema de conjunto de trayectorias, en el que las trayectorias del sistema dinámico bajo múltiples escenarios son consideradas de forma simultánea, y el problema de control óptimo bajo incertidumbre es así transformado en un problema convencional que puede ser tratado mediante métodos existentes en control óptimo. A continuación, empleamos este método para resolver el problema de la planificación de vuelo robusta. La incertidumbre en el viento y la probabilidad de ocurrencia de condiciones convectivas son modeladas mediante el uso de previsiones de conjunto o *ensemble*, compuestas por múltiples predicciones en lugar de una única previsión determinista. Este método puede ser empleado para maximizar la eficiencia esperada de los planes de vuelo de acuerdo a la estructura de costes de la aerolínea; además, la predictibilidad de la trayectoria y la exposición a la convección pueden ser incorporadas como objetivos adicionales. El *trade-off* entre estos objetivos puede ser evaluado mediante la metodología propuesta.

La segunda parte de la tesis presenta una solución para reconducir aviones en escenarios tormentosos en un horizonte táctico. La evolución de las células convectivas es representada con un modelo estocástico basado en las proyecciones de Rapidly Developing Thunderstorms (RDT), un sistema determinista basado en imágenes de satélite. Este modelo es empleado por un método de control óptimo numérico, basado en un modelo de masa puntual en el que se modela la dinámica de viraje, con el objetivo de maximizar la eficiencia y predictibilidad de la trayectoria en presencia de incertidumbre sobre la evolución futura de las tormentas. Finalmente, el proceso de optimización es inicializado por un método heurístico aleatorizado que genera múltiples puntos de inicio para las iteraciones del optimizador. Esta combinación permite explorar y explotar el espacio de trayectorias solución para proporcionar al piloto o al controlador un conjunto de trayectorias propuestas, así como información útil sobre su coste y el riesgo asociado.

Los métodos propuestos son probados en escenarios de ejemplo basados en datos reales, ilustrando las diferentes opciones disponibles de acuerdo a las prioridades del planificador y demostrando que las soluciones descritas en esta tesis son adecuadas para los problemas que se han formulado. De este modo, es posible enriquecer el proceso de planificación de vuelo para incrementar la eficiencia y predictibilidad de las trayectorias individuales, lo que contribuiría a mejoras en el rendimiento del sistema ATM.

Glossary

Δt	Step size
Σ_w	Covariance matrix of the motion of the storm
Ω	Space of possible outcomes
β	Sideslip angle
ζ	Defect constraint
γ	Flight path angle
λ	Longitude
$\bar{\lambda}$	Costate vector
μ	Bank angle
μ_h	Lagrange multipliers associated to the equality constraints
μ_L	Lagrange multipliers associated to the lower inequality bounds
μ_U	Lagrange multipliers associated to the upper inequality bounds
ν	Lagrange multipliers associated to the boundary conditions
ξ	Random parameters vector
ρ	Air density
σ_w	Amplitude of the random component of the motion of the storm
ψ	Initial and final conditions function
φ	Mayer (terminal cost) term
ϕ	Latitude
χ	True heading
χ_f	Storm heading
ω	Outcome
ω_w	Angular speed of the wind-fixed reference frame
\mathbb{E}	Expectation operator
\mathbb{I}	Identity matrix
\mathbb{P}	Probability measure

\mathbb{R}	Real line
\mathcal{F}	σ -algebra of events
\mathcal{H}	Hamiltonian
\mathcal{H}°	Augmented Hamiltonian
\mathcal{I}	Set of forecast steps
\mathcal{J}	Set of possible storm phases
\mathcal{L}	Lagrangian (running cost) function
\mathcal{T}	Time interval of interest
\mathbf{A}	Aerodynamic force vector
\mathbf{D}	Drag vector
\mathbf{F}	Force vector
\mathbf{K}	Feedback policy
\mathbf{K}_p	Parametrized feedback policy
\mathbf{L}	Lift vector
\mathbf{T}	Thrust vector
\mathbf{W}_t	Multidimensional Wiener process
\mathbf{X}_t	State vector of a stochastic process
\mathbf{Y}	Lateral force vector
\mathbf{c}	Position of the center of the storm
\mathbf{f}	Dynamical function
\mathbf{g}	Inequality constraints vector
\mathbf{g}_L	Lower bounds of the inequality constraints
\mathbf{g}_U	Upper bounds of the inequality constraints
\mathbf{h}	Equality constraints vector
\mathbf{i}	First basis vector
\mathbf{j}	Second basis vector
\mathbf{k}	Third basis vector
\mathbf{p}	Parameters vector
\mathbf{u}_L	Feedback law for the control vector
\mathbf{u}	Control vector

\mathbf{u}_q	Tracked part of the control vector
\mathbf{u}_r	Untracked part of the control vector
\mathbf{v}	Airspeed vector
\mathbf{v}_s	Storm speed vector
\mathbf{v}_w	Wind vector
\mathbf{x}	State vector
\mathbf{x}_q	Tracked part of the state vector
\mathbf{x}_r	Untracked part of the state vector
\mathbf{y}	Vector of reference variables
\mathbf{z}	Algebraic variables vector
\mathbf{z}_q	Tracked part of the algebraic vector
\mathbf{z}_r	Untracked part of the algebraic vector
CR	Convective Precipitation
D	Drag force
L	Lift force
LP	Linear predictor
T	Thrust force
TT	Total Totals
Y	Lateral force
C_D	Coefficient of drag
C_L	Coefficient of lift
C_T	Coefficient of thrust
EC	Exposition to convection
E_u	Tracking matrix associated to the control vector
E_x	Tracking matrix associated to the state vector
E_z	Tracking matrix associated to the algebraic variables
H	Geopotential altitude
H_p	Geopotential pressure altitude
$I_{-i,j}$	Indicator functions
J	Cost functional

- M Mach number
- P Pressure
- P_k Shifted Legendre polynomial of degree k
- R_E Mean radius of the Earth
- R_M Ellipsoid radius of curvature in the meridian
- R_N Ellipsoid radius of curvature in the prime vertical
- S Wetted wing surface
- S_W Wind-fixed reference system
- T Temperature
- V Value function
- W_t 1-dimensional Wiener process
- X Random variable
- Z_k Standard Gaussian variable
- a Equatorial radius of the Earth (6378.137 km)
- a_0 Speed of sound
- a_k Weighting sequence coefficient of degree k
- b Polar radius of the Earth (6356.752 km)
- c Probability of convective conditions
- d Control degrees of freedom
- e_o Eccentricity of the Earth
- f Ellipsoidal flattening
- f_c Fuel consumption
- g Gravitational acceleration
- h Geodetic altitude
- m Aircraft mass
- n_ξ Dimension of the uncertain parameters vector
- n_Ψ Number of boundary conditions
- n_g Number of inequality constraints
- n_h Number of equality constraints
- n_u Dimension of the control vector

n_x	Dimension of the state space
n_z	Dimension of the vector of algebraic variables
q	Dynamic pressure
p	Storm probability field
r_p	Risk penalty
t	Time variable
t_0	Time at the start of the trajectory
t_f	Time at the end of the trajectory
v	True airspeed
v_{CAS}	Calibrated airspeed
v_g	Ground speed
v_f	Forecasted storm speed
w_x	Wind speed, West-to-East component
w_y	Wind speed, South-to-North component
$w_{i,j}$	Statistical weights of the indicator functions
u_i	i -th control variable
x_i	i -th state variable
N	Number of discrete scenarios

Acronyms

2PBVP	Two-Point Boundary Value Problem
ANS	Air Navigation Services
ANSP	Air Navigation Services Provider
APM	Aircraft Performance Model
ATC	Air Traffic Control
ATM	Air Traffic Management
AUC	Area Under Curve
BDT	Business Developed Trajectory
BADA	Base of Aircraft DATA
CAS	Calibrated Airspeed
CCO	Continuous Climb Operations
CDO	Continuous Descent Operations
CI	Cost Index
CIWS	Corridor Integrated Weather System
CNS	Communications, Navigation and Surveillance
CONUS	Conterminous United States
CoSPA	Consolidated Storm Prediction for Aviation
CWAM	Convective Weather Avoidance Model
DASA	Dynamic Airspeed Adjustment
DAE	Differential-Algebraic Equation
DP	Dispersion Penalty
DST	Decision Support Tool
DROCT	Discretized Robust Optimal Control Problem with Tracking
EAS	Equivalent Airspeed
ECEF	Earth-Centered, Earth-Fixed
ECMWF	European Centre for Medium-Range Weather Forecasts
EEC	Eurocontrol Experimental Centre

EGM2008	Earth Gravitational Model 2008
EGM96	Earth Gravitational Model 1996
EPS	Ensemble Prediction System
FAA	Federal Aviation Authority
FMS	Flight Management System
GFS	Global Forecast System
GLM	Generalized Linear Model
gPC	Generalized Polynomial Chaos
GPU	Graphics Processing Unit
GRASP	Greedy randomized adaptive search procedure
GRS	Geodetic Reference System
HDF5	Hierarchical Data Format version 5
HJB	Hamilton-Jacobi-Bellman
HP	High Predictability
IAS	Indicated Airspeed
ICAO	International Civil Aviation Organization
ISA	International Standard Atmosphere
LAM	Limited Area Model
MAE	Maximum Average Efficiency
MDP	Markov Decision Process
MINLP	Mixed-Integer Nonlinear Programming
MSG	Meteosat Second Generation
NCAR	National Center for Atmospheric Research
NLP	Non-Linear Programming
NOAA	National Oceanic and Atmospheric Administration
NWCSAF	Satellite Application Facility on support to Nowcasting and Very Short-Range Forecasting
NWP	Numerical Weather Prediction
OCP	Optimal Control Problem
ODE	Ordinary Differential Equation

PDE	Partial Differential Equation
PEARP	Prévision d'Ensemble ARPEGE
RBT	Reference Business Trajectory
RDE	Randomly-parametrized Differential Equation
RDT	Rapid Developing Thunderstorms
ROC	Receiver Operating Characteristic
ROCT	Robust Optimal Control Problem with Tracking
RRT	Rapidly-exploring random trees
SBT	Shared Business Trajectory
SESAR	Single European Sky ATM Research
SEVIRI	Spinning Enhanced Visible and Infrared Imager
SDE	Stochastic Differential Equation
SOCP	Stochastic Optimal Control Problem
SQR	Stochastic Quadrature Rule
SWIM	System Wide Information Management
TAS	True Airspeed
TBO	Trajectory-Based Operations
TIGGE	The International Grand Global Ensemble
TP	Trajectory Prediction
TRACON	Terminal Radar Approach Control
UAS	Unmanned Aerial Systems
UQ	Uncertainty Quantification
VHP	Very High Predictability
VIL	Verrtically Integrated Liquid
WGS84	World Geodetic System 1984 reference system
WMO	World Meteorological Organization

Contents

Acknowledgements	iii
Published and Submitted Content	v
Other Research Merits	vii
Abstract	ix
Resumen	xi
Glossary	xiii
Acronyms	xix
1 Introduction	1
1.1 Motivation	1
1.2 Objectives	7
1.3 Contributions	8
1.4 Methods	9
1.4.1 Ensemble Forecasts	9
1.4.2 Rapidly Developing Thunderstorms	10
1.4.3 Numerical Optimal Control	11
1.5 Outline of the Thesis	12
2 Trajectory Optimization	15
2.1 Dynamic Optimization Paradigms	16
2.1.1 Pre-computed Control Laws	16
2.1.2 Receding Horizon Control	18
2.2 Algorithmic Trajectory Planning	19
2.2.1 Heuristic and Geometric Methods	19
2.2.2 Dynamic Programming	19
2.2.3 Path-Planning Algorithms	20
2.2.4 Combinatorial Optimization	21
2.2.5 Metaheuristic Methods	21

2.2.6	Optimal Control	22
2.3	Deterministic Optimal Control	22
2.3.1	Deterministic Differential Equations	22
2.3.2	Formulation of the Optimal Control Problem	23
2.3.3	Pontryagin's Minimum Principle	24
2.3.4	The Hamilton-Jacobi-Bellman Equation	26
2.4	Numerical Methods in Deterministic Optimal Control	27
2.4.1	HJB methods	27
2.4.2	Indirect Methods	28
2.4.3	Direct Methods	29
2.5	Uncertainty and Differential Equations	32
2.5.1	Random Variables	32
2.5.2	Randomly-parametrized Differential Equations	33
2.5.3	Stochastic Differential Equations	34
2.5.4	Discretization of Random Variables	35
2.6	Dynamic Optimization under Uncertainty	38
2.6.1	Dynamic Programming	38
2.6.2	Stochastic Optimal Control	39
2.6.3	Uncertainty Quantification Optimal Control	40
2.6.4	Robust Optimal Control	41
2.7	Flight Planning Applications	41
2.7.1	Deterministic Flight Planning	41
2.7.2	Wind-optimal Aircraft Trajectories	43
2.7.3	Aircraft Flight Planning under Uncertainty	44
2.7.4	Convective Environments	45
3	Modeling	47
3.1	Aircraft Dynamics	48
3.1.1	Reference Systems	48
3.1.2	Model of Forces	53
3.1.3	The 4D Equations of Motion	56
3.1.4	Reduced Models	60
3.1.5	Aircraft Performance Limitations	64
3.2	Meteorological Modeling	65
3.2.1	Environmental and Weather Model	65
3.2.2	Ensemble Forecasts	67
3.3	Convection	69
3.3.1	Planning	70
3.3.2	Tactical	71

4	Robust Optimal Control	75
4.1	The State-Tracking Formulation	76
4.2	Definition of the Tracking Problem	79
4.3	Stochastic Quadrature Rules	80
4.4	The Trajectory Ensemble	81
4.5	Formulation of the Discretized Problem	82
4.6	Feedback Policies	84
5	Robust Flight Planning	87
5.1	Concept	88
5.1.1	Dynamical Model	88
5.1.2	Coordinate Transformation	90
5.1.3	Trajectory Ensemble	92
5.1.4	Formulation of the Discretized Problem	93
5.1.5	Dynamic Airspeed Adjustment	94
5.2	Results	97
5.2.1	Fixed Airspeed	98
5.2.2	Variable Airspeed	101
5.2.3	Dynamic Airspeed Adjustment (DASA)	105
6	Flight Planning under Convection	111
6.1	Concept	112
6.2	Results	114
6.2.1	Constant Airspeed	114
6.2.2	Variable Airspeed	119
7	Modeling Uncertainty in Convective Nowcasts	129
7.1	Deterministic Input	130
7.2	Stochastic Motion	132
7.3	Statistical Model	133
7.4	Model Fitting	134
8	Optimal Stochastic Storm Avoidance	141
8.1	Problem Modeling	142
8.1.1	Dynamical Model	142
8.1.2	Optimal Control Model	143
8.2	Initialization	144
8.3	Results	146
8.3.1	Computational Setup	146
8.3.2	Global Optimum	147

8.3.3	Local Optima	148
9	Conclusions and future work	155
9.1	Future Work	157
9.1.1	Validation	157
9.1.2	Extended Aircraft Dynamics	157
9.1.3	Multiple Uncertainty Sources	158
9.1.4	Fast Quasi-Optimal Heuristics	158
9.1.5	Estimation of the Cost of Convection	159
9.1.6	Stochastic Convection Extensions	159
A	Tracking Feasibility for Linear Systems	183

Introduction

1.1 Motivation

Air traffic in Europe and all over the world is expected to expand substantially in the coming decades. According to Eurocontrol's *Challenges of Growth 2018* report [15], the total amount of movements in European airspace under the most likely scenario is expected to grow at an average rate of 2% per year between 2017 and 2040, leading to a total increment of 50% in 2040 with respect to the current traffic levels. In the most optimistic forecast, corresponding to strong economic performance in Europe and the rest of the world, the total amount of flights could nearly double in this period.

It is, therefore, considered likely that the growth in air traffic will be high enough to generate substantial capacity challenges for the whole European Air Traffic Management (ATM) system in the future. On the ground, a substantial shortfall in airport capacity is forecasted, with 1 to 4 million flights remaining unaccommodated in 2040 under current airport expansion plans. At the network level, flow management regulations are projected to become more frequent in order to deal with denser traffic, producing around 5 times more flow management-originated delays during summers (6.2 minutes per flight, on average, compared to 1.2 minutes per flight) and additional reactionary delay. Looking beyond the average numbers, the Eurocontrol projections paint a starker picture in the extremes, with the amount of flights delayed between 1 and 2 hours increasing sevenfold, affecting $\sim 470,000$ passengers (compared to 50,000 today) and increasing flight cancellations.

In addition, the ATM community will have to face the numerous challenges posed by climate change. In first place, environmental policies oriented towards mitigation will put pressure on the aviation industry, which in 2010 accounted for 2.4% of climate change-related emissions [16], to adopt technologies and operating procedures with reduced environmental impact. Nowadays, the main channels through which aviation contributes to climate change are CO_2 and NO_x emissions related to fuel burn, but condensation trails also play an important role [17]. Furthermore, aviation will need to adapt to the shifting weather patterns brought along with climate change. Extreme heat, precipitation or snow events, convective activity, wind and temperature changes

and sea-level rise have the potential to disrupt operations¹ and shift passenger demand between origin-destination markets. Finally, other environmental concerns, such as noise near aerodromes, will likely continue to play a role in the future.

These critical issues are not alone within the set of priorities that have been put forward by policymakers. Maintaining or improving the excellent safety record in aviation in the face of increased traffic and new technologies (such as the incorporation of Unmanned Aerial Systems (UAS) into the airspace) is a major concern. Additionally, the trend towards an increasingly interconnected ATM infrastructure will require the provision of physical security and cybersecurity against traditional and emerging threats. Finally, as an industry of key economic importance, improving the cost-effectiveness of ATM services will remain a highly sought-after objective.

As these challenges began to be recognized, major policy initiatives were put in place during the last decade in the most complex airspaces in the world. The European Union created the Single European Sky ATM Research (SESAR) project to provide the technological solutions to overhaul and unify the European airspace. Similarly, the Federal Aviation Authority (FAA) launched NextGen in the US with similar objectives; indeed, an agreement was signed in 2010 to promote interoperability between the technologies employed in both programs.

These initiatives aim to improve the performance of the ATM system in a wide range of aspects. The 2015 edition of the European ATM Master Plan has established high-level goals for 2035, which are summarized in Table 1.1. In order to attain these objectives, these programs develop, implement and leverage new technologies, concepts of operations, regulations, and organizational paradigms in every part of the ATM system.

Area	Relative improvement target
ANS cost-efficiency	50% reduction in unit costs
Operational efficiency	5-10% reduction in fuel burn
	3-6% reduction in flight time
ATM capacity	3-fold increase
Environmental impact	10% reduction in emissions per flight
Safety	Improvement by a factor of 3-4
Security	No incidents leading to traffic disruption

Table 1.1: Single European Sky performance ambition for 2035.

¹See *European Aviation in 2040: Challenges of Growth* (2018), “Annex 2: Adapting Aviation to a Changing Climate” (<https://www.eurocontrol.int/sites/default/files/publication/files/challenges-of-growth-annex-2-01102018.pdf>)

One of the core pillars of the future ATM system under the SESAR vision is the Trajectory-Based Operations (TBO) concept. It represents a new, more flexible paradigm in organizing and managing air traffic that seeks to enhance the efficiency and capacity of the airspace while maintaining or increasing its safety. Within the TBO concept, the 4D trajectory² becomes the centerpiece of ATM operations, thanks to the introduction of technologies enabling precise navigation, accurate trajectory prediction and fast data-sharing between the aircraft and the control centers.

In the future TBO concept, the planned trajectory will be elaborated in an iterative, collaborative process. In first place, the airline calculates its preferred trajectories at a long-term, strategic horizon, producing a Business Developed Trajectory (BDT). At some point, months to days before the trajectory is executed, the BDT is shared with the Air Navigation Services Provider (ANSP), turning into a Shared Business Trajectory (SBT). The ANSPs employ this information to allocate resources to sectors and determine route availability. The Network Manager is then able to compare the demand from all SBTs to the capacity that the ANSPs can provide in order to assess network imbalances and determine appropriate corrective measures by modifying, delaying or approving all the SBTs for the time period under study. A negotiation with all the involved stakeholders then leads to an agreed Reference Business Trajectory (RBT) to be flown by the airline. Finally, the trajectory is flown according to the RBT and possibly modified again in response to unplanned events.

The successful implementation of the TBO concept, and the attainment of the desired improvements in Key Performance Areas, rests upon appropriate management of uncertainty. As a complex social and technological system, ATM features a significant amount of uncertainty, which impacts all areas where progress is sought. All systems, organizations and people involved in ATM operations face uncertainty both directly from external sources and indirectly from other parts of the ATM system, exhibiting an interconnected pattern of influence that extends into multiple domains for any individual source of uncertainty. A flight being delayed on arrival because of meteorological or equipment issues can influence airport resource management and create reactionary delay in subsequent flights, creating deviations in sector demand with respect to the planned capacity and contributing to the installment of regulations that affect seemingly unrelated flights.

Therefore, it is essential to take uncertainty into account when designing the systems and concepts that will form the basis of the future ATM system. Specifically, it is necessary to reduce, mitigate and deal with uncertainty in order to achieve the desired improvements in capacity and efficiency without compromising safety. Furthermore, as we move towards a more trajectory-centric paradigm, improving our understanding of

²Within the context of this work, “4D” refers to the 3 spatial dimensions plus time, i.e., a “4D trajectory” represents not just a spatial path but also a temporal schedule for flying it.

the sources of uncertainty at the trajectory level becomes critical for the processes of trajectory prediction, planning, optimization, execution, monitoring and synchronization, which are the pillars of the TBO concept. In particular, the ability to design trajectories that are efficient, but also more predictable, must be developed in order to reach the targeted improvements in capacity, cost-efficiency and environmental impact. By reducing the uncertainty in trajectory computation and generation, it would be possible to reduce not just uncertainty at the individual trajectory level, but also downstream, system-level effects, allowing for increased capacity.

At the trajectory level, uncertainty proceeds from many sources, which may directly influence from a single flight to multiple interacting flights or even the entire network. According to the taxonomy presented in [18, Chapter 4], these sources can be classified in the following categories:

- **Data uncertainty**, arising from inexact or unavailable knowledge of relevant variables, such as mass [19] or aircraft position, or models, such as aerodynamic or engine performance [20].
- **Equipment uncertainty**, resulting from malfunction or failure of aircraft systems or ground systems, leading to deviations from normal modes of operation. Sensors, actuators, CNS equipment, airport and ATC facilities may face issues leading to abnormal behaviour.
- **Weather uncertainty**, produced by imperfect knowledge of current and future values of key atmospheric variables and the locations of adverse meteorological conditions.
- **Operational uncertainty**, generated by the lack of complete predictability of the choices of decision-makers in charge of aircraft, airports and airspace operations. Varying adherences to the flight plan or procedures [21], conflict resolution patterns and choices of regulations can be grouped in this category.

One uncertain factor of particular interest is meteorology. Weather influences many aspects of aviation, from aircraft design to air traffic planning, and it manifests itself as a wide range of relevant phenomena at multiple spatiotemporal scales. Forecasting weather is, additionally, an inherently uncertain endeavour: unavoidable inaccuracies due to incomplete knowledge of the current state of the atmosphere, computational limitations, and errors in the modeling of physical processes (particularly at the subgrid scale) are amplified by the nonlinear chaotic dynamics of the atmosphere. From the point of view of trajectory predictability, two channels that make a significant contribution to uncertainty can be identified:

- The magnitude and direction of wind has a direct impact on groundspeed. Therefore, as the realized wind deviates from its forecasted value, the actual trajectory will deviate in time from the prediction produced from the flight plan, introducing unpredictability into the trajectory prediction process. This is not only relevant for the airline, which faces costs from the uncertainty in fuel consumption and arrival times, but also for airports and controllers, which will deal with aircraft sooner or later than expected. As a result, the airspace is not used as efficiently as it would be in an ideal deterministic situation, in order to accommodate these deviations.
- Convective weather entails numerous hazards to aircraft, including strong turbulence, wind shear, downbursts, icing, lightning or hail. Therefore, pilots generally avoid regions of strong convection by deviating from the flight plan, if possible. However, convection is particularly challenging to forecast, and thus it is not currently possible to plan these deviation routes at lead times beyond the minutes or tens of minutes with solid assurances. As a consequence, at the timescale of hours there is significant uncertainty in trajectories and traffic patterns whenever and wherever convection is possible.

Thus, the spatiotemporal variability patterns of these phenomena creates interlinked challenges at different stages of the flight planning and traffic planning processes. Indeed, the estimation and modeling of uncertainty in individual aircraft trajectories due to wind [22, 23, 24, 25, 26] and its impact on conflict detection and resolution [27, 28] represents an active area of research; similarly, the uncertainty due to convection is being studied at both the trajectory level [29, 30, 31] and the network level [32]. The interactions between weather and other uncertainties are also being studied in recent works [33, 34].

Understanding the effect of uncertainty in trajectory prediction, while essential to attain the desired improvements in ATM performance under the future concept of operations, must be accompanied by the development of flight planning systems and algorithms that take uncertainty into account. However, the kind of algorithms that are currently employed in practice and in the literature to produce flight plans do not pay enough attention to these uncertainties. In order to limit the scope of this thesis, we will focus on the trajectory planning process at two scales:

- A *planning horizon* of 2 or 3 hours before departure where the flight plan associated to the RBT is elaborated. At this stage, the uncertainty due to wind is significant, and particularly so for the later stages of medium- and long-haul flights: for those flight legs, the lead time of the weather forecast can climb to more than 10 or 12 hours. In addition, the error in arrival times due to wind accumulates over the whole flight. Individual convective events cannot be forecasted at this stage;

however, it is possible to identify regions where the conditions for convection are met or where convection is more likely to develop.

- A *tactical horizon* that comprises scenarios where the flight is already taking place and where the trajectory replanning scenarios last from around 15 minutes to 2 hours. In flights encountering convective events, the uncertainty proceeding from convective developments becomes more important than the uncertainty due to wind in relative terms, since the aircraft derives its own wind data from sensor readings and the prediction error does not have enough time to accumulate to such an extent that it rivals uncertainty in the route.

At the planning horizon, a key Numerical Weather Prediction (NWP) technique to represent meteorological uncertainty is *ensemble forecasting* [35]. An Ensemble Prediction System (EPS) does not produce a single deterministic forecast; instead, it generates a set of multiple forecasts representing possible realizations of the future state of the atmosphere. EPS forecasts are being employed in some ATM studies for purposes such as developing algorithms for flow management under uncertainty [36] or estimating the potential reduction of airspace capacity due to adverse weather [37]. Nevertheless, its use is still not widespread among the wider ATM community, where deterministic forecasts such as the Global Forecast System (GFS) are employed to a greater extent.

At the trajectory level, the interest in employing EPS forecasts to estimate the uncertainty in flight time has been growing in recent years. Some authors [38, 39] have studied the problem of employing filtering techniques to blend aircraft observations with ensemble forecasts in order to produce more accurate estimations of the wind field. In the SESAR project IMET, the information from probabilistic meteorological forecasts was combined with a deterministic trajectory prediction tool in order to create a probabilistic trajectory prediction (PTP) system [40]. It was shown that uncertainty information about timeliness and fuel consumption could then be produced at multiple look-ahead times for a given flight plan, and that this information could be employed for decision-making in flight planning [41]. However, the question of how to explicitly incorporate ensemble forecasts in a fully probabilistic flight planning framework was left for future research and had not been addressed in other works.

At a tactical horizon, there is a greater recognition of the role of adverse weather and convective developments in particular. As the current cockpit technology is not enough to guarantee safe operations in heavy weather [42], multiple families of trajectory planning algorithms have been published to address this problem; we will present a brief survey of the corresponding literature in Section 2.7.4. However, the existing methods for trajectory planning do not generally take both aircraft dynamics and the uncertainty in the location and evolution of convective cells into account, and those that do tend to exhibit serious computational limitations, as discussed in Chapter 2. Thus, there is a

need to develop flexible methods that address the problem of realistic trajectory planning in convective environments under uncertainty with manageable computational costs.

The motivation of the research described in this thesis is, therefore, to enhance our understanding of the impact of meteorological uncertainty in flight planning within the incoming ATM paradigm and to contribute towards techniques that allow for uncertainty-aware optimization of flight plans, facilitating the enhancement of predictability and the understanding of the trade-offs involved.

The research interest on these themes has been increasing more recently. Indeed, concurrently with the works that form the basis of this thesis, several authors have published articles on related topics. In [43], a probabilistic trajectory prediction algorithm based on meteorological probabilistic forecasts is built, based on analytical propagation of the distribution aircraft mass and the flyover time; a similar approach is employed in [44] to assess aircraft conflict severity. At a tactical horizon, the problem of trajectory planning in uncertain convective environments is addressed in [6] with a scheme based on reachability analysis and dynamic programming. Finally, moving from an individual trajectory focus towards traffic-scale planning, the work in [5] presents an algorithm to reduce forecasted conflicts in a region by slightly modifying the flight plans many hours before execution using the information from ensemble forecasts.

1.2 Objectives

The main goal of the current dissertation is the development of algorithms and techniques for flight planning at the planning horizon and the tactical horizon under uncertain meteorological environments. To that end, the following intermediate objectives have been pursued:

- Reviewing the existing literature on methods for aircraft trajectory optimization.
- Modeling the flight planning problem under uncertainty.
- Processing relevant meteorological data and modeling uncertainty and convective activity.
- Developing dynamic optimization methods under uncertainty that allow us to address the formulated flight planning problems.
- Applying these methods to build a numerical tool for robust flight planning.
- Studying test scenarios to gain insight into the available options for improving predictability through flight planning.

1.3 Contributions

The original contributions presented in this thesis include:

1. In first place, we introduce a numerical framework for the solution of general nonlinear optimal control problems under uncertainty. It is based on an ensemble trajectory scheme, where the trajectories of the system under multiple scenarios are considered simultaneously within a single augmented dynamical system. The framework we propose generalizes previous, more limited concepts proposed in the literature, which are based on an “open-loop” concept that is not adequate for a wide range of dynamical systems. It transforms an uncertain optimal control problem into a conventional optimal control problem, with a size proportional to both the dimension of the state space and the number of uncertainty samples, that can be then solved by standard, well-studied direct methods in optimal control.
2. At the planning horizon, we propose a robust flight plan optimization methodology based on contribution #1, which can be employed to produce flight plans under uncertainty at the planning horizon. It relies on EPS forecasts for the modeling of the uncertainty in the wind and regions of potential convective activity. It can be used to optimize flight plans for maximum average efficiency according to the cost structure of the airline, but it also allows for the consideration of predictability and exposure to convection as additional objectives; the inherent tradeoff between these objectives can be assessed with this methodology. An analysis of how the flight plans are computed by the proposed method depending on the priorities of the user is offered.
3. Finally, at the tactical horizon, we present a solution for the tactical rerouting of aircraft in uncertain convective weather scenarios (again based on direct methods for optimal control). It combines three different components to propose a set of optimized avoidance trajectories: a stochastic model of the motion of the storms based on a deterministic extrapolation-based nowcast, a numerical optimal control framework for optimizing for efficiency and predictability in the face of convection, and a randomized heuristic procedure to generate multiple starting points for the optimization. The resulting method is able to explore as well as exploit the space of solution trajectories and provide the pilot or the air traffic controller with a set of different proposed avoidance trajectories, as well as the differences in expected cost and risk between them.

1.4 Methods

We now proceed to provide an overview of the main pillars of the work presented in this thesis. More detailed descriptions will be offered in Chapter 2, in the case of optimal control and trajectory optimization, and in Chapter 3, for meteorological data sources.

1.4.1 Ensemble Forecasts

NWP faces numerous uncertainties, from incomplete or imprecise knowledge of the state of the atmosphere at the time of the forecast to model uncertainties and computational limitations [45]. Because these uncertainties are propagated through nonlinear and chaotic atmospheric dynamics and error compounds with the lead time, simple statistical characterization is inadequate for the representation of the forecast uncertainty. Ensemble forecasting produces instead probabilistic forecasts by running a certain number (generally between 10 and 100) of different simulations or “members”, each one with different initial conditions, models or parameters. Ensemble forecasting has been one the drivers behind the improvement of weather prediction techniques over the last decades [46].

Several methods are used for producing the different members of an EPS forecast (see Figure 1.1). The most important include ensemble data assimilation, where the initial conditions are perturbed according to their uncertainty; singular and breeding vectors, where the perturbation is chosen to excite the fastest-growing dynamical instabilities; stochastic parametrizations, where the parameters of models of the atmospheric processes at subgrid spatiotemporal scales are perturbed; and multiphysics or multi-model schemes, where the members are drawn from the output from different models. Each NWP center maintains an Ensemble Prediction System that implements a different combination of these techniques. We note that these methods aim to be “economic” by producing a good representation of the uncertainty from a limited number of simulations, since each one is computationally expensive.

There are three main classes of EPS. In decreasing scope and maturity and increasing resolution, they are global scale EPS, regional scale (Limited Area Model or LAM) EPS and convective scale EPS. For the purposes of considering wind uncertainty in medium-haul flights, both global and LAM EPS can be used. However, since LAMs are usually produced in order to forecast surface weather of a specific region (such as Europe or North America), their oceanic coverage is limited. Therefore, in this work we will rely on global EPS, which can be used for intercontinental flights. Historical global EPS forecasts from major NWP centres around the world can be found at the

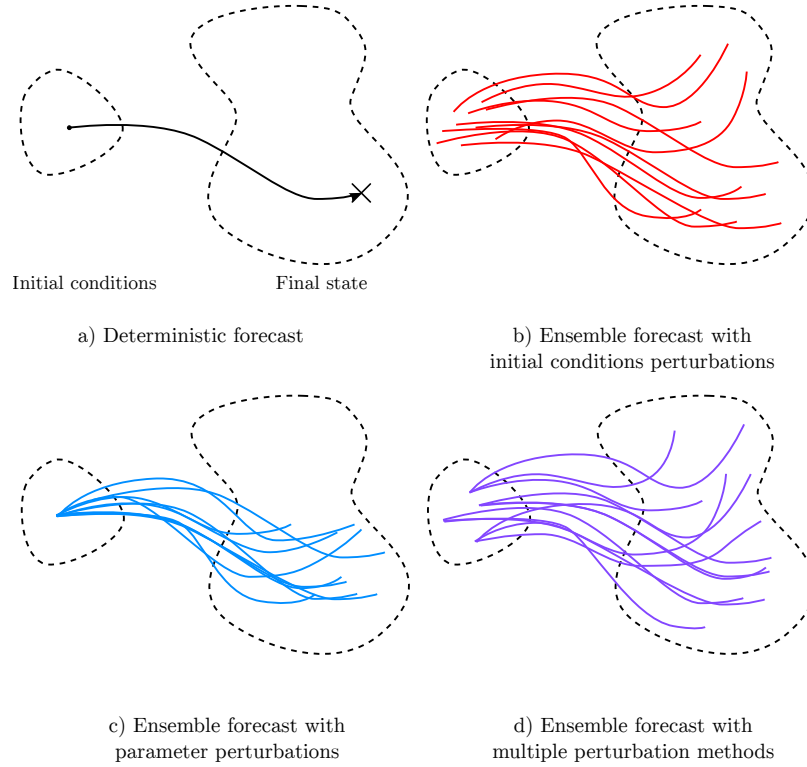


Figure 1.1: Conceptual representation of the generation of EPS members.

TIGGE dataset³, hosted at the European Centre for Medium-Range Weather Forecasts (ECMWF) website. Output from European LAM models is similarly compiled in the TIGGE LAM dataset.

Ensemble forecasts provide us with useful uncertainty information at a planning horizon. In first place, the estimated uncertainty in relevant variables for aircraft performance, such as wind, temperature and geopotential at barometric altitude, can be extracted from ensemble forecasts. Additionally, other variables can be employed to determine the likelihood that the necessary conditions for convection are present at each time and place. We will cover this process in Chapter 3.

1.4.2 Rapidly Developing Thunderstorms

At a planning horizon, it is hard to forecast the onset and location of convective activity with precision, creating numerous challenges for flight dispatchers, air traffic controllers and flow management authorities. Within the regions where conditions that

³<http://apps.ecmwf.int/datasets/data/tigge/>

allow for convective development are forecasted, it is not easy to predict the timing and location of convective initiation. The characteristic sizes and lifespans of convective phenomena are small compared with the spatiotemporal resolution of medium-range NWP models, and the chaoticity of the atmospheric dynamics compounds the challenge. As a consequence, both thunderstorm weather forecasting and avoidance take place at shorter timescales, in the form of deterministic nowcasting⁴ based on extrapolation of radar echoes or satellite imagery. This technique offers better accuracy than NWP in timescales below 1 or 2 hours.

In this work, we will employ the satellite-based Rapid Developing Thunderstorms (RDT) product as a data source for convective nowcasts. This nowcasting system takes as input data from multiple instruments aboard Meteosat Second Generation satellites and uses it to identify the locations and shapes of cloud formations, discriminate the convective cells, and forecast their motion. The resulting RDT output then provides the generated information in an object-oriented fashion, which can be used for trajectory planning; however, the forecast of the evolution of the storms is deterministic and the forecast error is therefore not represented. A more detailed description is offered in Section 3.3.2; we will enrich this nowcast with a probabilistic model in Chapter 7.

1.4.3 Numerical Optimal Control

A wide range of algorithms have been historically employed for the purposes of flight planning, including geometric, heuristic and metaheuristic methods, path-planning algorithms, dynamic programming approaches and optimal control techniques. As interest in 4D trajectory paradigms began to grow and computational power increased, researchers started to pay more attention to optimization techniques targeting the complete trajectory optimization problem [47]. Optimal control represents a natural framework for this purpose, and multiple classes of numerical methods can be employed for the practical solution of optimal control problems. In particular, direct methods represent a computationally attractive solution for aircraft trajectory optimization problems, where practical restrictions make indirect methods inconvenient and high-dimensional state spaces make dynamic programming computationally expensive.

However, while direct methods are well-developed for deterministic nonlinear optimal control problems, they are not as mature for problems with uncertainty. Some concepts, based on the aggregation of multiple deterministic trajectories (each one representing a realization of the uncertain parameters, corresponding to quadrature points in an uncertainty quantification rule) have been put forward to solve the uncertain problem

⁴Within a meteorological context, “nowcasting” refers to forecasting for short ranges (6 hours at most, according to the World Meteorological Organization (WMO)) at the mesoscale (weather phenomena of sizes between a few kilometers and a few hundred kilometers), and relying more on extrapolation of current sensor data (radar echoes, mainly) than on NWP techniques.

[48, 49, 50, 51]. In the mentioned works, a single optimal control sequence is obtained, which results in different state trajectories depending on the realization of the uncertain parameters. Because all the state variables are then uncertain at the final time, exact boundary conditions and partial state trajectories cannot be applied, thus rendering this approach inadequate for flight planning. Therefore, this research gap will be addressed in this thesis in order to achieve its objectives.

1.5 Outline of the Thesis

This document is structured as follows:

- Chapter 1, the current chapter, describes the motivation of the thesis, as well as the specific objectives and contributions.
- Chapter 2 provides further background information on the trajectory optimization methods employed in the thesis, briefly reviewing the literature on dynamic optimization and flight planning.
- Chapter 3 introduces the physical models that will be employed throughout the remainder of the thesis, covering both the aircraft dynamics and the atmospheric behaviour.
- Chapter 4 develops a robust optimal control methodology for general nonlinear problems.
- Chapter 5 employs the framework introduced in Chapter 2 to address the robust flight planning problem in the presence of uncertain winds. A test scenario is solved and a discussion of the results is offered.
- Chapter 6 extends the methodology of Chapter 5 to include convection at the planning horizon and discusses the results from a flight planning scenario.
- Chapter 7 describes a probabilistic model for the evolution of convective cells based on the deterministic RDT nowcasts.
- Chapter 8 describes an optimal control-based method to address the problem of tactical trajectory optimization in the presence of uncertain convective cells.
- Chapter 9 sums up the implications of the work, as well as discussing potential future avenues of research.

Figure 1.2 illustrates the structure of this dissertation, highlighting the data sources, algorithms and contributions of the thesis and associating them to the corresponding chapters.

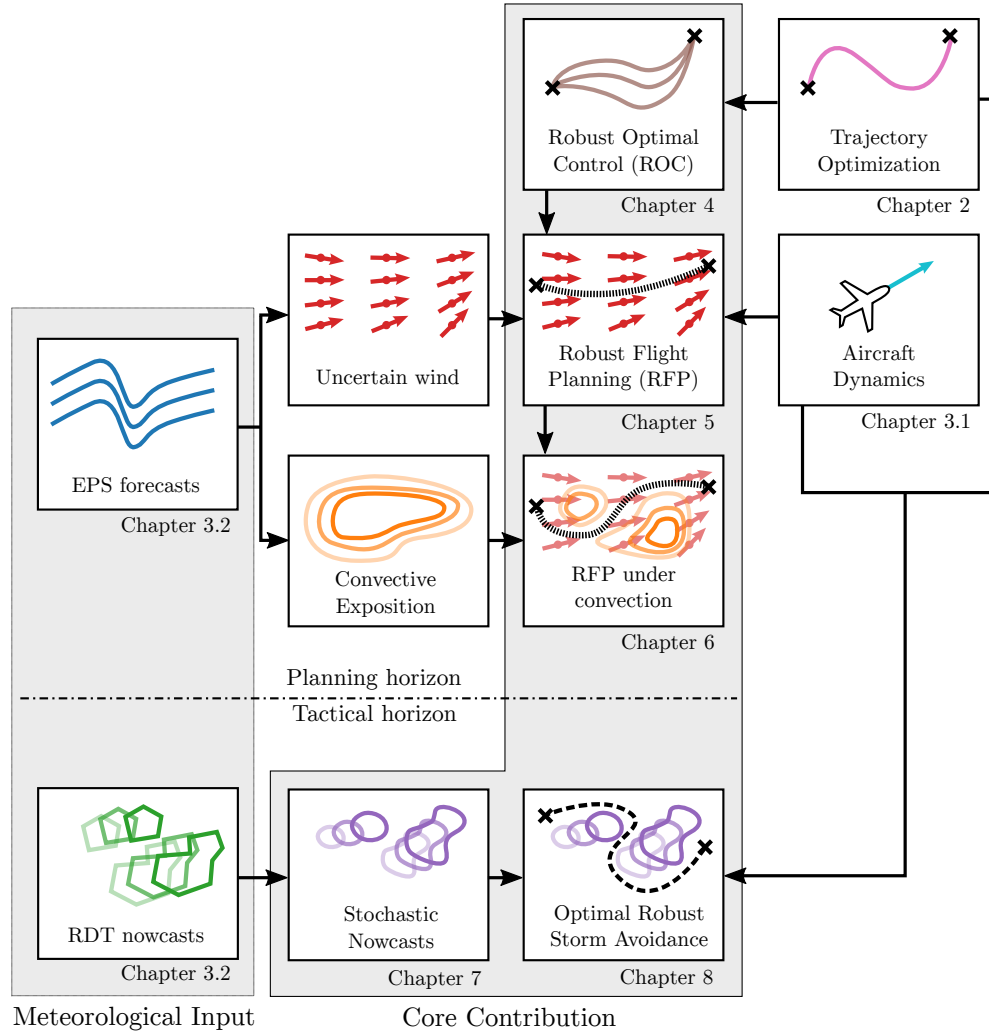


Figure 1.2: Thesis structure

Trajectory Optimization

The goal of flight planning is to produce a feasible sequence of aircraft operations that best meets the objectives of the flight planner while fulfilling practical constraints. Thus, from the mathematical point of view, a 4D flight planning problem represents a *dynamic optimization* problem, an optimization problem with a dynamical system at its core. A large number of engineering problems can be framed as dynamic optimization problems, and thus a wide range of solution methods has been studied. These methods differ not only on their computational performance, features and approach (*how* they solve dynamic optimization problems), but also in the practical concept that corresponds to the solution (*what* kind of practical problem is being solved). In this chapter, we aim to offer a brief overview of the literature on deterministic and stochastic dynamic optimization; the focus will be on aerospace applications and, in particular, flight planning.

This chapter is divided in the following sections:

- Section 2.1 describes the main paradigms in dynamic optimization.
- Section 2.2 introduces the main families of dynamic optimization methods employed in flight planning.
- Section 2.3 introduces the formulation of the optimal control problem.
- Section 2.4 reviews the existing methodologies for the solution of deterministic optimal control problems.
- Section 2.5 introduces the mathematical characterization of uncertainty and its integration into differential equations.
- Section 2.6 discusses the solution of optimal control problems under uncertainty.
- Section 2.7 reviews relevant work in the field of aircraft trajectory optimization employing the surveyed methods.

2.1 Dynamic Optimization Paradigms

Practical problems in dynamic optimization face wildly different requirements in different contexts. Some of them have particularly important implications in terms of the choice of methods, such as:

- When can we change the control input of the system?
- What information about the state of the system is available and when?
- How much time can we employ for expensive computations, both before and during operation?
- When do we know the objectives and constraints?

In a fully deterministic context, where the evolution of the system under consideration can be computed exactly for any control input, one might think that these issues are not critical. After all, if a satisfactory solution can be found before the start of the operation of the system, the blind application of the computed control should be enough to lead the system in the desired form.

However, even if uncertainty is not explicitly considered in theory, it is always present in practice, where no model is exact, no variable can be measured with arbitrary precision and no actuator can be activated with arbitrary precision and frequency. Thus, unless the impact of these factors is minimal, it is not enough to obtain an optimal control sequence and apply it; a non-trivial scheme is needed to deal with deviations. We can classify the operational paradigms¹ one might want to consider for this purpose in two broad families: *pre-computed control laws* and *receding horizon control*. Figure 2.1 illustrates these families.

2.1.1 Pre-computed Control Laws

In this concept the *expensive* optimization-related computations are completed *offline before* the system starts to operate. A set of functions or laws that can be evaluated in a relatively fast manner is then obtained and transferred to the controller of the system. Then, once the system starts to operate, these functions or laws are employed by the automatic controller to determine the control inputs at each point in time.

This approach is more suitable when the optimization objective and the dynamical system is known a priori and the computational capacities of the on-board controller

¹Note that, due to the scope of this thesis, we do not consider complications such as adversarial or game-theoretic contexts (i.e. where an intelligent decision-maker or system is working against the objectives of the planner), non-fully-observable state spaces or distributed control problems.

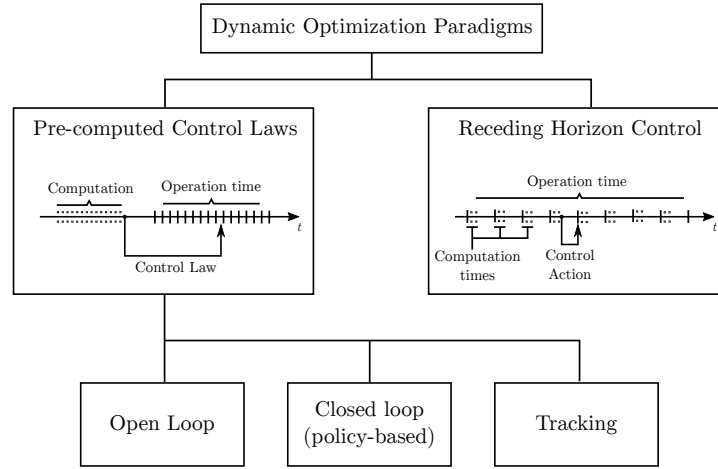


Figure 2.1: Paradigms in dynamic optimization.

are more limited or require certified code. One clear example of this approach are space exploration missions: the trajectory is designed before launch, and followed by the spacecraft afterwards (with a few corrections by the ground team).

Depending on the form of the computed control laws, several subtypes can be identified:

- In an *open loop* approach, the output of the optimization algorithm is a control sequence that, when applied, produces the optimal trajectory of the system. It represents the least robust concept, as any deviation from the nominal trajectory will be amplified sooner or later and thus a tight control of the uncertainty is required in order to avoid large errors.
- In a *tracking* paradigm, the trajectory of the state space (or a subset of the state variables) is obtained from the optimization process. Then, a lower-level, higher frequency and possibly non-optimizing controller is charged with the task of operating the controls in order to lead the system through the desired trajectory. This concept represents a more robust approach, as the disturbances on some of the states can be minimized or eliminated; however, in a high-uncertainty context, these disturbances may be redirected to the untracked states, which might not be desirable.
- In a *closed loop* or *policy-based* paradigm, the output of the optimization algorithm is a function that maps the state space (or the space of observable variables) and the time interval to a control action. At each point in time, the controller uses the measurements from the sensors to feed the control law or policy and compute

the corresponding controls. It represents the most general and flexible paradigm, and it achieves higher performance, but it might demand more computational resources (as the entirety of the state-action space must be explored, or at least a large region) and does not “shield” state variables of interest from uncertainty unless they are explicitly targeted through the objective function or policy shape.

In an ATM context, the most interesting approach for trajectory planning is the *tracking* one. Closely following spatial routes and temporal schedules is desirable in order to forecast conflicts, traffic and airport resource utilization in advance; on the other side, the existence of uncertainty makes open-loop solutions impractical. It is conceivable that, as we move towards higher levels of automation in the future, policy-based paradigms will replace current operational procedures; indeed, research activity under this framework is quickly growing. Nevertheless, while we will offer some thoughts on closed-loop approaches in Section 4.6, this thesis will mainly focus on developing methods under the tracking paradigm.

2.1.2 Receding Horizon Control

Alternatively, the on-board systems might be expected to optimize or re-optimize the trajectory and generate the optimized controls *online*. This might be interesting for multiple reasons: the objective might not be known a priori or might be shifting; the knowledge of the system or the context is limited or nonexistent in advance; the computational resources of the system might be large; the optimization algorithm is reliably fast (for example, if the system is linear and the optimization method only requires a fixed number of easy matrix operations). Thus, *real-time optimization* is carried out in this concept. It is usually performed through a *receding horizon* strategy: only a limited “window” of future time segments is considered to compute the values of the controls for the current time segment; then, the control is applied, the look-ahead window is shifted forward in time and the computations are performed again.

While this approach can be employed for certain airport or flight operations (for example, for adjusting the arrival time to a waypoint or runway), it is unsuitable for general flight planning under the current operational paradigm for the same reasons exposed previously.

2.2 Algorithmic Trajectory Planning

The second question that must be addressed (often in conjunction with paradigm selection) is *how* the operation of the system will be optimized, i.e., what kind of mathematical technique will be employed in order to produce a solution. Figure 2.2 illustrates the most relevant families of methods; we now proceed to overview their main features.

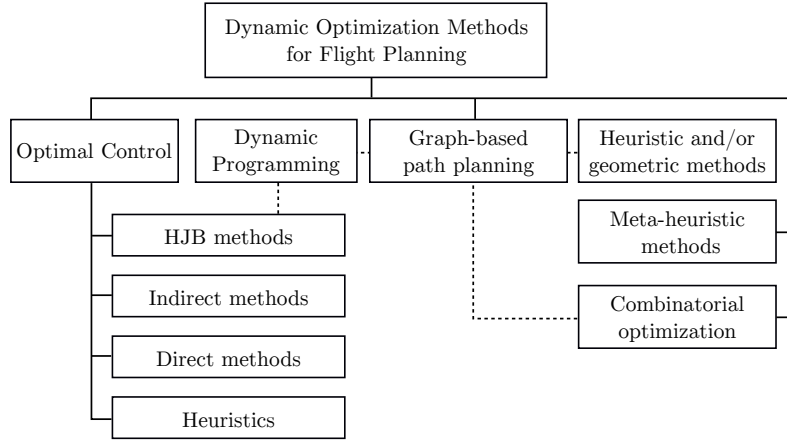


Figure 2.2: Methods in trajectory optimization.

2.2.1 Heuristic and Geometric Methods

These methods represent the simplest kind of flight planning methods, producing tracking-style control laws. An optimization principle is not explicitly employed and flight dynamics are not modeled; instead, heuristic or “greedy” decision rules that are likely to produce a “good enough” result are used. Thus, these methods are fast and reliable, even if they do not produce the best flight plans and trajectories. Most of these heuristics rely on geometric considerations, especially for deconfliction and weather avoidance scenarios.

2.2.2 Dynamic Programming

Dynamic programming produces policy-like control laws. In dynamic programming, the dynamical system is modeled in discrete time. A value function, mapping the state space and the time domain into a scalar field, is computed in backwards recursive fashion from the last time node to the first; then, the optimal action at each stage can be computed by evaluating the value function of subsequent states and aggregating any running costs. The output of a dynamic programming algorithm is then a globally

optimal policy. Its main drawback is the high computational cost for systems with more than a few states: it exhibits the “curse of dimensionality”, an exponential increase in the required computational resources with the dimension of the state.

In order to avoid this unfavourable computational scaling, some techniques have developed and grouped under the banner of *approximate dynamic programming* [52, 53]. These methods approximate one of several of the components of the DP problem. Some of these techniques involve approximating the cost function, discretizing the value function or the policy, limiting the lookahead horizon, aggregating or discretizing states or outcomes may be aggregated or discretized, or reducing the dimensionality of the problem.

2.2.3 Path-Planning Algorithms

Path-planning methods are employed mainly to optimize the route or the horizontal profile of the trajectory; thus, they represent a tracking concept. In these methods, the space is represented as a graph and a path is represented as a sequence of connected nodes and links of this graph that the aircraft will traverse. Each link has an associated cost (for example, the geographical distance between the nodes) and the goal is to find the least-cost traversal path from the beginning to the end node.

Historically, the most popular algorithm is Dijkstra’s shortest path algorithm [54], which employs an approach similar in spirit to dynamic programming. While dynamic programming principles are still the basis of most subsequent path-planning algorithms, the incorporation of heuristics has led to the development of faster algorithms such as A* [55] or D* [56, 57] over the years. These algorithms are useful for structured airspaces, where route-finding problems translate naturally to shortest path problems; however, it is hard to consider time-varying costs or additional performance characteristics (such as fuel burn or varying airspeed profiles) within them.

In non-structured airspaces, it is possible to discretize the continuous position space into a discrete graph and employ the aforementioned algorithms to solve the resulting problem. However, this procedure leads to suboptimal solutions, since the optimal trajectory cannot be arbitrarily approximated, in general, by a sequence of segments from a regular grid. In order to address this difficulty, “any-angle path-planning” algorithms based on A*-like methods have been proposed [58]. In particular, an important development from the robotics field are randomized sampling-based planning methods [59] and, in particular, methods based on Rapidly-exploring random trees (RRT). The introduction of the seminal RRT* algorithm [60] led soon afterwards to the development of RRT* variants seeking to address its shortcomings and incorporate additional features [61].

2.2.4 Combinatorial Optimization

When framing a flight planning problem as a shortest path problem on a graph, an alternative approach to conventional path planning algorithms is to formulate an equivalent mixed-integer optimization problem. The increase in computational capacity and solver performance over the last decades has increased the appeal of these methods. In this setting, the decision variables represent the chosen graph edges, while the continuity of the path can be ensured by a linear constraint. Compared to path planning methods, this approach features similar advantages and drawbacks, with the addition of a higher computational complexity; in exchange, certain advanced features such as the consideration of uncertainty can be incorporated in a more natural fashion [62].

2.2.5 Metaheuristic Methods

If the flight plan is represented as a series of discrete and continuous variables and a fast trajectory prediction algorithm that simulates the flight plan is available, generic gradient-free optimization algorithms can be employed to optimize it. One possible choice to solve the resulting problem are metaheuristic optimization algorithms, which apply combinations of randomized heuristic procedures (often inspired by natural and evolutive processes) in an iterative fashion in order to improve the candidate solution or set of solutions.

Many algorithms have been put forward within this field. Among them, some of the most popular are simulated annealing [63], particle swarm optimization [64, 65], evolution strategies [66], differential evolution [67], genetic algorithms [68], variable neighborhood search [69] and GRASP [70].

While these procedures often do not guarantee convergence to an optimal solution, may require more computational time than alternative methods and often struggle to incorporate constraints, they have several attractive features. Metaheuristic algorithms perform both exploration and exploitation of the search space, making them appropriate for global optimization in problems that may have multiple local minima. They do not require gradient information, and thus they can be integrated with already-existing trajectory prediction software. They are easy to implement, and it is frequently possible to parallelize them. Finally, some of them allow for multiobjective optimization. In consequence, some trajectory planning research has focused on these algorithms, as we discuss in Section 2.7.

2.2.6 Optimal Control

Optimal control methods work in a continuous-time framework and take into account the complete dynamics of the system. Multiple families of numerical methods, covered in Section 2.4 can be employed, with different properties. If the controls from the solution are implemented, optimal control represents an open-loop paradigm; if some states from the solution are extracted and set as reference trajectories for a lower-level controller, optimal control represents a tracking scheme. Finally, optimal control can also be implemented in a *receding horizon control* fashion, often under the label of “Model-Predictive Control” (MPC).

In this thesis, the focus will be on the optimal control paradigm. We will now proceed to provide a short introduction to this field.

2.3 Deterministic Optimal Control

2.3.1 Deterministic Differential Equations

We start from the simplest deterministic differential equation, an uncontrolled Ordinary Differential Equation (ODE) describing a system that evolves or operates on a time interval of interest $\mathcal{T} = [t_0, t_f]$. We assume that the system can be described by an n_x -dimensional vector of state variables $\mathbf{x} : \mathcal{T} \rightarrow \mathbb{R}^{n_x}$:

$$\mathbf{x}(t) = \begin{bmatrix} x_1(t) \\ \vdots \\ x_{n_x}(t) \end{bmatrix} \quad (2.1)$$

In an ODE, the derivative of the state vector at any point in the time interval is given by the dynamical function $\mathbf{f} : \mathbb{R}^{n_x} \times \mathbb{R} \rightarrow \mathbb{R}^{n_x}$:

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), t) \quad (2.2)$$

For clarity of notation, we will make the dependency on time implicit:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, t) \quad (2.3)$$

Certain systems described by an ODE can be influenced by the actions of a user, operator, or automatic control system. This action is described by a control vector of dimension $\mathbf{u} : \mathcal{T} \rightarrow \mathbb{R}^{n_u}$:

$$\mathbf{u}(t) = \begin{bmatrix} u_1(t) \\ \vdots \\ u_{n_u}(t) \end{bmatrix} \in \mathbb{R}^{n_u} \quad (2.4)$$

and the controlled ODE is then described by a dynamic function $\mathbf{f} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R} \rightarrow \mathbb{R}^{n_x}$:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}, t) \quad (2.5)$$

In the systems that we are interested in, we often face multiple constraints on both state space and controls. We will distinguish between equality constraints $\mathbf{h} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R} \rightarrow \mathbb{R}^{n_h}$ and inequality constraints $\mathbf{g} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R} \rightarrow \mathbb{R}^{n_g}$:

$$\mathbf{g}_L \leq \mathbf{g}(\mathbf{x}, \mathbf{u}, t) \leq \mathbf{g}_U \quad (2.6)$$

$$\mathbf{h}(\mathbf{x}, \mathbf{u}, t) = 0 \quad (2.7)$$

where the equality and inequality signs are interpreted in an element-wise fashion. Equality constraints are often paired in optimal control with algebraic variables $\mathbf{z} : \mathcal{T} \rightarrow \mathbb{R}^{n_z}$ (with the number of algebraic variables being at most equal to the number of equality constraints, i.e., $n_z \leq n_h$). These variables do not follow a differential equation; instead, their value is set implicitly by the algebraic constraints \mathbf{h} . With those additions, the ODE becomes a constrained Differential-Algebraic Equation (DAE):

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{z}, \mathbf{u}, t) \quad (2.8)$$

$$\mathbf{h}(\mathbf{x}, \mathbf{z}, \mathbf{u}, t) = 0 \quad (2.9)$$

$$\mathbf{g}_L \leq \mathbf{g}(\mathbf{x}, \mathbf{z}, \mathbf{u}, t) \leq \mathbf{g}_U \quad (2.10)$$

A DAE is more complex to integrate than an ODE; however, they can be dealt with in a very similar manner through direct methods for numerical optimal control. In this work, we will deal with both ODE and DAE models.

2.3.2 Formulation of the Optimal Control Problem

At the conceptual level, the field of control engineering seeks to find a way to generate a control sequence $\mathbf{u}(t)$ (or a closed-loop law/policy $\mathbf{u}_L(\mathbf{x}, t)$ that generates the sequence) such that the resulting trajectory of the system has properties that are desirable for the designer; depending on the type of problem, these properties may include stability, robustness, efficiency, tracking performance or other considerations. In an optimal control setting, the goal is to minimize a *cost functional* (defined in Equation (2.11)), while fulfilling the constraints and initial and final conditions.

$$J(\mathbf{x}, \mathbf{u}, t_f) = \varphi(\mathbf{x}(t_f), t_f) + \int_{t_0}^{t_f} \mathcal{L}(\mathbf{x}, \mathbf{u}, t) dt \quad (2.11)$$

Here, φ is a “terminal cost” term² and \mathcal{L} is called the *Lagrangian term*³. This form

²In a general formulation, the Mayer term can also depend on the initial state and time; this is infrequent in practice but certainly possible: consider, for example, the problem of finding an optimal launch time for a space vehicle

³Note that the dependency in time is again implicit for compactness, i.e. $\mathcal{L}(\mathbf{x}, \mathbf{u}, t) = \mathcal{L}(\mathbf{x}(t), \mathbf{u}(t), t)$

of the cost functional is known as the *Bolza form* of the optimal control. If $\varphi = 0$, then the problem is said to be in *Lagrange form*, and if $\mathcal{L} = 0$ the problem is said to be in *Mayer form*. The terminal cost term represents a cost associated with the final state, while the Lagrange term represents a cost accumulated along the state trajectory.

We denote by \mathbf{x}^* and \mathbf{u}^* the state and controls along the optimal trajectory. Optimality can be defined in a global ($J(\mathbf{x}^*, \mathbf{u}^*, t_f^*) \leq J(\mathbf{x}, \mathbf{u}, t_f)$ for all feasible trajectories (\mathbf{x}, \mathbf{u})) or in a local sense ($J(\mathbf{x}^*, \mathbf{u}^*, t_f^*) \leq J(\mathbf{x}, \mathbf{u}, t_f)$ for all neighbouring trajectories). From the theoretical point of view, the two most relevant mathematical results characterize optimal trajectories through either the *necessary* conditions for optimality (Pontryagin's Minimum Principle) or the *sufficient* conditions for optimality (the Hamilton-Jacobi-Bellman equation).

2.3.3 Pontryagin's Minimum Principle

Pontryagin's Minimum principle (PMP) [71] (alternatively, Pontryagin's Maximum Principle) characterizes the necessary conditions for a trajectory to be an optimum of the OCP. Consider the unconstrained problem:

$$\min J(\mathbf{x}, \mathbf{u}, t_0, t_f) = \varphi(\mathbf{x}(t_f), t_f) + \int_{t_0}^{t_f} \mathcal{L}(\mathbf{x}, \mathbf{u}, t) dt \quad (2.12)$$

subject to:

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}, \mathbf{u}, t) \text{ (dynamic equations)} \quad (2.13)$$

$$\psi(\mathbf{x}(t_0), \mathbf{x}(t_f), t_0, t_f) = 0 \text{ (initial and final conditions)} \quad (2.14)$$

We define the adjoint variables or *costates* $\bar{\lambda}(t) : \mathcal{T} \rightarrow \mathbb{R}^{n_x}$ and the Lagrange multipliers $\nu \in \mathbb{R}^{n_\psi}$ associated to the initial and final conditions. We also define the *Hamiltonian* of the problem as :

$$\mathcal{H}(\mathbf{x}, \bar{\lambda}, \mathbf{u}, t) = \mathcal{L}(\mathbf{x}, \mathbf{u}, t) + \bar{\lambda}^T \mathbf{f}(\mathbf{x}, \mathbf{u}, t) \quad (2.15)$$

The PMP now lists the necessary conditions for a trajectory to be an extremal of the cost functional:

Theorem 2.1 (Pontryagin's Minimum Principle). Let $(\mathbf{u}^*(t), \mathbf{x}^*(t))$ be an optimal trajectory for the problem introduced above. Then, there exist functions $\mathbf{u}(t)$, $\bar{\lambda}^*(t)$ such that:

- The states and costates obey the Euler-Lagrange equations:

$$\frac{d\mathbf{x}^*}{dt} = \frac{\partial \mathcal{H}}{\partial \bar{\lambda}} \quad (2.16)$$

$$\frac{d\bar{\lambda}^*}{dt} = -\frac{\partial \mathcal{H}}{\partial \mathbf{x}} \quad (2.17)$$

- The control minimizes the instantaneous Hamiltonian

$$\mathcal{H}(t, \mathbf{x}^*, \mathbf{u}^*, \bar{\lambda}^*) = \min_{\mathbf{u}(t) \in \mathbb{R}^{n_u}} \mathcal{H}(\mathbf{x}^*, \mathbf{u}, t, \bar{\lambda}^*) \quad \forall t \in [t_0, t_f] \subset \mathcal{T} \quad (2.18)$$

- The final costates satisfy the *transversality conditions*:

$$\bar{\lambda}(t_f) = \left[\frac{\partial \varphi}{\partial \mathbf{x}} + \nu^T \frac{\partial \psi}{\partial \mathbf{x}} \right]_{t=t_f} \quad (2.19)$$

- For problems where the final time t_f is not fixed, the Hamiltonian at the final time also verifies:

$$\left(\frac{\partial \phi}{\partial t} + \mathcal{H} \right) \Big|_{t=t_f} = 0 \quad (2.20)$$

These equations, along with the initial and final conditions, define a Two-Point Boundary Value Problem (2PBVP). The dynamics of the states and costates are coupled; however, where the states are fixed by an initial or final condition, the corresponding transversality condition for the corresponding costate becomes uninformative (as the corresponding Lagrange multiplier ν_j becomes non-zero). Since none of the endpoints of the temporal domain feature a complete set of boundary conditions, it is not possible to solve this problem by simply integrating forwards or backwards, as is usually done for the solution of initial value problem. Therefore, solving a 2PBVP requires a different approach; Section 2.4.2 discusses numerical methods for the solution of this problem.

Constrained problems

The addition of equality constraints (Equation 2.21) and inequality constraints (Equation 2.22) results in an increase in complexity, but similar principles can be applied.

$$\mathbf{h}(\mathbf{x}, \mathbf{u}, t) = 0 \quad (\text{equality constraints}) \quad (2.21)$$

$$\mathbf{g}_L \leq \mathbf{g}(\mathbf{x}, \mathbf{u}, t) = 0 \leq \mathbf{g}_U \quad (\text{inequality constraints}) \quad (2.22)$$

We define the time-dependent multipliers associated to the path constraints $\mu_U : \mathcal{T} \rightarrow \mathbb{R}^{n_g}$, $\mu_L : \mathcal{T} \rightarrow \mathbb{R}^{n_g}$ and $\mu_h : \mathcal{T} \rightarrow \mathbb{R}^{n_h}$, and the augmented Hamiltonian as follows (we have omitted the dependencies on $\mathbf{x}, \bar{\lambda}, \mathbf{u}, t$ for clarity):

$$\mathcal{H}^\circ = \mathcal{L} + \bar{\lambda}^T \mathbf{f} + \mu_L^T (\mathbf{g} - \mathbf{g}_L) + \mu_U^T (\mathbf{g}_U - \mathbf{g}) + \mu_h^T \mathbf{h} \quad (2.23)$$

The PMP conditions are now valid with \mathcal{H}° taking the role of \mathcal{H} and the Lagrange multipliers fulfilling the usual complementary slackness conditions at every t :

$$\mu_L^T (\mathbf{g} - \mathbf{g}_L) = 0 \quad (2.24)$$

$$\mu_U^T (\mathbf{g}_U - \mathbf{g}) = 0 \quad (2.25)$$

$$\mu_L^T \leq 0 \quad (2.26)$$

$$\mu_U^T \leq 0 \quad (2.27)$$

However, inequality conditions generate additional issues, because they may be binding or not at different points in the trajectory. At certain times (called *junction points*), they may switch from non-binding to binding, potentially creating discontinuities in the trajectory of the control and the costates. We refer the reader to [72] for a detailed survey of the theory of optimal control with state and mixed state and control constraints.

2.3.4 The Hamilton-Jacobi-Bellman Equation

The Hamilton-Jacobi-Bellman (HJB) equation describes the sufficient conditions for the trajectory to be an optimum. Consider again the unconstrained problem:

$$\min J(\mathbf{x}, \mathbf{u}, t_f) = \varphi(\mathbf{x}(t_f), t_f) + \int_{t_0}^{t_f} \mathcal{L}(\mathbf{x}, \mathbf{u}, t) dt \quad (2.28)$$

subject to:

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}, \mathbf{u}, t) \text{ (dynamic equations)} \quad (2.29)$$

$$\psi(\mathbf{x}(t_0), \mathbf{x}(t_f), t_0, t_f) = 0 \text{ (initial and final conditions)} \quad (2.30)$$

We define the value function $V : \mathbb{R}^{n_x} \times \mathbb{R} \rightarrow \mathbb{R}$ as:

$$V(\mathbf{x}, t) = \min_{\mathbf{u} \in \mathbb{R}^{n_u}} \left\{ \varphi(\mathbf{x}(t_f), t_f) + \int_t^{t_f} \mathcal{L}(\mathbf{x}, \mathbf{u}, t) dt \right\} \quad (2.31)$$

Theorem 2.2 (Hamilton-Jacobi-Bellman). The optimal control is given by the relationship:

$$\mathbf{u}^*(\mathbf{x}, t) = \arg \min_{\mathbf{u}} (\mathcal{L}(\mathbf{x}, \mathbf{u}, t) + \nabla_{\mathbf{x}} V(\mathbf{x}, t)^T \mathbf{f}(\mathbf{x}, \mathbf{u}, t)) \quad (2.32)$$

where the value function satisfies the Hamilton-Jacobi-Bellman PDE:

$$-\frac{\partial V}{\partial t}(\mathbf{x}, t) = \min_{\mathbf{u}} \{ \mathcal{L}(\mathbf{x}, \mathbf{u}, t) + \nabla_{\mathbf{x}} V(\mathbf{x}, t)^T \mathbf{f}(\mathbf{x}, \mathbf{u}, t) \} \quad (2.33)$$

with the boundary condition:

$$V(\mathbf{x}(t_f), t_f) = \varphi(\mathbf{x}(t_f), t_f) \quad (2.34)$$

The Hamilton-Jacobi-Bellman PDE represents, thus, the continuous-time equivalent of the dynamic programming principle. It produces a control policy that generates the globally optimal trajectory. However, from the practical point of view, there are

multiple challenges involved in the numerical solution of the HJB PDE. The appearance of a minimum operator in the PDE and the lack of smoothness of the solutions for certain problems are some of them, but the most salient one is the *curse of dimensionality*: a state space discretization in a standard tensor-product grid with N points per dimension is composed by N^{n_x} points, which is a prohibitive cost for system where the dimension of the state space is not low.

2.4 Numerical Methods in Deterministic Optimal Control

There are three main families of numerical methods in deterministic optimal control: indirect methods, direct methods and dynamic programming (or HJB-based methods). In this section, we proceed to review their basis, features and practical issues.

2.4.1 HJB methods

As mentioned in Section 2.3.4, one approach that might be employed to obtain the solution of the optimal control problem in a policy-based form is to numerically solve the Hamilton-Jacobi-Bellman PDE. In practice, a conceptually analogous approach that is more often employed in practice is the transformation of the continuous-time system into a discrete-time system, which then generates a dynamic programming problem. In both approaches, the state space is discretized and the value function is obtained by backwards recursion in time (in analogous form as a numerical method for PDEs marching forward in time). There are three main advantages of these methods:

- Ability to find global optima (as all the state space is exhaustively searched).
- Obtention of a closed-loop control rule or policy instead of an open-loop control trajectory.
- Natural extension to uncertain and stochastic problems.

As discussed previously, the main drawback of dynamic programming is its computational cost: the size of the grid required scales exponentially with the dimension of the state space (“curse of dimensionality”). Therefore, memory and operation requirements of standard dynamic programming grow very quickly with the number of state variable and become impractical at 4 or even 3 state variables. The direct application of dynamic programming is therefore limited in practice to problems with low state-space dimensionality.

For specific versions of the HJB equation, efficient methods are sometimes available. For example, *Ordered upwind* methods [73] can solve certain problems in an efficient, one-pass manner that is similar to Dijkstra’s algorithm.

2.4.2 Indirect Methods

Indirect methods solve the Hamiltonian 2PBVP by a numerical scheme. In an indirect method, the necessary conditions for optimality are derived through the application of Pontryagin’s Minimum Principle (“first optimize”) and then the 2PBVP is numerically solved (“then discretize”). As discussed in Section 2.3.3, a 2PBVP is substantially more challenging than an initial value problem, as only part of the boundary conditions for the states and costates are given at the initial time and the remaining ones are given at the final time.

The main techniques in indirect methods are indirect single shooting and indirect multiple shooting [74]. Indirect single shooting is an iterative technique that involves integrating the trajectory forward for an initial guess of the initial costates and then using the gradient of the final states and costates with respect to the initial values to refine them in a manner such that the final values are closer to the target values in each consecutive iteration. Multiple shooting subdivides the time interval into subintervals and then uses an initial guess for the initial values in each subinterval to integrate the ODE or DAE system, with each iteration seeking to reduce the distance between the final, computed values in each subinterval and the initial values of the next subinterval (as well as the final values in the last subinterval) with each iteration.

Indirect methods have considerable drawbacks [75] for certain application:

- The optimality conditions must be derived, which requires nontrivial amounts of work and expertise with optimal control methods. In particular, any change in the model of the system or the problem formulation requires re-deriving these conditions from scratch.
- The “switching structure” (the sequence of active path constraints and regular, singular and saturated control arcs) must be derived or known a priori, which may be a hard task for complex problems.
- An initial guess for the trajectory must be provided. While this is common to most numerical methods, indirect methods also require an initial guess for the costates. Since there are usually no available “intuitive” values for the costates (while state variables can usually be initialized at an actual trajectory), this is a challenging task. Indirect methods are also highly sensitive to this initial guess and ill-conditioned.

- Dealing with “empirical” models (for example, if a function is approximated by interpolating tabular data) can be a rather difficult or impossible task.

These difficulties limit the applicability of indirect methods to certain problems.

2.4.3 Direct Methods

Direct methods operate in reverse compared with indirect methods: the problem is “discretized first” to a nonlinear programming problem (NLP) and then solved by an NLP algorithm (“then optimize”). Their simplicity, as well as the ability to avoid the complications that arise in alternative methods, have bolstered the adoption of direct methods; in exchange, their theoretical foundation is usually weaker, they generally present less accuracy than indirect methods, and proofs of convergence to a solution of the continuous problem are limited for most methods. The NLPs produced by direct methods are usually solved by gradient-based optimization methods, such as the Sequential Quadratic Programming algorithm SNOPT [76] or the interior point algorithm IPOPT [77].

There are two main families of direct methods: direct shooting and direct transcription.

Direct shooting

Single shooting and multiple shooting have a direct version [75, 78]. In the direct single shooting method, the control history is discretized as a set of NLP variables⁴ and the state is propagated using numerical integration. The cost function can then be computed from the state and control history: it becomes the objective function of the NLP. Direct multiple shooting methods divide the temporal domain into several subintervals and act on a similar fashion, with the states at the edges of the subintervals becoming additional decision variables. By increasing the size of the NLP, multiple shooting improves reliability, because it reduces the high sensitivity of the single shooting method to perturbations in the initial conditions and improves the conditioning of the Jacobian of the problem.

Direct transcription

If the time domain is subdivided into subintervals that match the integration steps, the resulting method is called a transcription method. In direct transcription methods, the values of the state variables at each time subinterval become decision variables together with the control values. The differential equations are transformed into discrete *defect constraints*, which relate the values at the beginning of the subinterval to the

⁴Thus, direct shooting methods are sometimes referred to as “control parametrization” schemes.

values at the end. Different methods are characterized by the choice of quadrature rule to approximate the differential equation between each two subintervals. For example, consider the Euler method in numerical integration of ODEs:

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \Delta t \cdot \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, t_k) \quad (2.35)$$

where Δt is the step size ($\Delta t = t_{k+1} - t_k$). In an Euler-based transcription method, this exact equation is enforced as a constraint $\zeta_k = 0$ on variables \mathbf{x}_k , \mathbf{u}_k and \mathbf{x}_{k+1} for each subinterval k .

$$\zeta_k = \mathbf{x}_{k+1} - \mathbf{x}_k - \Delta t \cdot \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, t_k) \quad (2.36)$$

The most popular transcription methods are collocation methods, where the evolution of the state inside each subinterval is represented by a polynomial whose derivatives are equated to the values of the dynamic equation \mathbf{f} at a set of interior nodes within each subinterval. The Euler method represents the simplest of the collocation methods.

Collocation methods can be classified into two families. In local collocation methods, the degree of the interpolating polynomial is fixed and accuracy is improved by increasing the number of subintervals (thus reducing the step-size). Higher-order analogues of the Euler method, such as the Hermite-Simpson [79] and the Gauss-Lobatto collocation schemes [80], are the most commonly used among local collocation methods. From the computational point of view, an important advantage of these methods (though not unique to them) is the high sparsity of the Jacobian of the Lagrangian. Because each constraint only relates two adjacent subintervals, only the entries near the diagonal of this matrix will be populated. Certain NLP algorithms are thus able to take advantage of this sparsity to speed up computation and obtain good performance.

Alternatively, global collocation or pseudospectral methods have emerged as a powerful alternative to local collocation methods [81, 82]. In a pseudospectral method, the entire state trajectory may be represented by a high-order interpolating polynomial in a single subinterval and accuracy is improved by increasing the degree of the polynomial. The differential equations are then collocated at a set of collocation points determined with a Chebyshev-Gauss or Legendre-Gauss quadrature rule. Inspired by spectral collocation methods in fluid dynamics, pseudospectral methods in optimal control started being studied in the late 1990s. Key early milestones are the formulation of the Legendre Pseudospectral Method (LPM), based on Legendre-Gauss-Lobatto quadrature [83], and the Chebyshev method [84], later reworked on Clenshaw-Curtis quadrature in [85]. Later developments include the Gauss Pseudospectral Method (GPM), based on Legendre-Gauss quadrature [86], and the Radau Pseudospectral Method (RPM) [87, 88], based on Legendre-Gauss-Radau quadrature.

The main advantage of pseudospectral methods compared to usual direct collocation methods is the superior rate of convergence for problems that are smooth enough [86].

While local methods feature polynomial convergence rates on the number of nodes, pseudospectral methods can achieve *spectral accuracy* [89]: a rate of convergence higher than any polynomial rate and, in particular, exponential for analytic functions. Additionally, there is a stronger theoretical understanding of these methods when compared to other direct methods [86, 88, 90, 91, 92]. It is worth noting that these methods can also be applied in a local collocation fashion, fixing the degree of the approximating polynomial to a low number and increasing the number of subintervals instead of the degree of the polynomial [87].

Nevertheless, pseudospectral methods are not entirely without drawbacks:

- The individual sections of the Jacobian of the Lagrangian (corresponding to the differentiation matrix and the collocation constraints for each variable) are no longer sparse, but dense. This can impose significant computational burdens once the number of nodes becomes high enough.
- Pseudospectral methods do not handle non-smooth problems well. Because the solution is represented as a global polynomial, the existence of points where there is a discontinuity in one of derivatives of the state deteriorates the properties of these methods. This phenomenon is usually encountered in practice because of the activation of path constraints and the appropriate way to deal with it is to reformulate the problem as a multiphase problem [14, 93] or to employ an adaptive method.

Adaptive methods in direct transcription

A large number of practical optimal control problems feature optimal trajectories that may display features at different timescales. For example, in a simple optimal cruise problem, the optimal solution is often composed by an initial phase at either maximum or minimum throttle until an efficient velocity is reached, a central “singular arc” phase where the aircraft slowly decelerates as the optimal velocity is adjusted to the changing mass of the aircraft, and a final acceleration or deceleration phase to match the terminal velocity (or the minimum allowed velocity, if no final condition is imposed on the velocity) [94, 95]. This presents a challenge for most direct methods, since choosing a large step size that is efficient for the long central phase leads to inaccuracies in the initial and final phases, and choosing a small step size for accurate representation of the solution at the endpoints is computationally uneconomical for the central singular arc phase, leading to an oversized NLP.

To alleviate this phenomenon, the numerical optimal control community has proposed the usage of adaptive methods. These algorithms are built on top of other direct methods and follow this general scheme:

Algorithm 1: Main loop of an adaptive method in optimal control

```

1 Select a basic initial grid with an initial step size
2 Solve the problem on the selected grid
3 if solution has the demanded precision then
4   | Accept the current solution and terminate the algorithm
5 else
6   | Proceed to step 7
7 Use an error criterion to determine which parts of the grid require refinement
8 Design the new grid according to the criterion and go to step 2

```

Adaptive algorithms differ on the underlying method employed, on the choice of error or smoothness criterion and the grid-refinement rule. Some relevant adaptive algorithms are [96, 97, 98, 99], for local collocation methods, and [100, 101] for pseudospectral methods, which are particularly sensitive to non-smooth trajectories such as trajectories with changes in the active constraint set.

A related technique that can be employed in place of the error criterion for grid design is the subdivision of the problem according to constrained and unconstrained arcs. This was first considered by in the context of low-order direct methods [102, 103], and later on for pseudospectral methods [14, 104].

2.5 Uncertainty and Differential Equations

2.5.1 Random Variables

Uncertain parameters are mostly modeled within the framework of mathematical probability or within the framework of interval arithmetic. While the latter approach has been successfully employed for some problems within the field of robust control theory [105, 106], it is the former that is predominantly used for this purpose.

The mathematical pillar of probability theory in the common Kolmogorov formulation is the *probability space*. A probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is composed by three elements:

- An abstract sample space Ω of possible outcomes. Every possible realization of the uncertain parameters is a single point $\omega \in \Omega$ in this space.
- A σ -algebra of events \mathcal{F} . This is a collection of subsets of Ω that are called “events”, that is, groups of outcomes. If an outcome $\omega \in \Omega$ is contained within an event $A \in \mathcal{F}$, the event has taken place in this outcome.

- A probability measure \mathbb{P} that assigns a probability $\mathbb{P}(A) \in [0, 1]$ to each event $A \in \mathcal{F}$. It is normalized so that $\mathbb{P}(\Omega) = 1$.

A random variable is an \mathcal{F} -measurable function $X : \Omega \rightarrow \mathbb{R}^n$. Its distribution μ_X can then be defined as a probability measure on \mathbb{R}^n instead of the underlying abstract probability space by the relationship $\mu_X(B) = \mathbb{P}(X^{-1}(B))$, with $B \in \mathbb{R}^n$.

Uncertainty can be included in differential equations in different forms. Two main types of differential equations under uncertainty can be identified: the Randomly-parametrized Differential Equation (RDE) and the Stochastic Differential Equation (SDE), depending on whether the uncertainty is assumed to be described by random but static parameters or by randomly-varying dynamical variables.

2.5.2 Randomly-parametrized Differential Equations

In an RDE, we consider a differential equation influenced by a set of uncertain constant parameters whose probabilistic distribution is known. This class of systems are called *tychastic* dynamical systems in [51] and related work. The uncertain parameters of the system are modeled as a continuous random variable $\xi : \Omega \rightarrow \mathbb{R}^{n_\xi}$ that we assume to be *constant in time*. For each possible outcome $\omega \in \Omega$, the random variables take a different value $\xi(\omega)$.

Therefore, for each outcome $\omega_0 \in \Omega$, there exists a unique trajectory path $t \rightarrow (\mathbf{x}(\omega_0, t), \mathbf{z}(\omega_0, t), \mathbf{u}(\omega_0, t))$ that corresponds to the realization of the random variables $\xi(\omega_0)$. The dynamics of the system are given by the functions $\mathbf{f} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_z} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_\xi} \times \mathbb{R} \rightarrow \mathbb{R}^{n_x}$, $\mathbf{h} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_z} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_\xi} \times \mathbb{R} \rightarrow \mathbb{R}^{n_h}$, and $\mathbf{g} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_z} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_\xi} \times \mathbb{R} \rightarrow \mathbb{R}^{n_g}$, such that valid trajectories fulfill the conditions almost surely (i.e. with probability 1):

$$\frac{d}{dt}\mathbf{x}(\omega, t) = \mathbf{f}(\mathbf{x}(\omega, t), \mathbf{z}(\omega, t), \mathbf{u}(\omega, t), \xi(\omega), t) \quad (2.37)$$

$$\mathbf{h}(\mathbf{x}(\omega, t), \mathbf{z}(\omega, t), \mathbf{u}(\omega, t), \xi(\omega), t) = 0 \quad (2.38)$$

$$\mathbf{g}_L \leq \mathbf{g}(\mathbf{x}(\omega, t), \mathbf{z}(\omega, t), \mathbf{u}(\omega, t), \xi(\omega), t) \leq \mathbf{g}_U \quad (2.39)$$

where the \leq sign applies in an element-wise fashion in Equation (2.39) and analogous equations, $\mathbf{g}_U - \mathbf{g}_L > 0$ elementwise and $n_z \leq n_h$. Thus, given a policy or open-loop control sequence that determines the value of the controls, the trajectory will follow the deterministic differential-algebraic equations (2.37) and (2.38) for each possible realization of the random variable $\xi(\omega)$ (as long as the constraints (2.39) allow). We have employed the notation $\mathbf{x}(\omega, t)$ to emphasize the dependence of the trajectories on the random variables; henceforth, we may again use the abbreviated notation \mathbf{x} when working with RDEs for improved clarity.

2.5.3 Stochastic Differential Equations

Within the framework of stochastic differential equations, the state variables are not seen as deterministic functions from \mathcal{T} to \mathbb{R}^{n_x} , but instead as a collection of random variables on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$ which are parametrized by time:

$$\{X_t\}_{t \in \mathcal{T}} \quad (2.40)$$

That is, for each $t \in \mathcal{T}$ there is a random variable

$$\omega \rightarrow X_t(\omega), \omega \in \Omega \quad (2.41)$$

and, conversely, for each outcome ω there is a *path* of the process defined as the function

$$t \rightarrow X_t(\omega), t \in \mathcal{T} \quad (2.42)$$

The basic building block of an SDE is the *Wiener process* W_t (often called *Brownian motion* as well). Such a process is characterized by the following properties:

- $W_0 = 0$ (within this section, we will assume that $t_0 = 0$)
- The increments of W are independent, i.e. $W_s - W_t$ (with $t < s$) is independent of the history of W before t .
- The increments of W are standard Gaussian variables with variance equal to the time differential, i.e., $W_s - W_t$ (with $t < s$) has a probability distribution function of the form:

$$p(y) = (2\pi|s - t|)^{-1/2} \exp\left(-\frac{y^2}{2|s - t|}\right) \quad (2.43)$$

- The paths of W_t are almost surely continuous.

Thus, the Wiener process can be viewed as the continuous time limit of a Gaussian process. Similarly, we can define the n -dimensional Wiener process \mathbf{W}_t as a vector whose components are independent 1-dimensional Wiener process. The probability distribution function of $\mathbf{W}_s - \mathbf{W}_t$, given the history of the process up to t , is given by:

$$p(\mathbf{y}) = (2\pi|s - t|)^{-n/2} \exp\left(-\frac{\|\mathbf{y}\|^2}{2|s - t|}\right) \quad (2.44)$$

An uncontrolled SDE (the only kind that will be considered in Section 3.3.2) in the Itô formalism is written in the form:

$$d\mathbf{X}_t = \mathbf{B}(t, \mathbf{X}_t) + \Sigma(t, \mathbf{X}_t)d\mathbf{W}_t \quad (2.45)$$

where $\mathbf{B} : \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ represents the “deterministic” component of the dynamics while $\Sigma : \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ represents the amplitude and shape of the random perturbations at any point in time and the state space. We say that the stochastic process \mathbf{X}_t is a solution of Equation 2.45 if it fulfills the integral form of the equation, namely:

$$\mathbf{X}_t = \mathbf{X}_0 + \int_0^t \mathbf{B}(s, \mathbf{X}_s) ds + \int_0^t \Sigma(s, \mathbf{X}_s) d\mathbf{W}_s \quad (2.46)$$

where the second integral is an Itô integral (see Chapters 3-5 of [107]).

The mathematical characteristics of stochastic differential equations introduce new challenges from the point of view of numerical approximation [108], as the numerical methods for integrating ordinary differential equations do not necessarily extend automatically to the stochastic version.

As the “solution” of a SDE initial value problem is not a single trajectory but a random trajectory that depends on the underlying source of randomness, the usual approach to solve SDEs numerically is to simulate a large number of individual paths of the solution in a Monte Carlo fashion and then obtain statistics such as the mean or the variance from this solution. The individual realizations of the solution are simulated using numerical schemes such as the Euler-Maruyama method [109], the Milstein method [110] or other Taylor or Runge-Kutta⁵ approximations [111].

2.5.4 Discretization of Random Variables

Given a random variable X with a probability measure μ_X , most of the magnitudes of interest related to X can be expressed or computed as integrals of the form:

$$\mathbb{E}[f(X)] = \int_{\mathbb{R}^n} f(x) d\mu_X(x) \quad (2.47)$$

for some function f . For example:

- The expected value of X is $\mathbb{E}[X]$
- The variance of X is $\text{Var}(X) = \mathbb{E}[(X - \mathbb{E}[X])^2]$
- The probability that $X \in A$ (where $A \subset \mathbb{R}^n$) is $\mathbb{E}[I_A(X)]$, where $I_A(x)$ is the indicator function of the set A . This case includes quantities such as “the probability that certain variable is equal or greater than a critical value”.

Therefore, from the practical point of view, dealing with random variables often involves solving or approximating integrals of the form (2.47), often by transforming them to discrete random variables. There are multiple approaches for carrying out this task; we summarize in this section some of the most important.

⁵Despite the name, stochastic Runge-Kutta schemes (as named by [111]) are not just heuristic adaptations of the deterministic RK methods.

Monte Carlo

Monte Carlo methods and its derivatives are widely used in science and engineering. The basic Monte Carlo scheme approximates the integral 2.47 as:

$$\int_{\mathbb{R}^n} f(x) d\mu_X(x) \approx \frac{1}{N} \sum_{i=1}^N f(\xi_i) \quad (2.48)$$

where the values ξ_i have been randomly generated according to the distribution $d\mu_X$.

The main advantage of this method lies in its simplicity of implementation. Its main drawback lies in the fact that the convergence rate is slow. Under mild assumptions on f , the error converge at an $\mathcal{O}(1/\sqrt{N})$ rate. Therefore, a high N might be required to provide an accurate result. If the evaluation of f is computationally expensive (for example, if it requires running a costly simulation), then it might be infeasible to use a Monte Carlo approach. On the other hand, this convergence rate is independent of the dimension of the random variable X , thus making it relatively more efficient for high-dimensional problems than more sophisticated alternatives.

Depending on the complexity of the distribution of X , several variants of Monte Carlo can be employed to improve its performance. Examples of these techniques are importance sampling, stratified sampling and the Markov Chain Monte Carlo (MCMC) family of algorithms [112].

Quasi-Monte Carlo

Quasi-Monte Carlo methods work in a similar fashion to standard Monte-Carlo. However, instead of generating the evaluation points $\{\xi_i\}$ randomly, quasi-MC methods use a deterministic low-discrepancy sequence that samples the outcome space in a more even manner than raw Monte-Carlo [113]. While the upper bound for the approximation error incurred by quasi-MC, $\mathcal{O}\left(\frac{(\log N)^n}{N}\right)$, can be worse than standard Monte Carlo, empirical results from the quantitative finance field (starting with [114]) showed that quasi-MC often converges 1) substantially faster than what the bound seems to indicate and 2) at a rate of $\mathcal{O}(1/N)$, asymptotically faster than Monte Carlo. This discrepancy has been theoretically studied with the concepts of *weighted classes of functions* [115] and *effective dimension* [116]; these ideas rely on the fact that, for certain problems, not every dimension of the random is equally influential in the quantities of interest. Nevertheless, there is still no complete understanding of the phenomenon, as researchers have identified classes of functions with high effective dimension for which quasi-Monte Carlo still outperforms Monte Carlo with a $\mathcal{O}(1/N)$ convergence rate [117].

Generalized Polynomial Chaos

Generalized Polynomial Chaos (gPC) techniques rely on the expansion of the random inputs and outputs on an orthogonal polynomial basis [118, 119]. This allows for recovery of some statistical quantities directly from the expansion coefficients (for example, the expected value being the first basis coefficient). They come in two forms:

- In the *Galerkin* form, the variable X is replaced by its gPC expansion in the formulation of the problem, thus creating a problem on the coefficients of the expansion. The new problem is then solved to obtain the gPC characterization of the output.
- In the *stochastic collocation* form, the deterministic problem is solved for different values of X . These values are chosen as the nodes of a quadrature rule. The output is then characterized as a Lagrange interpolant at this set of nodes (therefore, stochastic collocation can be seen as a form of building surrogate models). The gPC coefficients can then be recovered with the quadrature rule.

Both variants of gPC have their own strengths and weaknesses. The Galerkin approach is generally more accurate, but it requires the modification of the underlying simulator of f (it is *intrusive*), which might be a difficult or unfeasible task for some problems. On the other hand, stochastic collocation only requires the deterministic evaluation of f at a pre-computed set of nodes, building the output approximation in a cheap post-processing step; it can also be integrated with kriging techniques [120].

Compared to Monte Carlo, gPC has been shown to perform favorably for problems where the random variables are smooth and low-dimensional [121]. This advantage vanishes for high-dimensional problems, because the number of expansion coefficients or function evaluations needed for gPC grows quickly with the number of dimensions n . A full tensor product grid will scale as $\mathcal{O}(e^n)$; a more efficient sparse grid [122] will still scale at a high polynomial rate $\mathcal{O}(n^l)$, where l is the level of the quadrature rule. This effect limits the accuracy of the gPC expansion for high-dimensional problems, as increasing the level of the quadrature beyond a low value will quickly make the method computationally unfeasible. Another drawback of polynomial chaos is the fact that estimates for the approximation error are usually not available [123].

Cubature

Cubature is a numerical method for the approximation of n -dimensional integrals, a high-dimensional analog of quadrature rules for one-dimensional integrals. Given a region $B_n \subset \mathbb{R}^n$, a weight function $w(x)$ (in our case, the probability density function) and a function $f(x)$, a cubature rule approximates the integral as:

$$\int_{B_n} w(x)f(x)dV = I_N f + R_N f$$

where $I_N f$ is computed as a linear combination of values of the function at the cubature points:

$$I_N f = \sum_{i=1}^N w_i f(x_i)$$

and $R_N f$ is the residual or error of the cubature, which is usually demanded to be zero for linear combinations of monomials [124]. There are numerous available cubature rules that differ in the region of computation B_n (usually n -dimensional cubes, spheres, simplexes or the whole space) and in the choice of cubature points $\{x_i\}$ and weights $\{w_i\}$. A compilation of cubature rules can be found in [125] and [126]. For stochastic processes (see Section 2.5.3), specific high-dimensional cubature has also been developed [127]. Finally, some cubature concepts have been put forward for specific robust optimal control applications [51].

2.6 Dynamic Optimization under Uncertainty

Uncertainty can be integrated into the optimal control problem in several forms, depending on:

- The nature of the modeled random variable (static or dynamic)
- The component of the problem that depends on the uncertainty (the dynamics, the cost functional or the constraints).
- The dynamic optimization paradigm.

We will group the main approaches employed in this field in four families: dynamic programming, stochastic optimal control, the “uncertainty quantification” optimal control approach, and the robust optimal control approach.

2.6.1 Dynamic Programming

Dynamic programming can be extended to uncertain problems, as it is a policy-based approach that deals with uncertainty in a natural form. However, the challenges related to the curse of dimensionality are worsened by the consideration of the dimensionality of the random parameters, and thus sub-optimal approximate dynamic programming techniques need to be applied in practice to produce a solution with limited computational resources [53, Chapter 6].

One particular framework that is often employed in uncertain contexts is the Markov Decision Process (MDP). In an MDP, a finite set of actions is available at each of the finite states of the system; the action taken determines the transition of the state at each stage in a probabilistic manner. The choice of a policy, i.e., a function selecting an action for each state of a system, reduces the system to a Markov chain. The goal is, then, to find the policy that maximizes the desired performance index.

Depending on the form of the approximation techniques, the methods for solving the resulting problem (often grouped under the label of *reinforcement learning*) can be classified in multiple families. Traditionally, the value function (assigning scores to states) or the Q-factors (which maps state-action pairs to scores instead) were approximated, leading to methods such as Q-learning [128] or Temporal Differences (TD) [129]. More recently, techniques based on the optimization of parametrized policies have become more popular, as they can be applied in model-free contexts; however, this generality comes at a significant cost, as these methods are often significantly less efficient in theory and practice than model-based approaches [130].

2.6.2 Stochastic Optimal Control

Stochastic optimal control is a policy-based dynamic optimization paradigm. In a Stochastic Optimal Control Problem (SOCP), uncertainty is represented by a dynamic stochastic process and the system dynamics are modeled with the aid of a continuous-time SDE. Thus, stochastic optimal control is the continuous-time equivalent of dynamic programming. Using an explicit closed-loop feedback law for the controls (in the form $\mathbf{u} = \mathbf{u}(t, \mathbf{X}_t)$) has several advantages, since the implementation and simulation of this controller are straightforward, computationally cheap processes (which is desirable for real-time controllers).

The linear-quadratic version of the problem (linear dynamics and quadratic costs) has been explored and used extensively in practice [131], just like its deterministic analogue. However, the general nonlinear problem SOCP can exhibit a more complex behaviour than the deterministic OCP due to its stochastic nature [132]. Stochastic analogues to the deterministic Pontryagin's Maximum Principle [133] and Hamilton-Jacobi-Bellman equations [134] have been explored in the literature.

Numerical methods in stochastic optimal control are more challenging than those in the deterministic case. The traditional approach for solving a SOCP is to discretize the state space and convert the problem into a discrete Markov Decision Process [135], which is then solved by dynamic programming techniques. The main drawback of this methodology is the fact that it suffers from the curse of dimensionality, as the computational size of the discretized state space grows exponentially on the dimension of the problem.

Some ideas have been introduced in the literature in order to avoid the curse of dimensionality in the context of stochastic optimal control. For certain problems, the Hamilton-Jacobi-Bellman equation can be transformed to a linear PDE [136], which can then be solved in an efficient manner. The path integral approach [137, 138, 139, 140] draws on ideas from physics and approximate dynamic programming in order to produce a reinforcement learning algorithm that scales to high dimensional control systems. The algorithm presented in [141] is able to locally solve an approximation to the SOCP and scales linearly on the number of dimensions, in addition to being parallelizable.

In [142], the state and control trajectories are expanded in a Fourier-like fashion in an orthonormal polynomial basis spanning the underlying *Wiener chaos space* using Malliavin calculus. The SOCP is therefore converted to a deterministic OCP on the coefficient functions of the expansion, instead of the states. This OCP can now be solved using mature deterministic optimal methods and the stochastic solution can be recovered from the solution coefficients.

2.6.3 Uncertainty Quantification Optimal Control

This approach combines a non-intrusive Uncertainty Quantification (UQ) methodology with a solver of the deterministic OCP. Under this approach, the deterministic problem is solved for different values of the uncertain parameters and the statistics of the solution are built using the corresponding probabilistic representation (usually from a gPC stochastic collocation scheme). When the uncertain variables are realized, this probabilistic representation can be used to obtain the deterministic solution corresponding to the final values of the uncertain variables.

This approach is useful to study the probabilistic characteristics of the deterministic problem when some parameters are unknown. It is also a convenient computational solution for systems where computation is done before the system starts operating and uncertainty is solved before the startup point: the solution of the slow and computationally expensive OCP is stored after calculation and the solution is built in real-time with a fast interpolation operation.

This methodology is, nevertheless, not applicable for problems where uncertainty is not completely solved before the system starts operating. In that case, there is no clear interpretation of the solution, as the probabilistic representation of the trajectory does not provide a rule to choose the control inputs. The most likely trajectory in this solution is not only not optimal in general (as optimality does not, in general, commute with the “most likely” property⁶): it is almost surely unfeasible.

⁶As an illustrative counterexample, consider the problem of finding u such that $(x - u)^2$ is minimum, where x is a random variable $x \in [-1, 1]$ with a probability density function $p(x) = (1 + \varepsilon x)/2$ and $0 < \varepsilon < 1$. It can be verified that the most likely solution of the deterministic problem is $u = 1$, but the solution that minimizes the average value of $(x - u)^2$ is $\varepsilon/3$, which is closer to 0

2.6.4 Robust Optimal Control

In the robust (or *tychastic*, using the terminology of [51]) optimal control approach, the dynamics of the problem are modeled with an RDE and an open-loop precomputed control law is employed. In robust optimal control, the fact that the uncertainty has not been resolved yet at the point of execution is explicitly taken into account. Instead of a single execution trajectory (as in deterministic optimal control), the solution of a robust OCP is a set or “tube” of trajectories in state-space that minimizes the expected value of a cost functional that may not only include an “average” cost, but can also include a term penalizing the variance or dispersion of the solution. The deterministic dynamical system is replaced by an augmented “virtual” dynamical system that represents the dynamics of all the trajectories, so it is converted to a larger deterministic optimal control problem.

In the particular case where the uncertainty only influences the cost functional and not the dynamics, the trajectory tube collapses to a single trajectory and thus it is not necessary to expand the state space. A simpler strategy for solving the problem with a direct method is replacing the cost functional by an approximation to the probabilistic cost functional [143].

While this approach is useful for some engineering problems, its reliance on an open-loop paradigm prevents it from being applied widely; as discussed in Section 2.1, open-loop concepts are not practical for a lot of problems (and, in particular, flight planning). In order to apply this concept successfully in more general settings, it is necessary to extend and generalize this robust optimal control approach to the tracking and policy-based paradigms. We will perform this task in Chapter 4.

2.7 Flight Planning Applications

We end the chapter with a brief survey of some examples of application of the described methodologies for dynamic optimization in the context of flight planning.

2.7.1 Deterministic Flight Planning

Flight planning can be defined as the procedure of generating a valid flight plan with the goal of minimizing fuel burn, flight time and overflight charges while taking into account weather, traffic, and ATC constraints. In order to deal with the regulation and structure of the current airspace structure, commercial flight planning tools usually subdivide the problem into a 2D route optimization problem and a vertical profile and speed optimization procedure [144].

However, initiatives such as SESAR [145] and NextGEN [146] are promoting the implementation of free route airspaces, an intermediate step towards the 4D trajectory

concept. This move towards Trajectory-Based Operations is renewing interest in optimal control techniques that can solve complete 4D trajectory optimization problems. This trend has also been fostered by the growing adoption of Continuous Climb Operations (CCO) [147] and Continuous Descent Operations (CDO) in major regions. Indeed, recent surveys [47, 148] find that optimal control currently represents the most popular framework for the solution of this class of problems.

The usage of analytical optimal control techniques to optimize fuel savings started several decades ago (see [149, 150] and references therein for some early works on vertical profile and speed optimization and [151] on three-dimensional trajectories). More recently, the work in [152] used Green's theorem to obtain the solution of the constant altitude cruise problem; indirect methods have also been employed to obtain the minimum fuel cruise with a fixed time of arrival [94]. Complete profiles including climb, cruise, and descent were studied in [153] and extended for altitude-constrained profiles [154]. An optimization procedure based on discrete patterns was introduced in [155] and extended by [156] in order to generate ATC-compliant trajectories composed by segments of constant Mach and altitude.

Several researchers have also studied the aircraft trajectory optimization problem with direct methods. Efficient and reliable landing procedures using optimal control are developed in [95]. As the performance of direct methods is highly dependent on the initial guess, a method for generating initial guess trajectories for the same problem is introduced in [157]. In [158, 159], a complete trajectory is optimized using hybrid optimal control and Mixed-Integer Nonlinear Programming (MINLP), and a similar approach is employed in [160] in order to include contrail avoidance in the objective function. The flexibility of direct methods enables advanced concepts proposed in the literature, such as linear holding [161].

Meta-heuristic methods have been also employed in the literature to solve trajectory optimization problems. Some examples are bee colony [162] and ant colony [163] algorithms, simulated annealing [164], and beam search [165]. Genetic algorithms are also studied for trajectory optimization [166, 167] as well as multi-aircraft deconfliction [168, 169].

2.7.2 Wind-optimal Aircraft Trajectories

While wind and weather have an important influence on trajectory efficiency, a sizeable amount of works on this topic (including some of the studies cited in Section 2.7.1) do not take it into account. Nevertheless, the problem of optimal aircraft routing in general wind fields has indeed been specifically studied in the optimal control literature. The earliest works trace their roots to the “Zermelo problem” [170]. In this setting, a vehicle (in the original formulation, a boat) moves with constant relative speed with respect to an airmass, whose movement is described as a constant or time-varying wind field; the optimal heading can then be obtained as the solution of a differential equation by employing the optimal control framework.

One of the first modern studies on the topic was [171], which relied on the techniques of Neighbouring Optimal Control (NOC). A guidance method in a general wind field is developed, as well as a NOC-based algorithm, which solves a linearized system describing a perturbation around a nominal trajectory. In [172], this method (called Near-Optimal Wind Routing) is compared with an “Optimal Wind Routing” solution, which integrates a differential equation (obtained by analytical optimal control techniques) from the destination point with different values of a parameter in order to find a route that passes through the initial point. Additional work on this class of techniques has been carried out in [173], which considers time-varying wind fields and variable airspeed, and another analytical optimal control-based method is applied in [174] to find optimal routes in the presence of winds while avoiding contrail formation. This method is combined with a vertical profile optimization method in [175].

Methods based on dynamic programming have also been employed for the calculation of wind-optimal paths. In [176], an efficient dynamic programming approach relying on an ordered upwind method is used to optimize the 2D trajectory of a constant-velocity aircraft. This method is extended in [177] to consider the flight on the surface of a sphere in order to adapt it to medium- and long-haul flights. Dijkstra’s algorithm is also employed for route-optimization problems in some works [178, 179].

In contrast with these works, which apply to general wind fields but usually make other simplifications (such as constant airspeed or altitude), the studies in [180, 181, 182] study the effect of *average* along-track wind speeds while considering variable speed, with a formulation based on indirect methods. In [153], hybrid multiphase optimal control is used to generate complete climb, cruise and descent profiles with the ability to consider general wind profiles.

Direct methods, which have the potential of dealing with more complex and complete problems as discussed in Section 2.4.3, have also been used to optimize trajectories on general wind fields. In [158], [183] and [159], the wind is modeled as a polynomial obtained by regression on tabular data and hybrid optimal control is used in conjunction

with a direct transcription to solve the optimal control problem; the wind approximation scheme is improved in [13] with a multiphase formulation.

Other methods on trajectory optimization on general wind fields without using optimal control were developed in [184, 185], which use a genetic algorithm and a performance database to obtain optimal 4D trajectories.

2.7.3 Aircraft Flight Planning under Uncertainty

Recently, interest has been growing on the quantification of the impact of uncertainty on all aspects of the ATM system. In particular, as mentioned in Chapter 1, the SESAR-sponsored IMET project was set up with the goal of utilizing ensemble meteorological data to study the impact of meteorological uncertainty and improve decision-making in ATM and pre-tactical flight planning operations. Some results were presented in [40], where the impact of meteorological uncertainty on a North Atlantic route was studied, and [41], where the information from a “Probabilistic Trajectory Planning” framework is employed to provide uncertainty information to the flight planner. In [186], ensemble data was also employed in a robust flight planning scheme based on mixed-integer linear programming.

However, the consideration of uncertainty as an *input* to the flight planning process instead of as a component of the *outcome* of a flight (i.e. as a robust or stochastic optimal control problem) plan has not been studied in depth on this field and the literature on the topic is sparse. There are some studies solving what we called the “uncertainty quantification OCP” (see Section 2.6.3). In [187], this approach is used in order to solve the problem of optimal aircraft routing through an environment with lethal threats of uncertain location, and it combines polynomial chaos for uncertainty quantification and pseudospectral methods for solving the deterministic OCP. A similar approach is employed in [188, 189] and related work in the context of conflict resolution.

In [50] (based on the method of [48]), the authors deal with what we call the “Robust OCP” in order to find the controls for the shortest climb of a supersonic aircraft, which used a polynomial chaos technique to characterize aerodynamic data uncertainty and integrates this representation into the problem formulation. The robust approach is compared with the uncertainty quantification approach. A robust scheme is also used in [49] in order to optimize the trajectory of an unmanned aerial vehicle based on dynamic soaring. Nevertheless, as noted in [50], the field of nonlinear robust dynamic optimization has not been explored in depth.

As a consequence, developing an appropriate methodology for 4D trajectory optimization under weather-related uncertainty is one of the main challenges addressed in this thesis.

2.7.4 Convective Environments

Turning our attention to convective weather avoidance, the flight planning literature describes multiple kinds of algorithms that have been employed for this purpose. The simplest of them are based on geometric-heuristic procedures [190, 191]. These methods do not generally consider thunderstorm evolution, uncertainty, or trajectory optimality; however, as discussed in Section 2.2.1, their simplicity results in fast computational times (thus being compatible with real-time usage) and easy implementation and integration with other tools and algorithms. Indeed, the Dynamic Weather Routes tool [192] implements the Autoresolver algorithm [193] with convective nowcasts in order to create a weather avoidance system, which has been enhanced to provide common avoidance trajectories to multiple aircraft in the Dynamic Multi-Flight Common Route Advisories system [194, 195].

A second class of methods is based on graph-search algorithms such as Dijkstra's shortest path algorithm, A* or D*. These methods feature good computational properties too, but it is harder to model time-varying costs (such as fuel burn) and aircraft dynamics. An A* algorithm is employed in [196] to generate weather avoidance routes in a TRACON impacted by convective activity. In [197], the authors combine a variant of Dijkstra's algorithm with a multi-objective genetic algorithm in order to produce a set of proposed reroutes around convective weather, as represented by the FAA's Convective Avoidance Weather Model (CWAM). In [198], CWAM weather forecasts were employed in combination with Dijkstra's algorithm to minimize a combination of fuel burn and expected cost of deviation due to weather. In [199], the any-angle path planning field D* algorithm is applied in a receding horizon fashion in order to minimize a combination of flight time and accumulated probability of bad weather conditions along the route.

Mathematical programming techniques have also been studied for route selection under convective weather [200]. In [201, 202], a global optimization method is employed for the solution of routing problems.

Alternatively, methods that consider aircraft dynamics explicitly can be employed. In [203], a receding horizon optimization scheme is employed in order to compute avoidance trajectories. A different approach involves formulating a stochastic reach-avoid problem [6, 10, 204], which is then solved through dynamic programming techniques. Markov Decision Process techniques have also been employed in related works [205]. In [206], which addresses the topic of for midair conflict resolution and bad weather avoidance, uncertainty in the dynamics and wind is modeled through a stochastic differential equation and the resulting optimal control problem is solved by discretizing the state space to obtain an MDP, which is solved through Jacobi iteration. In [207, 208, 209], algorithms for multi-aircraft routing and traffic flow management under uncertain convective weather conditions is proposed, again relying on MDP dynamic programming

techniques. As these works rely on state space discretization, they are therefore vulnerable to the “curse of dimensionality”, i.e., the exponential scaling of the required computational resources with the dimension of the problem.

None of the methods studied in the literature is yet able to generate realistic, model-based trajectories that take into account both aircraft dynamics and uncertain thunderstorm evolution with real-time or near-real-time computational times. Therefore, in this thesis we will also aim to develop a methodology for aircraft trajectory planning in the presence of convective thunderstorms whose evolution is considered uncertain.

Modeling

This chapter introduces the dynamical, physical and mechanical models that are employed in this thesis to describe the motion of the aircraft within an ATM context and the meteorological data sources employed. It is divided in three parts: Section 3.1 covers the behaviour and performance of the aircraft at the ATM scale. Section 3.2 introduces the models of the meteorological and atmospheric variables that be employed throughout this thesis, while and Section 3.3 describes the specific variables and data sources employed in the modeling of convection.

3.1 Aircraft Dynamics

The processes of trajectory simulation, prediction and optimization depend on accurate models of the aircraft dynamics that are able to replicate its behaviour in practice. At the same time, it is not necessary to consider phenomena that take place at spatiotemporal scales that are not of interest, as doing so would increase computational burdens and demand higher resolutions in time integration while not providing any practical benefit of significance.

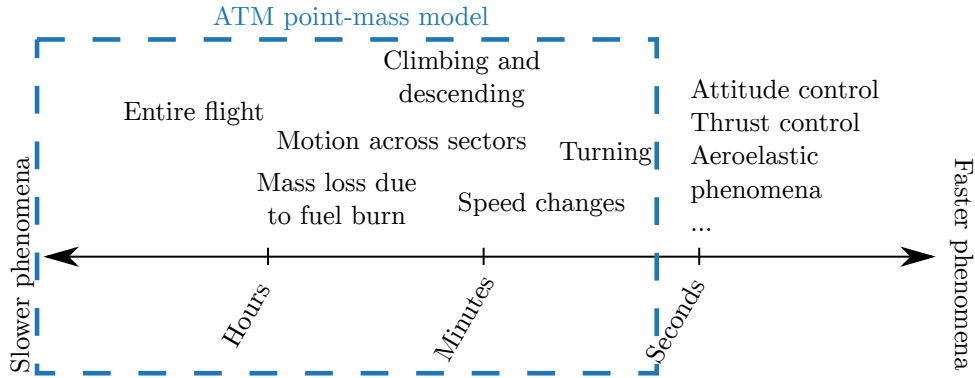


Figure 3.1: Characteristic times of aircraft dynamics.

In the ATM field, most studies employ a three-degrees-of-freedom (3-DoF) point-mass model. Figure 3.1 illustrates the main dynamic characteristics that are relevant for most ATM-related work and are, therefore, captured with this model. In this model, the aircraft is represented as a point mass moving in a three-dimensional space; therefore, its state at any point in time is characterized by its 3-dimensional position coordinates, a 3-dimensional velocity vector and its mass m .

3.1.1 Reference Systems

The position coordinates will be taken with respect to the WGS84 reference ellipsoid, as it is common in ATM studies. In this reference system, the surface of the Earth is represented as an oblate spheroid with an equatorial radius of $a = 6378.137$ km and a flattening of $1/298.257$; therefore, the polar radius has a length of $b = (1-f)a = 6356.752$ km; see Figure 3.2 for reference. The actual surface of the Earth, as modeled by the EGM2008, deviates from this reference ellipsoid by around ± 100 m.




Figure 3.2: The Earth-Centered, Earth-Fixed reference frame in the WGS84 reference ellipsoid

We start from a reference frame that rotates with the Earth (see Figure 3.2):

Definition 3.1.1. The Earth-Centered, Earth-Fixed (ECEF) reference frame is defined as having its origin at the center of gravity of the Earth and axes as follows:

- Its first axis X_{ECEF} points towards the intersection at latitude 0° of the Equator and the Greenwich meridian.
- Its third axis Z_{ECEF} points towards the North pole.
- Its second axis Y_{ECEF} completes a right-hand-oriented axes system.

Due to the translation and rotation of the Earth, the ECEF reference frame is not inertial. However, the associated non-inertial acceleration terms can be neglected at typical aircraft speeds, because they are several orders of magnitude smaller than that of gravity while the forces actuating on an aircraft are on the order of the gravitational attraction (in the vertical axis) and one order of magnitude below gravity (in the horizontal axis). Thus, we adopt the following assumption:

Assumption 3.1. *The ECEF reference frame can be considered inertial.*

The Cartesian coordinates in the ECEF system can now be employed to represent the position of the aircraft. In aviation, however, it is more convenient to employ Geodetic Reference System (GRS) coordinates; namely, latitude ϕ , longitude λ and geodetic altitude h . The conversion between ECEF position and GRS coordinates is given by Equations 3.1 – 3.3

$$x_{\text{ECEF}} = (R_N(\phi) + h) \cos \phi \cos \lambda \quad (3.1)$$

$$y_{\text{ECEF}} = (R_N(\phi) + h) \cos \phi \sin \lambda \quad (3.2)$$

$$z_{\text{ECEF}} = (R_N(\phi)(1 - e_o^2) + h) \sin \phi \quad (3.3)$$

where e_o is the *first eccentricity* and $R_N(\phi)$ is the ellipsoid radius of curvature in the prime vertical at a location of latitude ϕ :

$$e_o = \sqrt{f(2 - f)} \quad (3.4)$$

$$R_N(\phi) = \frac{a}{\sqrt{1 - e_o^2 \sin^2 \phi}} \quad (3.5)$$

We will also define $R_M(\phi)$, the radius of curvature in the meridian:

$$R_M(\phi) = \frac{a(1 - e_o^2)}{(1 - e_o^2 \sin^2 \phi)^{\frac{3}{2}}} \quad (3.6)$$

When working with the equations of motion in GRS coordinates, it is useful to employ a reference system that is aligned with the ellipsoid surface at each point. For that purpose, we define a *local horizon* reference system:

Definition 3.1.2. The local horizon reference system S_{LRS} is centered on the aircraft's center of gravity and has axes arranged as follows (see Figure 3.3):

- X_{LRS} parallel to the ellipsoid surface, pointing towards the North.
- Y_{LRS} parallel to the ellipsoid surface, pointing towards the East.
- Z_{LRS} orthogonal to the ellipsoid surface, pointing towards nadir.

We denote its associated basis vectors as $\{\mathbf{i}_{\text{LRS}}, \mathbf{j}_{\text{LRS}}, \mathbf{k}_{\text{LRS}}\}$. The evolution of the GRS coordinates when the aircraft moves with a velocity of $\mathbf{v}_{\text{ECEF}} = v_x \mathbf{i}_{\text{LRS}} + v_y \mathbf{j}_{\text{LRS}} + v_z \mathbf{k}_{\text{LRS}}$ with respect to the ECEF frame is given by:

$$\begin{aligned} \dot{\phi} &= \frac{v_x}{(R_M(\phi) + h)} \\ \dot{\lambda} &= \frac{v_y}{(R_N(\phi) + h) \cos \phi} \\ \dot{h} &= -v_z \end{aligned} \quad (3.7)$$




Figure 3.3: The Local Reference System

Finally, in order to represent the aerodynamic forces, we will employ two reference systems related to the concept of airspeed.

Definition 3.1.3. The *airspeed vector* \mathbf{v} , representing the velocity of the aircraft relative to the surrounding airmass, is defined as $\mathbf{v} = \mathbf{v}_{\text{ECEF}} - \mathbf{v}_w$, where \mathbf{v}_{ECEF} represents the velocity of the aircraft with respect to the Earth and \mathbf{v}_w represents the velocity of the wind field at the location of the aircraft.

Definition 3.1.4. The *airspeed* v is defined as the modulus of the airspeed vector, i.e., $v = \|\mathbf{v}\|$.

Definition 3.1.5. The *Mach number* M is defined as the ratio of the airspeed v to the speed of sound a_0 at a given position, i.e. $M = v/a_0$.

Definition 3.1.6. The *groundspeed* v_g is defined as the modulus of the projection of the absolute speed vector \mathbf{v}_{ECEF} onto the ellipsoidal plane, i.e.,

$$v_g = \|(\mathbf{v}_{\text{ECEF}} \cdot \mathbf{i}_{\text{LRS}})\mathbf{i}_{\text{LRS}} + (\mathbf{v}_{\text{ECEF}} \cdot \mathbf{j}_{\text{LRS}})\mathbf{j}_{\text{LRS}}\| \quad (3.8)$$

Assumption 3.2. The wind speed has a negligible component in the direction of Z_{LRS} ; thus, it can be expressed as $\mathbf{v}_w = w_y \mathbf{i}_{\text{LRS}} + w_x \mathbf{j}_{\text{LRS}}$

Definition 3.1.7. The *wind-fixed reference frame* S_W is defined with an origin at the center of mass of the aircraft and axes arranged as follows:

- X_W aligned with the airspeed vector
- Y_W orthogonal to the plane of symmetry of the aircraft, pointing towards star-board.
- Z_W contained in the plane of symmetry of the aircraft, closing a right-hand-oriented system.




Figure 3.4: The Wind-Fixed Reference Frame

This reference frame can be obtained from the local horizon by performing a sequence of three Euler rotations, illustrated in Figure 3.4:


1. A rotation of S_{LRS} around Z_{LRS} of magnitude χ (called the *heading angle*). This generates an intermediate reference system S_1 , with basis vectors $\{\mathbf{i}_1, \mathbf{j}_1, \mathbf{k}_1\}$.
2. A rotation of S_1 around Y_1 of magnitude γ (called the *flight path angle*). This generates an intermediate reference system S_2 , with basis vectors $\{\mathbf{i}_2, \mathbf{j}_2, \mathbf{k}_2\}$.
3. A rotation of S_2 around X_2 of magnitude μ (called the *bank angle*), generating the wind-fixed reference frame S_W .

3.1.2 Model of Forces


In order to derive the equations of motion of a vehicle, we need to consider the forces that are acting on it, as well as any source of variations to its mass. In the case of an aircraft, the relevant forces are the aircraft weight, the thrust from the engines and the aerodynamic forces, and its mass varies mainly due to the gradual consumption of fuel. In this work, we will employ an Aircraft Performance Model (APM) to supply the aerodynamic and propulsive functions associated to these forces: the Base of Aircraft Data (BADA) model, developed and maintained by the Eurocontrol Experimental Centre (EEC) in collaboration with aircraft manufacturers. There are currently two versions of BADA in use:

- The older BADA 3 family [210], which covers most of the aircraft types operating over the ECAC area.
- The newer BADA 4 family [211] covers fewer aircraft types (around 70% at the moment) and has a more restrictive licensing model, but it provides a more accurate modeling of several components (for example, by including compressibility effects in the drag polar).

Figure 3.5 represents the relationships between these forces; we now proceed to examine them.



(a) Longitudinal forces (for $\mu = 0$)



(b) Lateral forces (for $\gamma = 0$)

Figure 3.5: Forces diagram.

Weight

Due to the irregularity of the shape of the Earth and the inhomogeneity in its density, the gravitational field generated by the Earth is not regular and models such as EGM96 and EGM2008 can be used as approximations to it. Nevertheless, for ATM applications the error incurred by using a spherical approximation is negligible compared to other sources of error. Thus, the weight is given by:

$$\mathbf{W} = mg_0 \left(\frac{R_E}{R_E + h} \right)^2 \mathbf{k}_{\text{LRS}} \quad (3.9)$$

where R_E represents the mean radius of the Earth while $g_0 = 9.81 \text{ m}\cdot\text{s}^{-2}$ represents the average value of the gravitational attraction at sea level. Since, as explained in Section 3.1.4, we will be dealing with flight at an almost-constant altitude, the term $g_0 \left(\frac{R_E}{R_E + h} \right)$ can be replaced by a constant $g_h := g_0 \left(\frac{R_E}{R_E + h} \right)$. Thus, the weight force is given by:

$$\mathbf{W} = mg_h \mathbf{k}_{\text{LRS}} \quad (3.10)$$

For the remainder of this thesis, we will employ g as a synonym for g_h .

Thrust and fuel burn

When modeling the thrust force, the following simplification is valid for our purposes:

Assumption 3.3. *The thrust force \mathbf{T} is aligned with the airspeed vector \mathbf{v} ; i.e.,*

$$\mathbf{T} = T \mathbf{i}_w \quad (3.11)$$

The thrust is assumed to be regulated by the throttle or thrust parameter C_T , which we will model as a control of the system. Higher thrust production ($\uparrow C_T$) is associated with higher fuel burn f_c ; the relationship depends on the airspeed and other environmental variables:

$$f_c = f_c(C_T, v, \dots) \quad (3.12)$$

The variation of the mass of the aircraft can now be represented by:

$$\dot{m} = -f_c \quad (3.13)$$

Aerodynamic forces

As the aircraft moves in an airmass, it experiences an aerodynamic force \mathbf{A} . This force is commonly expressed in the wind-fixed reference frame as:

$$\mathbf{A} = \mathbf{L} + \mathbf{D} + \mathbf{Y} = -(\mathbf{L}\mathbf{k}_w + \mathbf{D}\mathbf{i}_w + \mathbf{Y}\mathbf{j}_w) \quad (3.14)$$

where \mathbf{L} is called the *lift*, \mathbf{D} is the *drag*, and \mathbf{Y} is the *lateral force*.

Assumption 3.4. *The aircraft is in symmetric, coordinated flight and the airspeed vector is contained in the plane of symmetry of the aircraft, i.e., the sideslip angle is zero.*

As a consequence of Assumption (3.4), the lateral force is negligible in this context. We now define the coefficients of lift C_L and drag C_D by normalizing by the product of the dynamic pressure $q = \frac{1}{2}\rho v^2$ and the wetted area S :

$$C_L = \frac{L}{\frac{1}{2}\rho S v^2} \quad (3.15)$$

$$C_D = \frac{D}{\frac{1}{2}\rho S v^2} \quad (3.16)$$

At a given aerodynamic configuration, the lift coefficient is determined by the angle of attack, which is controlled by the pilot or autopilot in order to generate enough lift to reach the desired rate of ascent or descent; thus, in this thesis we will treat the lift coefficient as a control of the dynamical system. An increase in the lift coefficient induces an increase in the drag coefficient; the relationship between both coefficients is known as the *drag polar*:

$$C_D = C_D(C_L, M) \quad (3.17)$$

As we will explain in Section 3.1.4, we will study the cruise phase of a flight, and thus we will only consider a clean aerodynamic configuration (i.e. with retracted landing gear and no deployed high-lift devices or speed brakes). The corresponding drag polars will be drawn from the BADA model.

3.1.3 The 4D Equations of Motion

The state vector of the aircraft is given by the three GRS position coordinates (latitude ϕ , longitude λ and geodetic altitude h), the three velocity parameters (airspeed v , flight path angle γ and heading χ) and its mass; meanwhile, the control vector is composed by the coefficient of thrust C_T , the coefficient of lift C_L and the bank angle μ . The airspeed vector \mathbf{v} can be expressed in the local horizon reference frame as:

$$\mathbf{v} = (v \cos \gamma \cos \chi) \mathbf{i}_{\text{LRS}} + (v \cos \gamma \sin \chi) \mathbf{j}_{\text{LRS}} - (v \sin \gamma) \mathbf{k}_{\text{LRS}} \quad (3.18)$$

Making use of Definition 3.1.3 and Assumption (3.2) to generate the absolute speed vector \mathbf{v}_{ECEF} and employing Equation (3.7), we can write the differential equations for the position coordinates as:

$$\begin{aligned}\dot{\phi} &= \frac{v \cos \gamma \cos \chi + w_y}{(R_M(\phi) + h)} \\ \dot{\lambda} &= \frac{v \cos \gamma \sin \chi + w_x}{(R_N(\phi) + h) \cos \phi} \\ \dot{h} &= v \sin \gamma\end{aligned}\tag{3.19}$$

At the same time, the evolution of the mass of the aircraft is given by Equation (3.13). In order to complete the equations of motion (EoM), we need to obtain differential equations for v , γ and χ ; for convenience, we will make use of additional assumptions to do so.

Assumption 3.5. (*Flat-Earth assumption*) *The rotation of the local horizon with respect to the ECEF frame can be disregarded at typical aircraft speeds, being on the order of $v_{\text{characteristic}}/R_E \leq 4 \cdot 10^{-5} \text{rad} \cdot \text{s}^{-1}$.*

Assumption 3.6. *The mass of the aircraft varies at a slow rate and the propulsive effect of mass loss is already included in the propulsive model, so that it can be safely assumed that the mass of the aircraft is constant for the derivation of the equations of motion.*

The derivative of the wind velocity as experienced by the aircraft with respect to time depends on the variation of the wind field with respect to time as well as its spatial gradient:

$$\frac{d\mathbf{v}_w}{dt} = \frac{\partial \mathbf{v}_w}{\partial t} + \mathbf{v}_{\text{ECEF}} \cdot \nabla \mathbf{v}_w\tag{3.20}$$

Assumption 3.7. *$\frac{d\mathbf{v}_w}{dt}$ is small enough that the wind can be assumed to be locally constant when deriving the equations of motion.*

In an inertial frame, Newton's second law for a point mass specifies that:

$$\sum \mathbf{F} = m \frac{d\mathbf{v}_{\text{ECEF}}}{dt}\tag{3.21}$$

Using Assumption (3.7) implies that the variation of the velocity of the aircraft is equal to the variation of the airspeed vector:

$$\frac{d\mathbf{v}_{\text{ECEF}}}{dt} = \frac{d\mathbf{v}}{dt} + \frac{d\mathbf{v}_w}{dt} \approx \frac{d\mathbf{v}}{dt} \Rightarrow \sum \mathbf{F} = m \frac{d\mathbf{v}}{dt}\tag{3.22}$$

Now, since the airspeed vector is expressed in a rotating reference frame as $\mathbf{v} = v\mathbf{i}_w$, expanding its derivative with the transport theorem leads to:

$$\sum \mathbf{F} = m \frac{d\mathbf{v}}{dt} = m (\dot{v}\mathbf{i}_w + \omega_w \times (v\mathbf{i}_w)) \quad (3.23)$$

where ω_w represents the rotation of the wind-fixed reference frame with respect to the inertial frame. Using Assumption (3.5) to disregard the terms associated to the slow rotation of the local horizon, we can express ω_w in terms of the wind-fixed system as:

$$\omega_w = \dot{\chi}\mathbf{k}_{\text{LRS}} + \dot{\gamma}\mathbf{j}_1 + \dot{\mu}\mathbf{i}_2 = \begin{bmatrix} \mathbf{i}_w & \mathbf{j}_w & \mathbf{k}_w \end{bmatrix} \begin{bmatrix} -\dot{\chi} \sin \gamma + \dot{\mu} \\ \dot{\chi} \cos \gamma \sin \mu + \dot{\gamma} \cos \mu \\ \dot{\chi} \cos \gamma \cos \mu - \dot{\gamma} \sin \mu \end{bmatrix} \quad (3.24)$$

and, therefore,

$$\omega_w \times (v\mathbf{i}_w) = v (\dot{\chi} \cos \gamma \cos \mu - \dot{\gamma} \sin \mu) \mathbf{j}_w + v (\dot{\chi} \cos \gamma \sin \mu + \dot{\gamma} \cos \mu) \mathbf{k}_w \quad (3.25)$$

Working on the left-hand side, we express the sum of forces in the wind-fixed frame:

$$\begin{aligned} \sum \mathbf{F} &= \mathbf{L} + \mathbf{D} + \mathbf{T} + \mathbf{W} = \\ &= -L\mathbf{k}_w + T\mathbf{i}_w - D\mathbf{i}_w + mg(-\sin \gamma \mathbf{i}_w + \cos \gamma \sin \mu \mathbf{j}_w + \cos \gamma \cos \mu \mathbf{k}_w) \\ &= (T - D - mg \sin \gamma) \mathbf{i}_w + mg \cos \gamma \sin \mu \mathbf{j}_w + (mg \cos \gamma \cos \mu - L) \mathbf{k}_w \end{aligned} \quad (3.26)$$

Gathering the terms on \mathbf{i}_w , \mathbf{j}_w and \mathbf{k}_w in both the right-hand and the left-hand side, we arrive to the following set of equations:

$$T - D - mg \sin \gamma = m\dot{v} \quad (3.27)$$

$$mg \cos \gamma \sin \mu = (\dot{\chi} \cos \gamma \cos \mu - \dot{\gamma} \sin \mu) mv \quad (3.28)$$

$$-L + mg \cos \gamma \cos \mu = (\dot{\chi} \cos \gamma \sin \mu + \dot{\gamma} \cos \mu) mv \quad (3.29)$$

Equation (3.27) provides the derivative of the airspeed. Multiplying Equation (3.28) by $\sin \mu$ and adding the product of Equation (3.29) by $\cos \mu$ leads to the following expression for $\dot{\gamma}$

$$mg \cos \gamma - L \cos \mu = -\dot{\gamma} mv \quad (3.30)$$

Similarly, multiplying Equation (3.28) by $\cos \mu$ and subtracting the product of Equation (3.29) by $\sin \mu$ leads to the following expression for $\dot{\chi}$

$$L \sin \mu = \dot{\chi} \cos \gamma mv \quad (3.31)$$

Collecting Equations (3.13), (3.19), (3.27), (3.30) and (3.31), we can build the 4D equations of motion for the full 4D point-mass model, which we will denominate the **4D- γ** model:

4D- γ model

State variables: $\phi, \lambda, h, v, \gamma, \chi, m$

Control variables: C_L, C_T, μ

Differential equations:

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

Note that some of the functional dependencies have been omitted for clarity; in particular, the environmental variables depend on the position, and the aerodynamic forces and propulsive relationships depend on the environmental variables.

A useful and common simplification is to ignore flight path angle dynamics (as they take place at a scale of seconds); employing this assumption, γ can be modeled as a control. In order to do that, we set $\dot{\gamma} = 0$ in Equation 3.30 and apply the resulting constraint either in an explicit form (as a constraint) or in an implicit form (by removing C_L as a control and replacing it with the expression resulting from solving Equation 3.30 for C_L).

4D-e modelState variables: $\phi, \lambda, h, v, \chi, m$ Control variables: C_L, C_T, γ, μ

Differential equations:

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

Algebraic constraints:

$$L(C_L) \cos \mu = mg \cos \gamma$$

4D-i modelState variables: $\phi, \lambda, h, v, \chi, m$ Control variables: C_T, γ, μ

Differential equations:

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

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

The first choice leads to the **4D-e** model, while the second choice generates the **4D-i** model. Both models are mathematically equivalent (and different from the **4D- γ** model) and generate equivalent optimization models when discretized with direct methods; however, they have different numerical properties, as the explicit model (4D-e) has more variables and constraints but lighter expression trees (and, therefore, faster function evaluation).

3.1.4 Reduced Models

In this work, we will only consider the cruise phase of a flight. The main reason is that the uncertainty due to wind is cumulative, and thus most of the uncertainty is “accumulated” during the cruise phase for any medium-haul or long-haul flight. While the methods presented in Chapters 4 and 5 can be applied to more complex scenarios, such a choice would complicate the implementation and exposition of the methodology while providing little additional insight or precision. Thus, such an extension falls outside of the scope of this thesis.

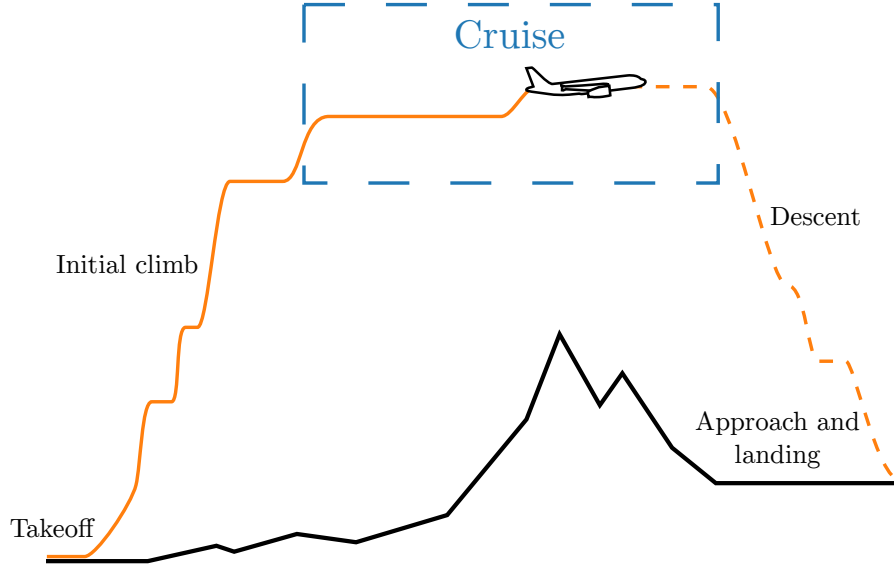


Figure 3.6: Phases of a flight.

Assumption 3.8. A flight level is defined as a spatio-temporal surface of constant geopotential pressure altitude (see Section 3.2.1 for the definition). A common designation is “flight level xyz” (often denoted as “FLxyz”), which describes a surface with $H_p^{FL} = 100 \cdot xyz$ feet.

Assumption 3.9. The flight takes place at a constant flight level.

Under Assumption 3.9, **4D-e** model and **4D-i** model can be employed directly with the addition of the following constraint:

$$H_p(\phi, \lambda, h, t) = H_p^{FL} \quad (3.32)$$

It is possible to simplify the model considerably by noting that $\gamma \approx 0$ under this assumption. Let us define the horizontal position vector $\mathbf{r}(t) = (\phi(t), \lambda(t))$. Then, Equation 3.32 implies that

$$\dot{H}_p(\phi, \lambda, h, t) = 0 \Rightarrow \frac{\partial H_p}{\partial h} \frac{dh}{dt} + \dot{\mathbf{r}} \cdot \nabla_{\mathbf{r}} H_p + \frac{dH_p}{dt} = 0 \quad (3.33)$$

Assumption 3.10. The derivatives of the geopotential pressure altitude $\nabla_{\mathbf{r}} H_p$ and dH_p/dt are negligible.

Since $\frac{\partial H_p}{\partial h} \approx 1$, Assumption 3.10 leads us to $\dot{h} \approx 0$ and thus $\gamma = \arcsin\left(\frac{\dot{h}}{v}\right) \approx 0$.

Assumption 3.11. *The flight path angle γ is 0*

This assumption leads to the explicit **2Dt-m μ -e** model and the implicit **2Dt-m μ -i** model:

2Dt-m μ -e model	2Dt-m μ -i model
State variables: $\phi, \lambda, v, \chi, m$	State variables: $\phi, \lambda, v, \chi, m$
Control variables: C_L, C_T, μ	Control variables: C_T, μ
Differential equations:	Differential equations:
$\begin{bmatrix} \dot{\phi} \\ \dot{\lambda} \\ \dot{v} \\ \dot{\chi} \\ \dot{m} \end{bmatrix} = \begin{bmatrix} \frac{v \cos \chi + w_y}{(R_M(\phi) + h)} \\ \frac{v \sin \chi + w_x}{(R_N(\phi) + h) \cos \phi} \\ \frac{T(C_T) - D(C_L)}{m} \\ \frac{g}{v} \tan \mu \\ -f_c(C_T) \end{bmatrix}$	$\begin{bmatrix} \dot{\phi} \\ \dot{\lambda} \\ \dot{v} \\ \dot{\chi} \\ \dot{m} \end{bmatrix} = \begin{bmatrix} \frac{v \cos \chi + w_y}{(R_M(\phi) + h)} \\ \frac{v \sin \chi + w_x}{(R_N(\phi) + h) \cos \phi} \\ \frac{T(C_T) - D(C_L(\mu))}{m} \\ \frac{g}{v} \tan \mu \\ -f_c(C_T) \end{bmatrix}$
Algebraic constraints:	
$L(C_L) \cos \mu = mg$	with $C_L(\mu) = \frac{2mg}{\rho v^2 S \cos \mu}$

If a problem does not involve turning dynamics to a significant extent (for example, in long-range planning), then these models can be simplified further by assuming that $\mu \approx 0$, removing turn dynamics $\dot{\chi}$ and letting χ be a control. Thus, we define the **2Dt-m-e** model and the **2Dt-m-i** model as:

2Dt-m-e modelState variables: ϕ, λ, v, m Control variables: C_L, C_T, χ

Differential equations:

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

Algebraic constraints:

$$L(C_L) \cos \mu = mg$$

2Dt-m-i modelState variables: ϕ, λ, v, m Control variables: C_T, χ

Differential equations:

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

$$\text{with } C_L = \frac{2mg}{\rho v^2 S}$$

Finally, if constant airspeed is assumed ($v = v_0$), it is useful to remove aerodynamic forces, fuel burn and mass evolution from the model. When building an optimal control problem, the conventional objective to minimize is a weighted average of flight time and fuel consumption, where the weight given to each objective depends on the cost structure of the airline. The main influence of this weight is the planned airspeed, as higher airspeeds shorten the duration of the flight but also require more thrust to compensate an increased drag. However, if the airspeed is fixed, then both objectives become heavily correlated, as a shorter flight will generally burn less fuel; thus, from a computational perspective, it is useful to reduce the problem to a minimum-time problem, as it will simplify and speed up the optimization progress. Fuel consumption can then be computed a posteriori by integrating the rest of the dynamics. Under these assumptions, the model simplifies to the **2Dt** model:

2Dt model

$$\begin{array}{l}
\text{State variables: } \phi, \lambda \\
\hline
\text{Control variables: } \chi \\
\hline
\text{Differential equations:} \\
\begin{bmatrix} \dot{\phi} \\ \dot{\lambda} \end{bmatrix} = \begin{bmatrix} \frac{v \cos \chi + w_y}{(R_M(\phi) + h)} \\ \frac{v \sin \chi + w_x}{(R_N(\phi) + h) \cos \phi} \end{bmatrix}
\end{array}$$

3.1.5 Aircraft Performance Limitations

The models introduced in Section 3.1.3 and Section 3.1.4 describe the essential differential-algebraic equations required to characterize the motion of the aircraft for our purposes. However, in order to obtain realistic trajectories and flight plans it is necessary to incorporate additional constraints into the model. These constraints represent the flight envelope and performance limits of the aircraft, and must be provided by the APM (in our case, the BADA specification). Table 3.1 represents the constraints that we will employ. Again, the functional dependencies of these variables in the states and controls have been omitted for clarity:

Description	Expression	Type
Maximum altitude	$h \leq h_{max}$	State
Maximum airspeed	$v_{CAS} \leq v_{CAS,max}$	State
Maximum Mach number	$M \leq v_{CAS,max}$	State
Max. and min. thrust	$C_{T,IDLE} \leq C_T \leq C_{T,max}$	State and control
Maximum bank angle	$ \mu \leq \mu_{max}$	Control
Maximum flight path angle	$ \gamma \leq \gamma_{max}$	State or control
Maximum lift coefficient	$C_L \leq C_{L,max}$	Control ¹
Mass limits ²	$m_{OEW} + m_{pl} \leq m \leq m_{MTOW}$	State

Table 3.1: Aircraft performance limitations

¹This limitation is a control constraint in explicit models, but a state and control constraint in implicit models.

²In this constraint, m_{OEW} represents the Operating Empty Weight, m_{pl} represents the mass of the crew, passengers, luggage, cargo, and crew equipment and m_{MTOW} represents the Maximum Takeoff

3.2 Meteorological Modeling

3.2.1 Environmental and Weather Model

The forces experimented by the aircraft are influenced by the characteristics of the air (temperature T , pressure P and density ρ) in which it moves, as well as the wind \mathbf{v}_w . These variables, depend, in general, on the position of the aircraft and on time:

$$\begin{aligned} T &= T(\phi, \lambda, h, t) \\ P &= P(\phi, \lambda, h, t) \\ \rho &= \rho(\phi, \lambda, h, t) \\ \mathbf{v}_w &= \mathbf{v}_w(\phi, \lambda, h, t) \\ g &= g(\phi, \lambda, h) \end{aligned} \tag{3.34}$$

In meteorology and aviation, it is convenient to express these relationships in terms of geopotential altitude, H , instead of geodetic altitude h . The geopotential altitude is defined in differential terms, starting from a value of 0 at the ellipsoid surface:

$$g_0 dH = g dh \tag{3.35}$$

where g_0 is the average gravitational acceleration at sea level and g is the gravitational acceleration at altitude.

The reference model for the temperature, pressure and density profiles is the International Standard Atmosphere (ISA) model, which has been adopted as an ICAO standard and is based on the assumption of hydrostatic equilibrium with a fixed temperature profile. For aviation purposes, it is sufficient to consider the troposphere and tropopause; in the ISA model, the boundary between both atmospheric layers takes place at an altitude of 11 km. In this model, the atmospheric variables can be computed as functions of the geopotential altitude:

$$\begin{aligned} T &= T_{\text{ISA}}(H) \\ P &= P_{\text{ISA}}(H) \\ \rho &= \rho_{\text{ISA}}(H) \end{aligned} \tag{3.36}$$

Weight.




Figure 3.7: The ISA model

The ISA model is particularly important for aviation, as it is employed to define certain reference variables. In particular, the geopotential pressure altitude, H_p , is defined as the altitude at which the pressure according to the ISA model corresponds to the measured pressure:

$$H_p = P_{\text{ISA}}^{-1}(P) \quad (3.37)$$

This measure is particularly important for aircraft navigation, since it can be computed just from barometer readings, facilitating air traffic control.

For simulation and optimization, it is not necessary to assume ISA profiles for atmospheric variables, as they can be obtained from meteorological forecast data. At the planning horizon, we will make use of global ensemble medium-range weather forecasts. The usage of global models allows us to consider intercontinental flights (as higher-resolution limited-area models have limited oceanic coverage), while the usage of ensemble forecasts instead of deterministic ones allows us to model the uncertainty in the forecasts. Section 3.2.2 describes turns to this topic.

3.2.2 Ensemble Forecasts

We will represent the uncertainty in meteorological forecasts at the planning horizon with EPS forecasts. The goal of an EPS is to produce a collection of forecasts (typically from 10 to 50) for the same prediction time that constitutes a representative sample of the possible future states of the atmosphere. The individual forecasts are called the *members* of the ensemble.

In order to generate different forecasts according to the uncertainty in the initial conditions, parametrized models of physical processes and approximation errors, different NWP centers employ combinations of several techniques. These methods include changing initial conditions in the most sensitive directions, changing the parameters of the simulation, combining different models or building time-lagged ensembles [212, 213, 214, 215].

Characteristics

For deterministic medium-range prediction, the typical setup of a forecasting system is a hydrostatic dynamical core with a horizontal resolution of around 15km and around 80 levels from surface to the top of the model, usually between 50 km and 80 km of altitude. Standard practices involve running the model 2 or 4 times a day, with output at a 1-, 3-, or 6-hour intervals.

For deterministic short-range prediction, the model is restricted to a limited area, non-hydrostatic effects are taken into account, and the resolution is improved to a horizontal grid of 1.5 to 5 km. Runs are also shorter (around a day of output) and run more frequently (from 4 to 8 times a day).

Since ensemble predictions require multiple model runs (10 to 50) for the same time interval, lower resolutions compared with the deterministic systems are used. The horizontal resolution is around twice the resolution for the deterministic prediction, and the vertical levels are also reduced. Depending on the area of interest and the target time interval, ensemble predictions can be classified in three categories:

- Global, medium-range forecasts are run for the whole globe and usually aim at forecasting 2 to 10 days ahead. They capture the uncertainty associated to planetary-level perturbations.
- Limited Area Model (LAM), short-range forecasts predict weather in a specific region at a 1-2 day timescale. The associated uncertainty captured by LAM forecasts is associated to mesoscale-alpha phenomena (200-2000 km characteristic scales).
- LAM, very short range forecasts produce predictions for a time-horizon of a few hours. The uncertainty is related to mesoscale-beta (20-200 km) and mesoscale-gamma (2-20 km) phenomena.

Medium-range ensemble forecasting

The World Meteorological Organization launched the THORPEX (The Observing System Research and Predictability Experiment) research initiative in 2004 to stimulate the development, usage, utility and accuracy of medium-range ensemble forecasts. Importantly, it created the The International Grand Global Ensemble (TIGGE) dataset³ [216].

TIGGE contains global medium-range ensemble forecast data produced by several NWP centers for the purposes of scientific research [217] in an homogeneous format. It is hosted at the website of the European Center for Medium-Range Weather Forecasts (ECMWF).⁴ The ECMWF [218], the Canadian Meteorological Center [219] and the National Center for Environmental Prediction (NCEP) [220], among others, develop ensemble prediction systems included in TIGGE. In 2014, the TIGGE-LAM (Limited Area Model) dataset was also launched with the aim of including regional ensembles on finer (from 2 to 10 km) grids.

These models run at a coarse resolution of 30 km or larger grid size in order to be able to forecast the weather at the whole globe for 2 to 10 days because of computational constraints. Nevertheless, as hardware becomes more efficient and less expensive, the NWP upgrade their systems in order to increase the resolution of their forecasts.

Limited Area Model ensemble forecasting

By limiting the area and period of interest, Limited Area Models (LAM) are able to use a higher resolution. Examples of these models are:

- The MOGREPS [221] ensemble (which contains both a global and a LAM model), developed by the UK Met Office. It is a 10-member ensemble that relies on perturbation of the initial conditions and the model physics.
- The SREPS [222] ensemble from the Spanish meteorological service AEMET, which combines five different LAM models (with a 25 km horizontal \times 40 vertical levels resolution) with five different initial and boundary conditions for a total of 25 ensemble members that predict up to 72h.
- The Norwegian Meteorological Institute's LAMEPS [223] is a 21-member ensemble that has a 28 km horizontal \times 31 vertical levels resolution with the aim of forecasting precipitation in northern Europe in a time horizon of 0 to 3 days.
- The COnsortium for Small-scale MOdelling's COSMO-LEPS [224] is a 16-member ensemble with a 10 km horizontal resolution and 40 levels in the vertical that is based on a downscaling of the global ECMWF ensemble runs.

³<http://tigge.ecmwf.int/documents/>

⁴<http://apps.ecmwf.int/datasets/>

- Météo France’s Pr vision d’Ensemble ARPEGE (PEARP) [225] is an EPS based on the ARPEGE deterministic model. It has a variable horizontal resolution that allows it to be more accurate over France and combines multiphysics and initial condition perturbations.
- The GLAMEPS (Grand Limited Area Model Ensemble Prediction System) is a pan-European EPS [226] in development resulting from the cooperation of the HIRLAM (The High Resolution Limited Area Model) and ALADIN (Aire Limit e Adaptation Dynamique D veloppement International) [227] consortia. It combines runs from a variant of the ECMWF EPS, HIRLAM and ALADIN models ran by different countries in order to produce a 52-member short-range high-resolution EPS.

Limited Area Models often focus on land regions in order to study the weather in areas of greater impact on human activities; as a consequence, their oceanic and maritime coverage is often limited. Therefore, as discussed in Section 1.4.1, we will employ global ensembles that allow us to model long oceanic flights (those that are most impacted by wind uncertainty) despite the higher resolution offered by LAMs. We employ the publicly available low-resolution output of these forecasts, which are run between 2 or 4 times a day (depending on the center), have a spatial resolution of 0.5 degrees of latitude and longitude (and 9 isobaric levels, plus the surface fields) and a temporal step size of 6 hours, while having a range of up to 16 days.

The relevant variables for our purposes are the values at isobaric levels of the temperature T , the geopotential height H and the zonal and meridional components of the wind \mathbf{v}_w . These fields will be spatially interpolated through cubic b-splines in order to produce a continuous and smooth field that can be employed in optimal control.

3.3 Convection

Convection occurs when the stratification of the atmosphere breaks down due to the presence of convective instability, with the formation of tower clouds that easily reach the tropopause. It is observed over a wide range of spatial and temporal scales [228]:

- At annual time periods and planetary-length scales, convection is modulated by meteorological events governed by astronomical and oceanic forcings, such as the Intertropical Convergence Zone (ITCZ), the monsoonal circulations and the El Ni o Southern Oscillation (ENSO).
- At the synoptic scale and middle latitudes, Rossby waves produce low and high pressure systems with length scales from 1000 to 6000 km and durations on the order of one or two weeks. While high pressure systems tend to produce a stable

stratification of the atmosphere that inhibits convection, low pressure systems tend to feature convection along cold, warm and occluded fronts and instability lines [229].

- At the mesoscale, we can find convective phenomena with characteristic lengths of 10 to 1000 km, such as isolated storms, squall lines, convective complexes and tropical cyclones, that can be embedded in bigger structures. [230].

While planetary and synoptic scale phenomena can be forecasted with global EPS [217, 231], these forecast systems lack the spatiotemporal resolution to resolve mesoscale phenomena. LAM EPS products can be employed at the mesoscale [226, 232], but they are not yet able to forecast individual cells with greater accuracy than nowcasts⁵ based on extrapolation.

Therefore, it is not possible to model individual convective cells at the planning horizon, since the temporal distance between the planning time and the potential convective weather encounter is in the order of several hours, which is greater than the range at which nowcasting is accurate (1 or 2 hours, depending on the size of the convective phenomena). We describe an alternative approach in Section 3.3.1.

At the tactical horizon, where we consider storm encounter scenarios on the sub-hour timescale, it is useful to employ a nowcasting system based on extrapolation instead of an NWP product. We employ the RDT system for this purpose, which we describe in Section 3.3.2.

3.3.1 Planning

While it is not possible to forecast the precise location of individual thunderstorm cells at the planning horizon, it is possible to identify higher regions where the conditions for convection are met, even if the precise location within that region where convection will trigger is not known. We can perform such an identification in a probabilistic manner, as described in [4].

We define a function $c_{t_p}(\phi, \lambda, t) : [-90, 90] \times [-180, 180] \times \mathbb{R} \rightarrow [0, 1]$ that represents the probability that the conditions for convection are met at the corresponding point in space and at time t , employing the forecasts available at planning time t_p . In order to compute this indicator, we make use of two variables included in the high-resolution ECMWF forecast:

- The total totals index (TT), which itself is the sum of two components: the vertical totals index (VT), defined as the temperature differential between the 850 hPa

⁵Within a meteorological context, “nowcasting” refers to forecasting for short ranges -6 hours at most, according to the WMO- at the mesoscale -weather phenomena of sizes between a few kilometers and a few hundred kilometers-, and relying more on extrapolation of current sensor data -radar echoes, mainly- than on NWP techniques.

and the 500 hPa levels, and the cross totals index (CT), the difference between the dewpoint temperature at 850 hPa and the temperature at the 500 hPa level. Higher values of TT imply higher intensity and likelihood of convective activity: some convective cells can appear at TT values of 45 K, while values higher than 50 K imply high probability of strong convective activity.

- The convective precipitation (CR) indicator, which represents the forecasted amount of precipitation coming from convective clouds.

Given thresholds TT_{TH} and CR_{TH} , we define c_{t_p} at a given point (ϕ, λ, t) as the fraction of the members where both of the two variables exceed their thresholds. Thus, c_{t_p} is an indicator that can be computed at the planning horizon to inform the flight planning process.

3.3.2 Tactical

At a tactical horizon, prediction is usually performed in the form of deterministic nowcasts. Instead of NWP based methods, thunderstorm nowcasting has been traditionally based on extrapolation of radar echoes [233], which has been improved by the usage of satellite data and Doppler radars. This is due to the fact that extrapolation is more effective than NWP at short timescales (1 to 2 hours). Nonetheless, the quality of extrapolation nowcasts degrades rapidly as the forecasting horizon increases (and particularly so for smaller convective cells). One extrapolation based nowcasting system of particular interest for aviation is the Corridor Integrated Weather System (CIWS)[234], which covers the Conterminous United States (CONUS) and is in use by the FAA.

More recently, extrapolation based nowcasting methods have started to include NWP forecasts in order to improve their accuracy beyond the 1 to 2-hour mark, where the statistical skill of extrapolation techniques in nowcasting drops below that of NWP forecasting due to the characteristic timescales of convection initiation, growth and decay [235, 236]. These combined systems blend the extrapolation and a high-resolution NWP forecast whose initial conditions have been improved by data assimilation. By placing higher weight in the extrapolation for short horizons and the NWP forecast for longer horizons, it is possible to combine the strengths of both approaches. One such combined system, designed for aviation purposes, is Consolidated Storm Prediction for Aviation (CoSPA) [237]. It integrates CIWS data with forecasts from the National Oceanic and Atmospheric Administration (NOAA) and the National Center for Atmospheric Research (NCAR). Additionally, the trend towards increasing computational power has increased the interest in higher resolution, convection-permitting NWP forecasts in the last years. In [238], the AROME ensemble forecast from Météo-France is employed in conjunction with statistical post-processing techniques in order to forecast reflectivity

(a variable that is highly correlated to convective weather hazards) in a probabilistic fashion. Nevertheless, the statistical performance of such blended systems is still limited and research is ongoing at meteorological centres around the world to improve in this aspect.

A different trend is the study of thunderstorm forecasting with data-driven methods with or without extrapolation, eschewing the use of NWP modeling. In [239], the authors conducted a characterization of the uncertainty in the movement of thunderstorms, as detected by a radar-based system, for usage in aircraft routing problems. Approaches to convection forecasting based on machine learning are also starting to be employed; for example, classifiers with Support Vector Machines (SVM) [240] or predictors with neural networks[241].

Finally, an alternative aviation-oriented approach focuses on analyzing and forecasting “weather avoidance fields”, i.e., regions of space that pilots will try to avoid. This can be performed by using statistical analysis to identify relationships between processed convection-related meteorological variables and weather-caused deviations of aircraft from the planned route. An important system of this kind is the Convective Weather Avoidance Model (CWAM) [242, 243, 244], covering the CONUS. The CWAM assigns a probability of deviation to each point in space according to spatially filtered values of the Vertically Integrated Liquid (VIL) metric and the difference between the flight altitude and the echo top of the storm. CWAM has been extensively studied and employed in US-based ATM research, but no model of this kind has been developed to the same extent for other regions of the world. Similarly, the work in [245] develops probabilistic deviation thresholds based on the NCWF-6 convective forecasts.

A different approach is taken in [246]; instead of employing flown trajectories as data source, a forecast-based scenario is presented to pilots, who are then asked to design avoidance trajectories in order to characterize pilot behaviour. It was found that pilots are willing to take more safety margins when it is easy to do so and fly closer to the storms when avoidance is difficult. The authors conclude that there might be a trade-off between efficiency and comfort that could be derived empirically.

In order to address the convective avoidance in Europe, we employ the RDT product. RDT is one of the products of the Satellite Application Facility on support to Nowcasting and Very Short-Range Forecasting (NWCSAF) consortium, participated by several European National Weather Services. It uses imagery collected from the SEVIRI instrument installed aboard the geostationary Meteosat Second Generation (MSG) satellites, with a horizontal resolution of 3 km, in order to characterize convective systems in an area around Europe (or another region with MSG coverage) every 15 minutes. The RDT system identifies convective cells and describes them as polygons on a latitude-longitude map in an object-oriented approach, while also determining certain attributes of the cells. The algorithm at the core of RDT operates in three steps:

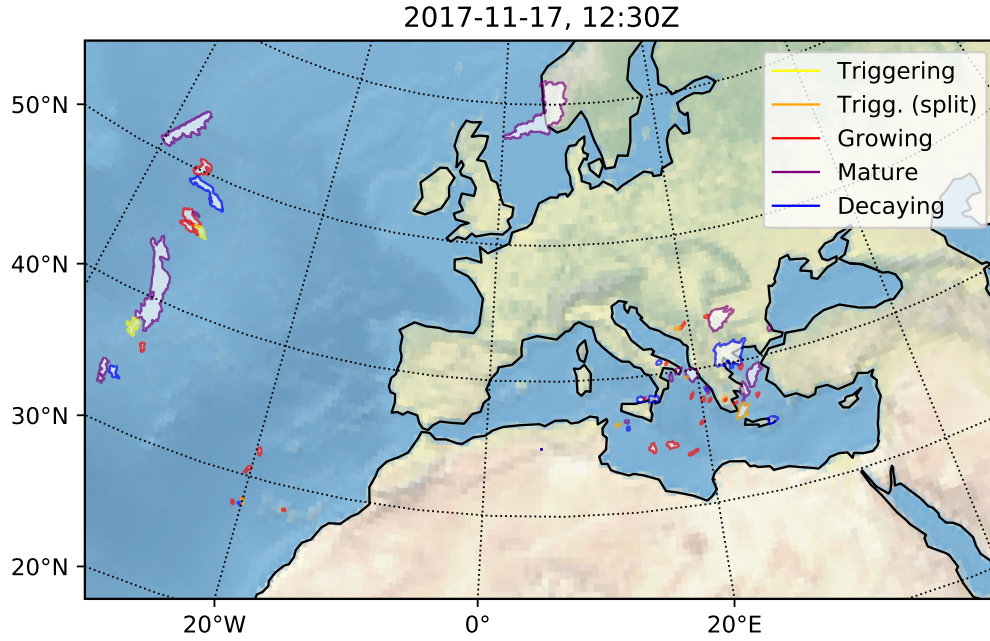


Figure 3.8: Convective systems over Europe at the indicated date, as identified by the RDT product.

- In first place, an *identification* of cells corresponding to cloud systems is performed using the infrared brightness temperature channel from the SEVIRI data. It identifies the cloud towers from the local minima of the infrared brightness temperature and then analyzes the local temperature pattern in the vicinity of the cloud tower to determine an adaptive temperature threshold that is employed to determine the extent of the cloud system.
- In second place, successive analyses are compared with a *tracking* algorithm that employs overlap and correlation methods in order to match previously determined cells with the newly identified objects and store their evolution and trajectories; this information is employed to generate new estimates of the speed and direction of the cells. Cloud splitting and merging events are also handled at this stage.
- Finally, a *discrimination* method is employed to distinguish convective objects from other cloud cells, which are much more numerous. It relies on data from multiple infrared and vapour channels and uses both spatial characteristic and temporal variations and trends. These parameters are stored in a learning database, which is used to train a statistical model that employs lightning occurrence data from the Météorage and EUCLID networks as ground truth.

The RDT output contains a list of the identified convective objects, along with characteristics such as the perimeter, the speed and direction of the motion or the cloud top pressure. One parameter of particular interest for our analysis is the *phase* of the convective cell, which classifies thunderstorm systems into five categories: triggering, triggering from split, growing, mature and decaying.

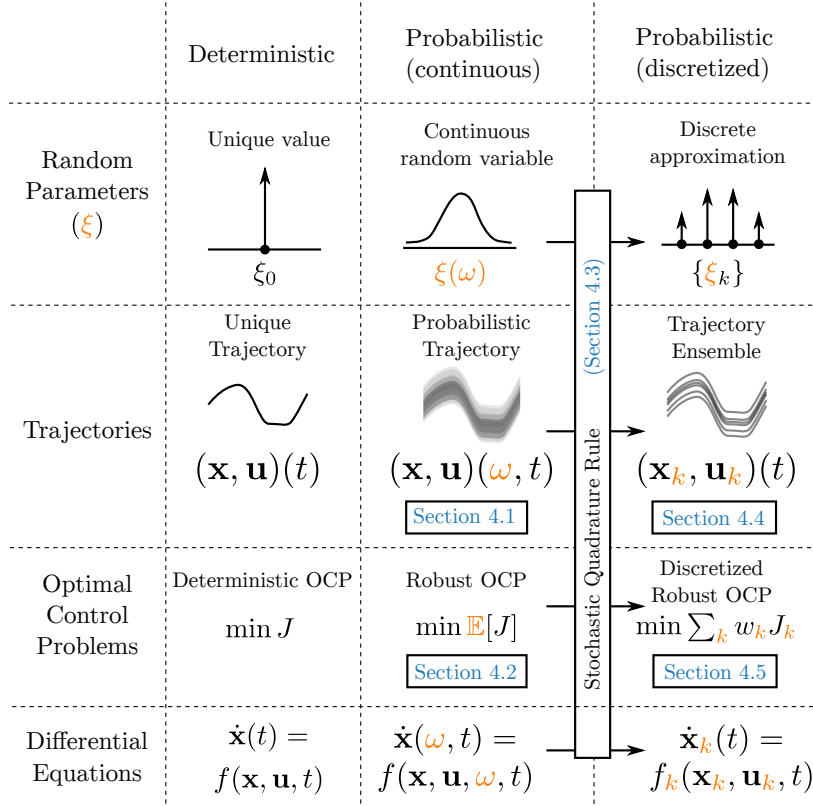
These parameters (in particular, the perimeter, speed and direction) can be employed to extrapolate the position of the storm in a deterministic manner (indeed, RDT output includes extrapolated nowcasts generated in this fashion). However, this extrapolation is deterministic and doesn't take into account the error inherent to the nowcast. Therefore, in Chapter 7 we will build a probabilistic nowcast based on the deterministic output generated by RDT.

Robust Optimal Control

Robust flight planning belongs to a class of engineering problems that cannot be solved in a practical sense with an open loop optimal control scheme, whether deterministic or stochastic. Therefore, it is necessary to develop a more general robust optimal control framework in order to build a methodology to optimize flight plans under uncertainty. In this chapter, we introduce such a formulation, which was first presented in [1, 12].

The proposed scheme rests on two pillars. The first one is the concept of *state tracking*, which allows part of the uncertainty in the dynamics to be transferred to the controls instead of “tracked” state variables; it is described in Section 4.1. This arrangement allows us to define optimal control problems in Section 4.2 whose solution can be employed in practice where open-loop trajectories are not desirable. The second pillar is a *stochastic quadrature rule* (Section 4.3), which is the denomination we employ for any scheme that discretizes the continuous probability distribution of the uncertain parameters into a number of finite discrete scenarios with associated weights. By employing a stochastic quadrature rule, the probabilistic trajectory can be discretized into a *trajectory ensemble* composed by multiple deterministic trajectories corresponding to each of the scenarios, as we describe in Section 4.4. The trajectory ensemble allows us to define a deterministic optimal control problem that approximates the robust optimal control problem (Section 4.5); this problem can then be solved with direct methods. Finally, in Section 4.6 we discuss an alternative formulation, based on feedback policies instead of state-tracking, that can be solved with the same tools.

Figure 4.1 illustrates the structure of the chapter, as well as the relationships between these concepts.

**Figure 4.1:** The robust optimal control methodology.

4.1 The State-Tracking Formulation

We start from a dynamical system represented with an RDE, as described in Section 2.5.2. An RDE is composed by a differential equation (Equation 4.1), the equality constraints (Equation 4.2) and the inequality constraints (Equation 4.3):

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{z}, \mathbf{u}, \xi, t) \quad (4.1)$$

$$\mathbf{h}(\mathbf{x}, \mathbf{z}, \mathbf{u}, \xi, t) = 0 \quad (4.2)$$

$$\mathbf{g}_L \leq \mathbf{g}(\mathbf{x}, \mathbf{z}, \mathbf{u}, \xi, t) \leq \mathbf{g}_U \quad (4.3)$$

In this setting, a realized trajectory is completely determined by two elements:

- The realization of the uncertain parameters ξ .
- The control or guidance law that determines the value of the controls \mathbf{u} at each point in time.

One possibility is to model the control law as a fixed control sequence, to be determined by the optimization process; i.e. the controls would only be dependent on time $\mathbf{u} = \mathbf{u}_F(t)$ and they would be equal in every scenario. This concept represents an “open-loop” scheme in practice, and it has been explored in previous literature that employs scenario-based approaches [48, 50, 51].

This “open-loop” formulation does not, however, represent a scheme that can be employed in practice for all general dynamic optimization problems. The main reason is that a fixed control law transfers all the uncertainty in the dynamics to the states according to the shape of the function f . This drift in the state space leads to multiple potential drawbacks:

- If the dynamical system is unstable, the trajectories might diverge towards undesirable regions of the state space.
- If the system needs to reach certain final conditions or some of the states need to follow unique trajectory, the state-space drift could make such a requirement infeasible.
- Similarly, if the cost functional places a heavy weight on diverging trajectories, the state-space drift leads to high efficiency losses.

In the context of trajectory planning for commercial aircraft, the second drawback is critical: not only we must enforce arrival of the aircraft at a fixed set of coordinates, but the horizontal path of the route must be unique under current and future concepts of operations. In addition, the airspeed schedule is nowadays assumed to be fixed, too.

A possible alternative, discussed in Section 4.6, is to replace the fixed control sequence by a *feedback policy*, i.e., a function $\mathbf{K} : \mathbb{R}^{n_x} \times \mathbb{R} \rightarrow \mathbb{R}^{n_u}$ such that the controls at every instant and scenario are determined by the relationship:

$$\mathbf{u}(\omega, t) = \mathbf{K}(\mathbf{x}(\omega, t), t) \quad (4.4)$$

However, this concept is not practical for ATM in the short-to-medium term, as it would require a paradigm shift towards a concept of operations based on “dynamic” flight plans. This would demand new Flight Management System (FMS), ATC and CNS technology built on top of this concept, as well as expensive retraining of controllers and other personnel.

Therefore, instead of looking for an optimal *control* or policy, then, we will look for an optimal *guidance*: we designate some of the state variables as “tracked” and we replace the unique controls $\mathbf{u}_F(t)$ that are applied identically in all scenarios by scenario-specific controls $\mathbf{u}(\omega, t)$. These controls attain the values that ensure that the tracked states follow a unique trajectory for all possible values of the random variables,

as long as it is feasible within the dynamics and constraints of the problem and the random variables are bounded. This is appropriate only for practical problems in which the actual system has low-level controllers that can track the desired trajectory in real time at a shorter timescale than that of the optimal control problem. In our practical context (long-range trajectory planning), the autopilot can compute the controls that are needed for the aircraft to follow a route at the calculated airspeeds and altitudes, and the omission of the short-term control dynamics (by assuming that the states are tracked exactly) does not introduce a significant error in performance.

We define the number of “control degrees of freedom” of the dynamical system as $d = n_z + n_u - n_h$. Let $q_x \leq \min\{n_x, d\}$ be the number of tracked states; without loss of generality, we can assume that the tracked states are the first q_x states (rearrange the state vector otherwise), i.e.

$$\mathbf{x} = \begin{bmatrix} x_1 & \dots & x_{q_x} & x_{q_x+1} & \dots & x_{n_x} \end{bmatrix}^T = \begin{bmatrix} \mathbf{x}_q \\ \mathbf{x}_r \end{bmatrix} \quad (4.5)$$

where \mathbf{x}_q is the tracked part of the state vector and \mathbf{x}_r is the untracked part. Let \mathbb{I}_n be the identity matrix of shape $n \times n$ and $0_{n_1, n_2}$ be the zero matrix (i.e. a matrix with zeroes in all its entries) of shape $n_1 \times n_2$. We define the matrix $E_x \in \mathbb{R}^{q_x \times n_x}$ as

$$E_x = \begin{bmatrix} \mathbb{I}_{q_x} & 0_{q_x, n_x - q_x} \end{bmatrix} \quad (4.6)$$

This matrix transforms the state vector into the “tracked states” vector $\mathbf{x}_q = E_x \mathbf{x}$ that contains only the states whose evolution is equal in all scenarios. In general, other combinations of states (represented by a general full-rank tracking matrix E_x that is not built as we have described, or even an analogous nonlinear transformation $E_x(\cdot)$) could also be tracked.

As emphasized earlier, controls are no longer unique for all scenarios. However, in order to completely determine the value of the controls at each moment, we need to close the remaining $d - q_x$ degrees of freedom. We can do that by selecting q_u controls and q_z algebraic variables to be equal in all scenarios (as in the “open-loop” problem), as long as $q_x + q_u + q_z = d$. In analogous fashion as E_x , we can define the tracking matrices E_z and E_u in order to select the tracked algebraic variables and controls. The tracking scheme is then completely defined by the tuple $\{q_x, q_z, q_u\}$ and a rearranging of the dynamic system in the manner described above (i.e. tracked variables before untracked variables).

With the aid of the tracking matrices, we can now define the tracking conditions (which apply almost surely, i.e., with probability one):

$$\begin{aligned} E_x(\mathbf{x}(\omega_i, t) - \mathbf{x}(\omega_j, t)) &= 0, \forall t, \forall \omega_i, \omega_j \in \Omega, \\ E_z(\mathbf{z}(\omega_i, t) - \mathbf{z}(\omega_j, t)) &= 0, \forall t, \forall \omega_i, \omega_j \in \Omega, \\ E_u(\mathbf{u}(\omega_i, t) - \mathbf{u}(\omega_j, t)) &= 0, \forall t, \forall \omega_i, \omega_j \in \Omega \end{aligned} \quad (4.7)$$

The tracking conditions enforce equality in the tracked variables between realizations: note that $E_x(\mathbf{x}(\omega_1, t) - \mathbf{x}(\omega_2, t))$ is the vector of differences between the tracked states in outcome ω_1 and the tracked states in outcome ω_2 . The other two conditions are analogous tracking conditions for the dependent variables and the controls.

For an arbitrary dynamical system with uncertain parameters, not every tracking scheme of this form is necessarily feasible. Indeed, it is trivial to build examples where there are no values of the untracked controls and algebraic variables that fulfill the dynamic equations along a tracked trajectory for certain values of the uncertain parameters, or examples where there are several or infinite values that do (therefore leaving the controls undetermined during operation of the system). We discuss this condition for linear systems in Appendix A.

Note that the “open-loop” robust control formulation employed in [48, 50, 51] and [247, Chapter 2] constitutes a special case of this formulation, with $q_x = q_z = 0$.

4.2 Definition of the Tracking Problem

Let $\mathbb{E}[\cdot]$ be the expectation operator associated to the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. We define the terminal cost or “Mayer term” $\Phi : \mathbb{R} \times \mathbb{R} \times \mathbb{R}^{n_x} \times \mathbb{R}^{n_x} \rightarrow \mathbb{R}$, the running cost or “Lagrange term” $\mathcal{L} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_z} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_\xi} \times \mathbb{R} \rightarrow \mathbb{R}$. We define the cost functional to minimize as:

$$J_R = \mathbb{E} \left[\Phi(t_0, t_f, \mathbf{x}(t_0), \mathbf{x}(t_f)) + \int_{t_0}^{t_f} \mathcal{L}(\mathbf{x}, \mathbf{z}, \mathbf{u}, \xi, t) dt \right] \quad (4.8)$$

We define the initial conditions function $\Psi_0 : \mathbb{R}^{n_\xi} \times \mathbb{R} \rightarrow \mathbb{R}^{n_x - q_x}$. This equation specifies the value of the untracked states as a function of the uncertain parameters, therefore allowing for uncertainty in the initial conditions:

$$\mathbf{x}_r(\xi, t_0) = \Psi_0(\xi, t_0) \quad (4.9)$$

If the initial conditions are known with certainty (as we will assume in Chapter 5), then $\Psi_0(\xi, t_0) = \mathbf{x}_{r,0}$.

We also define the function $\Psi : \mathbb{R} \times \mathbb{R} \times \mathbb{R}^{n_x} \times \mathbb{R}^{n_x} \rightarrow \mathbb{R}$ that contains the remaining boundary conditions:

$$\mathbb{E}[\Psi(t_0, t_f, \mathbf{x}(t_0), \mathbf{x}(t_f))] = 0 \quad (4.10)$$

While these conditions are imposed in average, the ones that depend only on tracked states collapse to boundary conditions that are imposed exactly (as the value of the tracked states at the endpoints is unique); otherwise, they remain probabilistic constraints.

The objective J_R and the boundary conditions Ψ are written in terms of mean value, but they can be easily generalized to other statistics under the expected value formulation. For example, the variance of a function $G(\xi)$ can be written as $\mathbb{E}[(G - \mathbb{E}[G])^2] = \mathbb{E}[G^2] - \mathbb{E}[G]^2$ using expected values.

We also group the differential-algebraic equations and constraints (2.37) (2.38) (2.39):

$$\begin{aligned} \frac{d}{dt}\mathbf{x} &= \mathbf{f}(\mathbf{x}, \mathbf{z}, \mathbf{u}, \xi, t), \\ \mathbf{h}(\mathbf{x}, \mathbf{z}, \mathbf{u}, \xi, t) &= 0, \\ \mathbf{g}_L &\leq \mathbf{g}(\mathbf{x}, \mathbf{z}, \mathbf{u}, \xi, t) \leq \mathbf{g}_U \end{aligned} \tag{4.11}$$

The robust optimal control problem with tracking (ROCT) can now be defined as:

$$\left. \begin{array}{ll} \text{minimize} & J_R(4.8) \\ \text{subject to} & \left. \begin{array}{l} \text{differential-algebraic equations (4.11)} \\ \text{boundary conditions (4.10)} \\ \text{tracking conditions (4.7)} \end{array} \right\} \end{array} \right\} \tag{ROCT}$$

4.3 Stochastic Quadrature Rules

In order to solve Problem (ROCT), we will approximate the uncertain parameters with a discrete probability distribution. We define a “stochastic quadrature rule” (SQR) as a procedure that generates a finite set of quadrature points $\{\xi_k\}$, $k \in \{1, \dots, N\}$ and weights $\{w_k\}$, $k \in \{1, \dots, N\}$, such that we can build an approximation¹ to a stochastic integral $I = \mathbb{E}[G(\xi)] = \int G(\xi)d\mathbb{P}$ with the sum:

$$QG = \sum_{k=1}^N w_k G(\xi_k) \tag{4.12}$$

where $G(\xi)$ is a well-behaved function. Basic statistical quantities, such as averages and variances, can be obtained with this integral by the corresponding function choices². There are a number of approaches with different approximation techniques that can provide a stochastic quadrature rule; depending on the problem and the desired approximation precision, different SQRs might be best. The ideal SQR can reach the desired approximation accuracy with a low N , because the size of the discretized robust optimal control problem grows linearly on N (as it can be observed in Section 4.5).

¹Naturally, we expect the discrete approximation to converge to the true value as $N \rightarrow \infty$

²e.g. the variance of $\gamma = G(\xi)$ can be computed as

$$\mathbb{E}[\gamma^2] - (\mathbb{E}[\gamma])^2 = \sum_{k=1}^N w_k G(\xi_k)^2 - \left(\sum_{k=1}^N w_k G(\xi_k) \right)^2.$$

As an example, consider Monte Carlo techniques. In this setting, the points ξ_k are randomly sampled from the probability distribution and weighted equally ($w_k = N^{-1}$). More advanced alternatives have been proposed in the literature; a brief summary has been provided in Section 2.5.4.

In certain practical settings, it might be the case that the uncertain parameters are already characterized as a discrete distribution: for example, if they are empirically estimated by a discrete set of measurements or simulated from a relatively small number of scenarios. EPS forecasts can be classified in this latter category, and we therefore use them “as given” with no other preprocessing (with equal weights $w_k = N^{-1}$ for each member).

4.4 The Trajectory Ensemble

Let $\mathbf{x}_q(t) : \mathbb{R} \rightarrow \mathbb{R}^{q_x}$ define a trajectory for the tracked states, and $\mathbf{z}_q(t)$ and $\mathbf{u}_q(t)$ analogously. Suppose an SQR has been chosen, with a number of points N . For each one of these points ξ_k , the tracking trajectory $(\mathbf{x}_q, \mathbf{z}_q, \mathbf{u}_q)(t)$ defines a unique trajectory given a full set of initial conditions; we will now collect every one of these N trajectories in a *trajectory ensemble*. We define the trajectory ensemble associated to a tracking trajectory $(\mathbf{x}_q, \mathbf{z}_q, \mathbf{u}_q)(t)$ as the set of trajectories $\{(\mathbf{x}_k, \mathbf{z}_k, \mathbf{u}_k)(t)\}$ with $k \in \{1, \dots, N\}$ such that the trajectory k is generated by the initial conditions $\mathbf{x}_k(t_0) = \mathbf{x}_0$ and the tracking trajectory with $\xi = \xi_k$, i.e.

$$\begin{aligned} \frac{d}{dt}\mathbf{x}_k &= \mathbf{f}(\mathbf{x}_k, \mathbf{z}_k, \mathbf{u}_k, \xi_k, t) \\ \mathbf{h}(\mathbf{x}_k, \mathbf{z}_k, \mathbf{u}_k, \xi_k, t) &= 0 \\ \mathbf{g}_L &\leq \mathbf{g}(\mathbf{x}_k, \mathbf{z}_k, \mathbf{u}_k, \xi_k, t) \leq \mathbf{g}_U \\ E_x \mathbf{x}_k(t) &= \mathbf{x}_q(t) \\ E_z \mathbf{z}_k(t) &= \mathbf{z}_q(t) \\ E_u \mathbf{u}_k(t) &= \mathbf{u}_q(t) \end{aligned} \tag{4.13}$$

Note that the last three equations equate the tracked variables to their value in the tracked trajectory and thus take the place of the tracking conditions.

We can now build an augmented dynamical system that comprises all the trajectories in the ensemble, whose state ($\mathbf{x}_E \in \mathbb{R}^{n_x N}$), control ($\mathbf{u}_E \in \mathbb{R}^{n_u N}$) and algebraic ($\mathbf{z}_E \in \mathbb{R}^{n_z N}$) vector contain the state, control and algebraic vectors of all the trajectories in the trajectory ensemble:

$$\mathbf{x}_E = \begin{bmatrix} \mathbf{x}_1 \\ \vdots \\ \mathbf{x}_N \end{bmatrix}; \quad \mathbf{z}_E = \begin{bmatrix} \mathbf{z}_1 \\ \vdots \\ \mathbf{z}_N \end{bmatrix}; \quad \mathbf{u}_E = \begin{bmatrix} \mathbf{u}_1 \\ \vdots \\ \mathbf{u}_N \end{bmatrix} \tag{4.14}$$

We define the differential equation, algebraic equations and inequality constraints of this augmented dynamical system as:

$$\mathbf{f}_E(\mathbf{x}_E, \mathbf{z}_E, \mathbf{u}_E, t) = \begin{bmatrix} \mathbf{f}(\mathbf{x}_1, \mathbf{z}_1, \mathbf{u}_1, \xi_1, t) \\ \vdots \\ \mathbf{f}(\mathbf{x}_N, \mathbf{z}_N, \mathbf{u}_N, \xi_N, t) \end{bmatrix} \quad (4.15)$$

$$\mathbf{h}_E(\mathbf{x}_E, \mathbf{z}_E, \mathbf{u}_E, t) = \begin{bmatrix} \mathbf{h}(\mathbf{x}_1, \mathbf{z}_1, \mathbf{u}_1, \xi_1, t) \\ \vdots \\ \mathbf{h}(\mathbf{x}_N, \mathbf{z}_N, \mathbf{u}_N, \xi_N, t) \end{bmatrix} \quad (4.16)$$

$$\mathbf{g}_E(\mathbf{x}_E, \mathbf{z}_E, \mathbf{u}_E, t) = \begin{bmatrix} \mathbf{g}(\mathbf{x}_1, \mathbf{z}_1, \mathbf{u}_1, \xi_1, t) \\ \vdots \\ \mathbf{g}(\mathbf{x}_N, \mathbf{z}_N, \mathbf{u}_N, \xi_N, t) \end{bmatrix} \quad (4.17)$$

Note that, while this augmented dynamical system represents (approximately) a system with uncertainty, it is a deterministic system. This means that we can use it to formulate a deterministic optimal control problem that approximates Problem (ROCT).

4.5 Formulation of the Discretized Problem

We now proceed to build such a deterministic approximant to Problem (ROCT), which we will call DROCT (Discretized ROCT). Using the trajectory ensemble and the SQR, we define the Mayer and Lagrange terms of the Problem (DROCT) as follows:

$$J_E = \Phi_E(t_0, t_f, \mathbf{x}_E(t_0), \mathbf{x}_E(t_f)) + \int_{t_0}^{t_f} \mathcal{L}_E(\mathbf{x}_E, \mathbf{z}_E, \mathbf{u}_E, t) dt \quad (4.18)$$

$$\Phi_E(\mathbf{x}_E(t_0), \mathbf{x}_E(t_f)) = \sum_{k=1}^N w_k \Phi(\mathbf{x}_k(t_0), \mathbf{x}_k(t_f)) \quad (4.19)$$

$$\mathcal{L}_E(\mathbf{x}_E, \mathbf{z}_E, \mathbf{u}_E, t) = \sum_{k=1}^N w_k \mathcal{L}(\mathbf{x}_k, \mathbf{z}_k, \mathbf{u}_k, \xi_k, t) \quad (4.20)$$

and discretize the boundary conditions as

$$\Psi_E(t_0, t_f, \mathbf{x}_E(t_0), \mathbf{x}_E(t_f)) = \sum_{k=1}^N w_k \Psi(t_0, t_f, \mathbf{x}_k(t_0), \mathbf{x}_k(t_f)) \quad (4.21)$$

For concise writing of the discretization of the tracking conditions (4.7), we will

define the matrix $E_x^N \in \mathbb{R}^{q_x(N-1) \times n_x N}$ as:

$$E_x^N = \begin{bmatrix} E_x & & \\ & \ddots & \\ & & E_x \end{bmatrix} \begin{bmatrix} \mathbb{I}_{n_x} & -\mathbb{I}_{n_x} & & \\ & \mathbb{I}_{n_x} & -\mathbb{I}_{n_x} & \\ & & \ddots & \ddots \\ & & & \mathbb{I}_{n_x} & -\mathbb{I}_{n_x} \end{bmatrix} \quad (4.22)$$

$E_z^N \in \mathbb{R}^{q_z(N-1) \times n_z N}$ and $E_u^N \in \mathbb{R}^{q_u(N-1) \times n_u N}$ can be defined in analogous fashion. These matrices map the ensemble state vector to the differences in the tracked states between trajectories.

Making use of Equations (4.13) - (4.17) as well as Equations (4.18) - (4.22), we can complete now the formulation of the deterministic approximant:

$$\left. \begin{array}{l} \text{minimize} \quad J_E \\ \text{subject to} \quad \dot{\mathbf{x}}_E = \mathbf{f}_E(\mathbf{x}_E, \mathbf{z}_E, \mathbf{u}_E, t) \\ \quad \quad \quad \mathbf{h}_E(\mathbf{x}_E, \mathbf{z}_E, \mathbf{u}_E, t) = 0 \\ \quad \quad \quad I_G \mathbf{g}_L \leq \mathbf{g}_E(\mathbf{x}_E, \mathbf{z}_E, \mathbf{u}_E, t) \leq I_G \mathbf{g}_U \\ \quad \quad \quad E_x^N \mathbf{x}_E = 0 \\ \quad \quad \quad E_z^N \mathbf{z}_E = 0 \\ \quad \quad \quad E_u^N \mathbf{u}_E = 0 \\ \quad \quad \quad \psi_E(t_0, t_f, \mathbf{x}_E(t_0), \mathbf{x}_E(t_f)) = 0 \end{array} \right\} \quad (\text{DROCT})$$

where $I_G = [\mathbb{I}_{n_g} \dots \mathbb{I}_{n_g}]^T \in \mathbb{R}^{n_g N \times n_g}$.

This optimal control problem can now be solved with standard deterministic direct methods. From the point of view of implementation, defining and modeling the tracked variables in every scenario is not necessary as they are equal; as a consequence, implementation of the tracking constraints $E_x^N \mathbf{x}_E = 0$, $E_z^N \mathbf{z}_E = 0$ and $E_u^N \mathbf{u}_E = 0$ can also be omitted. The resulting formulation constitutes the “compact form” of Problem (DROCT).

Note that the size of this problem is proportional to both N and the dimension of the original dynamical system minus the number of tracked variables (if $N \gg 1$), with the number of variables being $\mathcal{O}(N(n_x + n_h))$ (in the compact form) and the number of function evaluations (of f , g , h and \mathcal{L}) per node being $\mathcal{O}(N(n_x + n_h + n_g + 1))$. That is, a DROCT is approximately N times larger than a similar deterministic problem, but it preserves the linear scaling of the problem size on the dimension of the state space that represents a core advantage of numerical optimal control. Therefore, it avoids the “curse of dimensionality”, i.e., the exponential scaling of the problem size and cost on the dimension of the state space that appears in dynamic programming.

While we will not do so, it is possible to partially parallelize this problem in order to speed it up. The computational cost of an optimal control problem solved through direct

collocation can be divided in function evaluations (which includes the computation of derivative matrices) and NLP algorithm time. The former can be naturally parallelized in a scenario-wise fashion, but it is harder to replace NLP solvers by parallel versions; depending on which of the two parts of the cost dominates, parallelization would be more or less effective.

Finally, we note that using the SQR we can include higher moments (such as the variance) in the cost functional, as suggested in Section 4.2.

4.6 Feedback Policies

When the real-world system to be controlled does not enforce the usage of the tracking framework, a natural election is to employ a general feedback policy in Equation (4.4). This idea is employed in the *closed-loop* or *policy-based* dynamic optimization paradigm, using the taxonomy introduced in Section 2.1. A wide range of dynamic optimization methods, such as dynamic programming, Markov Decision Process optimization, reinforcement learning or linear-quadratic control, can be classified under this label. However, feedback policies have not been studied in depth within the realm of general nonlinear, continuous time optimal control.³ In this section, we show how to use the trajectory ensemble framework to model feedback laws in a form that generates optimal control problems that can be solved by direct transcription methods.

We start by introducing *reference variables* $\mathbf{y} : \mathbb{R} \rightarrow \mathbb{R}^{n_y}$ that are equal in all scenarios and depend on time, $\mathbf{y} = \mathbf{y}(t)$. These variables are chosen by the control designer and they can be employed to represent reference values for the states or controls⁴. Then, we introduce the feedback law $\mathbf{K}_p : \mathbb{R}^{n_x} \times \mathbb{R}^{n_y} \times \mathbb{R}^{n_z} \times \mathbb{R}^{n_p} \times \mathbb{R} \rightarrow \mathbb{R}^{n_u}$, which represents a general parametrized policy with n_p free parameters that determines the controls through the following relationship:

$$\mathbf{u} = \mathbf{K}_p(\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{p}, t) \quad (4.23)$$

where $\mathbf{x} = \mathbf{x}(\omega, t)$ and $\mathbf{z} = \mathbf{z}(\omega, t)$ represent the scenario-specific realization of the trajectory, $\mathbf{p} \in \mathbb{R}^{n_p}$ is a vector of parameters and $\mathbf{u} = \mathbf{u}(\omega, t)$ represents the scenario-specific computed values of the controls at time t . As an illustrative example, consider a system with a single state x and a single control u . One possible simple choice could be a linear policy that employs one reference variable y and one parameter p :

$$u = y + p \cdot x \quad (4.24)$$

³With the linear law for the coefficient of lift in [247, Chapter 3] representing the main exception, to the best of the author's knowledge.

⁴Because of the general shape of \mathbf{K}_p , it is not necessary, in principle, to use these reference variables as they could be modeled by parametrized functions of time within \mathbf{K}_p ; for example, as a Lagrange polynomial interpolation with the node values being represented by some of the parameters \mathbf{p} . However, its usage is convenient from the point of view of clarity and ease of implementation.

This policy represents a feedback law that can be decomposed into two terms: the quantity $y + p \cdot \mathbb{E}[x]$, which represents the value of the control at the average trajectory, and the proportional term $p(x - \mathbb{E}[x])$, which represents an extra increment or decrement of the control which is proportional to the difference between the realization of the state and the average value of the state. Thus, this scheme represents a classical proportional controller with a changing reference signal and the solution of the equivalent optimal control problem represents an integrated optimization of the reference trajectory and tuning of the P controller on a nonlinear system.

Note that other parametrizations of the control law are possible; for example, one might use an implicit parametrization $\mathbf{K}_p^i(\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{u}, \mathbf{p}, t) = 0$ or a constraint with dimension n_u that only relates to the states, thus making the value of the control implicit. We discuss here only the case of the explicit policy, as it is straightforward to implement the solution in practice without resorting to a potentially complex process of solving nonlinear algebraic systems of equations in real time.

The robust optimal control problem with the explicit feedback law can be defined as

$$\left. \begin{array}{ll} \text{minimize} & J = \mathbb{E} \left[\Phi(t_0, t_f, \mathbf{x}(t_0), \mathbf{x}(t_f)) + \int_{t_0}^{t_f} \mathcal{L}(\mathbf{x}, \mathbf{z}, \mathbf{u}, \xi, t) dt \right] \\ \text{subject to} & \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{z}, \mathbf{u}, \xi, t), \\ & \mathbf{h}(\mathbf{x}, \mathbf{z}, \mathbf{u}, \xi, t) = 0, \\ & \mathbf{g}_L \leq \mathbf{g}(\mathbf{x}, \mathbf{z}, \mathbf{u}, \xi, t) \leq \mathbf{g}_U \\ & \mathbf{u} = \mathbf{K}_p(\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{p}, t) \\ & \mathbb{E}[\Psi(t_0, t_f, \mathbf{x}(t_0), \mathbf{x}(t_f))] = 0 \end{array} \right\} \quad (\text{ROCF})$$

We can now proceed to discretize the problem with the aid of the trajectory ensemble, in analogous fashion as described in Section 4.5. Define the ensemble policy as:

$$\mathbf{K}_E(\mathbf{x}_E, \mathbf{y}, \mathbf{z}_E, \mathbf{p}, t) = \begin{bmatrix} \mathbf{K}_p(\mathbf{x}_1, \mathbf{y}, \mathbf{z}_1, \mathbf{p}, t) \\ \vdots \\ \mathbf{K}_p(\mathbf{x}_N, \mathbf{y}, \mathbf{z}_N, \mathbf{p}, t) \end{bmatrix}. \quad (4.25)$$

and the discretized version of Problem (ROCF) can be written as:

$$\left. \begin{array}{ll} \text{minimize} & J_E \\ \text{subject to} & \dot{\mathbf{x}}_E = \mathbf{f}_E(\mathbf{x}_E, \mathbf{z}_E, \mathbf{u}_E, t) \\ & \mathbf{h}_E(\mathbf{x}_E, \mathbf{z}_E, \mathbf{u}_E, t) = 0 \\ & I_G \mathbf{g}_L \leq \mathbf{g}_E(\mathbf{x}_E, \mathbf{z}_E, \mathbf{u}_E, t) \leq I_G \mathbf{g}_U \\ & \mathbf{u}_E = \mathbf{K}_E(\mathbf{x}_E, \mathbf{y}, \mathbf{z}_E, \mathbf{p}, t) \\ & \psi_E(t_0, t_f, \mathbf{x}_E(t_0), \mathbf{x}_E(t_f)) = 0 \end{array} \right\} \quad (\text{DROCF})$$

where J_E is defined as in Equations (4.18) – (4.20).

Choosing an appropriate shape for \mathbf{K}_p represents the critical research question when employing this formulation. Clearly, the optimality of the solution depends on the choice of \mathbf{K}_p , as the optimal solution of Problem (DROCF) contains the optimal policy among those policies that have the same shape at best, but not the optimal policy among arbitrary functions. This difference suggests that it would be ideal to use a general function approximation scheme that could approximate arbitrary functions: Taylor expansions, polynomial interpolators (with sparse grids in higher dimensions), or even neural networks (as it is done in “neuro-dynamic programming” [248]). Nonetheless, most of these choices could lead to additional issues:

- NLP software is designed to reliably produce good results with continuous, smooth functions. However, both theory and practice show that some problems have solutions with non-smooth, even discontinuous policies (such as bang-bang solutions). In addition, most classical function approximation schemes are efficient for smooth functions but generate problematic approximations to non-smooth functions (such as the “Gibbs phenomenon”). On the other hand, for some problems it might be the case that the existence of uncertainty “smooths out” the optimal policy.
- In order to model complex policies, it might be attractive to employ a very flexible form (such as a neural network). This could generate undesirable local minima; in addition, they could require a number of parameters that is big relative to the number of scenarios, thus running into potential overfitting issues where the policy ‘learns’ the discrete scenarios instead of the true, continuous uncertainty. Increasing the amount of scenarios would mitigate this issue, but the computational cost could become prohibitive; alternatively, some regularization scheme could be employed, which would require tuning and experimentation that does not necessarily generalize to other optimal control problems (the “hyper-parameter” question in machine learning).
- Furthermore, even when employing sparse grids, the number of parameters will grow quickly in high-dimensional problems when employing a classical function approximation method, leading to the mentioned issues again.
- The time resolution of optimal control problems is often coarser than the resolution that would be required to resolve fast control dynamics that might be implied by the optimal policy, leading to sub-optimal policies unless a fine grid (with its implications for computational cost) is employed.

As the main focus of this thesis is on the tracking formulation instead of the feedback formulation, these speculative questions will not be addressed in depth within the scope of this work and are left for future research to investigate.

Robust Flight Planning

In this chapter, a methodology for optimizing free-route flight plans under wind uncertainty at the planning horizon (2-3 hours before departure) is proposed. It is based on the robust optimal control framework introduced in Chapter 4 and the EPS forecasts described in Chapter 3. By considering the different weather scenarios contained in an ensemble forecast with the robust optimal control framework, it is possible to design trajectories that are optimal when considering multiple possible weather scenarios. Furthermore, it is also possible to trade off some average efficiency in order to increase predictability.

The chapter is divided in two sections. In Section 5.1, we describe the mathematical construction of the robust flight planning problem. Then, a transatlantic flight planning scenario is presented in Section 5.2 to demonstrate the proposed method. We show how the algorithm can produce different flight plans for different airline priorities, and how these trajectories achieve the desired combinations of objectives.

5.1 Concept

We consider the problem of flight planning at the planning horizon, i.e., 2-3 hours before departure. At this point, the weather forecasts carry significant uncertainties. For medium-haul and long-haul flights, in which uncertainty becomes more important due to the longer flight time, the weather in the latter parts of the flight has to be forecasted with a higher lead time, which implies that the forecast will have a higher prediction error.

In order to model the error in the forecast, we rely on EPS forecasts. Within the robust optimal control framework of Chapter 4, the uncertainty will be represented with an SQR where each member of the EPS forecast corresponds to a quadrature point. Each scenario will be weighted equally; if the EPS contains N members, then the weight of each member is $w_k = 1/N$.

5.1.1 Dynamical Model

As this is a routing problem where turning dynamics become relatively unimportant, we will work with a **2Dt-m-i** model from Section 3.1.4

2Dt-m-i model

State variables: ϕ, λ, v, m

Control variables: C_T, χ

Differential equations:

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

with $C_L = \frac{2mg}{\rho v^2 S}$

With this model, we consider a constant flight level for simplicity, but we will note that our methodology can be extended to full 4D problems.¹ We will also consider the

¹Direct methods are flexible enough that they can handle more complex problems; we choose our assumptions, which are comparable to most of the published routing algorithms, for simplicity.

constraints from Section 3.1.3 that are related to the airspeed and the thrust limits.

It is advantageous for computational purposes and for clarity of exposition, as we will explain in Section 5.1.2, to reformulate this dynamical system as a differential-algebraic system (DAE) with the addition of the ground speed v_G as an algebraic variable and the course ψ as a control variable, linked to the remaining variables by two new equality constraints (see Figure 5.1 for a graphical explanation). The reformulated system is given by:

2Dt-m- ψ model

State variables:	ϕ, λ, v, m
Control variables:	C_T, χ, ψ
Algebraic variables:	v_G
Differential equations:	
$\begin{bmatrix} \dot{\phi} \\ \dot{\lambda} \\ \dot{v} \\ \dot{m} \end{bmatrix}$	$= \begin{bmatrix} \frac{v_G \cos \psi}{(R_M(\phi) + h)} \\ \frac{v_G \sin \psi}{(R_N(\phi) + h) \cos \phi} \\ \frac{T(C_T) - D(C_L)}{m} \\ -f_c(C_T) \end{bmatrix}$
Algebraic equations:	
$\begin{bmatrix} v_G \cos \psi \\ v_G \sin \psi \end{bmatrix}$	$= \begin{bmatrix} v \cos(\chi) + w_x \\ v \sin(\chi) + w_y \end{bmatrix}$
with $C_L = \frac{2mg}{\rho v^2 S}$	

The inequality constraints apply in the same fashion as before, with the addition of

$$v_G \geq 0 \tag{5.1}$$

to ensure uniqueness of v_G and ψ (otherwise, $(-v_G^*, \psi^* + \pi)$ would produce the same left-hand side of the algebraic constraints as (v_G^*, ψ^*)).




Figure 5.1: Relationship between airspeed, groundspeed, wind, heading and course.

5.1.2 Coordinate Transformation

In the robust optimal control framework that we have presented, the independent variable (time) is unique in all scenarios (it varies in the same range). Therefore, the direct application of the state-tracking formulation would demand the position of the aircraft to be in a fixed schedule with respect to time in all scenarios, and thus the groundspeed would be fixed. This implies that, under this concept, the variability in wind speed would be fully compensated by airspeed modifications, with associated variations in fuel burn². There are multiple reasons why this is not convenient in cruise scenarios³:

- There would be a large spread in the airspeed values between scenarios, resulting in airspeeds that are inefficiently high or low in the extreme scenarios. Because the inefficiency of a certain airspeed level is approximately quadratic on the difference with respect to the optimal airspeed, a high spread in airspeeds leads to low efficiency in average.
- In some regions with high wind uncertainty, the average airspeed would be reduced below efficient ranges in order to have sufficient margin of variation for the most unfavourable scenarios: otherwise, the airspeed in these scenarios would go beyond the performance envelope.

²In descent scenarios, this is not necessarily the case, as the deviation can be translated into altitude deviations [249].

³Nevertheless, we will note that this type of solutions can also be obtained with our approach and the extension presented in Section 5.1.5; the formulation that we will present there generalizes and includes this kind of fixed-schedule solutions, allowing for intermediate concepts.

Thus, instead of applying the robust optimal control framework with time as the independent variable, we will adopt a formulation that is more consistent with current flight procedures and Flight Management System (FMS) technology, as they are not expected to change in the short-to-medium term (even after the introduction of free-routing). We will employ distance flown along the route (denoted as s) as the independent variable, because its initial and final value are the same in all scenarios that follow a unique route. As a consequence, the time t becomes a state variable and the new dynamical function can be obtained by dividing the time derivatives by $ds/dt = v_G$:

2Dt-s model

Independent variable:	s
State variables:	ϕ, λ, v, m, t
Control variables:	C_T, χ, ψ
Algebraic variables:	v_G

Differential equations:

$$\frac{d}{ds} \begin{bmatrix} \phi \\ \lambda \\ v \\ m \\ t \end{bmatrix} = \begin{bmatrix} \frac{\cos \psi}{(R_M(\phi) + h)} \\ \frac{\sin \psi}{(R_N(\phi) + h) \cos \phi} \\ \frac{T(C_T) - D(C_L)}{mv_G} \\ \frac{-f_c(C_T)}{v_G} \\ v_G^{-1} \end{bmatrix}$$

Algebraic equations:

$$\begin{bmatrix} v_G \cos \psi \\ v_G \sin \psi \end{bmatrix} = \begin{bmatrix} v \cos(\chi) + w_x \\ v \sin(\chi) + w_y \end{bmatrix}$$

with $C_L = \frac{2mg}{\rho v^2 S}$

All inequality constraints remain the same as in the untransformed system of differential-algebraic equations. Note that, by employing v_G as an algebraic variable (instead of a

function of wind, airspeed and heading), it will be computed only once at each node in the NLP iterations and produce sparser derivative matrices.

5.1.3 Trajectory Ensemble

As described in Section 4.4, the robust optimal control framework requires the formulation of the trajectory ensemble. An ensemble forecast contains a set of ensemble members, each one defining a different wind forecast (and, therefore, different functions w_x and w_y). If the ensemble contains N members, we define N scenarios, each one having weight $w_k = 1/N$ and the wind field that corresponds to the respective member; our stochastic quadrature rule is, therefore, a simple empirical average. We will write the compact form of the trajectory ensemble directly.

We choose to track the course ψ and the true airspeed v , i.e., the functions $\psi(s)$ and $v(s)$ are the same in every scenario (so we do not need to implement scenario-specific versions). As a consequence of the dynamics of the **2Dt-s** model, this implies that the evolution of the latitude ϕ is unique (as it only depends on the evolution of the unique variable ψ) and λ is also unique (as it only depends on ϕ and ψ). Therefore, the position variables act like tracked variables too, which is both relevant from the implementation point of view (because we do not need to create copies of them for each scenario) and a desired goal (since we want to obtain a unique route). We also define $a := dv/ds$ for practical purposes, in order to combine the derivative of v and its tracking condition in a single set of constraints. Its physical interpretation is the slope of the airspeed profile.

Taking advantage of these manipulations, we can define the dynamical system associated to the trajectory ensemble with the dynamical function:

$$\frac{d\mathbf{x}_E}{ds} = \frac{d}{ds} \begin{bmatrix} \phi \\ \lambda \\ v \\ t_1 \\ \vdots \\ t_N \\ m_1 \\ \vdots \\ m_N \end{bmatrix} = \begin{bmatrix} \cos(\psi)/(R_N + h) \\ \sin(\psi)/(R_N + h)/\cos\phi \\ a \\ 1/v_{G,1} \\ \vdots \\ 1/v_{G,N} \\ f_c T_1/v_{G,1} \\ \vdots \\ f_c T_N/v_{G,N} \end{bmatrix} \quad (5.2)$$

with the control vector:

$$\mathbf{u}_E = \begin{bmatrix} a & \psi & T_1 & \dots & T_N & \chi_1 & \dots & \chi_N \end{bmatrix}^T \quad (5.3)$$

the equality constraints:

$$\begin{bmatrix} v_{G,1} \cos(\psi) \\ \vdots \\ v_{G,N} \cos(\psi) \\ v_{G,1} \sin(\psi) \\ \vdots \\ v_{G,N} \sin(\psi) \\ a \cdot v_{G,1} \\ \vdots \\ a \cdot v_{G,N} \end{bmatrix} = \begin{bmatrix} v \cos(\chi_i) + w_{y,1} \\ \vdots \\ v \cos(\chi_i) + w_{y,N} \\ v \sin(\chi_i) + w_{x,1} \\ \vdots \\ v \sin(\chi_i) + w_{x,N} \\ (T_1 - D_1)/m_1 \\ \vdots \\ (T_N - D_N)/m_N \end{bmatrix} \quad (5.4)$$

and the inequality constraints:

$$\begin{aligned} v_{CAS, stall} &\leq v_{CAS}(v) \leq v_{CAS, max}, \\ M(v) &\leq M_{max}, \\ \left. \begin{aligned} T_{idle}(v) &\leq T_k \leq T_{max} \\ 0 &\leq v_{G,k} \end{aligned} \right\} \forall k \in \{1, \dots, N\} \end{aligned} \quad (5.5)$$

5.1.4 Formulation of the Discretized Problem

We will now finish the formulation of the Robust Flight Planning problem by adding an optimization criterion and boundary constraints. We define two additional scalar optimization variables: the *earliest arrival time* $t(s_f)_{\min}$ and the *latest arrival time* as $t(s_f)_{\max}$, which together define a “window of arrival”. We also define two user-specified parameters: the Cost Index CI^4 , which represents the user’s preferences for reduced flight time versus reduced fuel burn, and the “dispersion penalty” DP , which represents the user’s preferences for increased predictability versus average efficiency. We define the cost functional as:

$$J_{RFP} = -\frac{1}{N} \sum_{i=1}^N m_i(s_f) + CI \cdot \frac{1}{N} \sum_{i=1}^N t_i(s_f) + DP \cdot (t(s_f)_{\max} - t(s_f)_{\min}) \quad (5.6)$$

and the boundary conditions:

$$\begin{aligned} (\phi, \lambda, v)(0) &= (\phi_0, \lambda_0, v_0), \\ (\phi, \lambda, v)(s_f) &= (\phi_f, \lambda_f, v_f), \\ \left. \begin{aligned} t_k(0) &= 0 \\ m_k(0) &= m_0 \\ t_{f, \min} &\leq t_k(s_f) \leq t_{f, \max} \end{aligned} \right\} \forall k \in \{1, \dots, N\}. \end{aligned} \quad (5.7)$$

⁴Whenever not explicitly noted, the cost index CI is given in $(\$/hr)/(cents/lb)$.

The Problem (DROCT) associated to the Robust Flight Planning problem can now be defined as:

$$\left. \begin{array}{ll} \text{minimize} & J_{RFP} \text{ (5.6)} \\ \text{subject to} & \left. \begin{array}{l} \text{dynamical equation (5.2)} \\ \text{equality constraints (5.4)} \\ \text{inequality constraints (5.5)} \\ \text{boundary conditions (5.7)} \end{array} \right\} \end{array} \right\} \quad (\text{DRFP})$$

Note that, with minimal changes to the cost functional and boundary conditions, similar problems can be solved. For example, finding the fuel-optimal trajectory that arrives in a specified time window or identifying the required fuel that needs to be loaded on the aircraft before departure.

5.1.5 Dynamic Airspeed Adjustment

In Section 5.1.4, we have described a formulation whose solution generates a flight plan that is consistent with current flight procedures and existing FMS and Trajectory Prediction (TP) technology: a lateral path (that can be discretized to a sequence of segments delimited by waypoints) and an airspeed profile (that can also be discretized to a sequence of airspeeds associated to each leg). In this section, we consider a potential future concept where the airspeed is dynamically adjusted in response to time leads or lags in order to increase adherence to the predicted flyby times at each point in the trajectory, and the adjustment is done according to a feedback law that is incorporated to the flight plan. We will name this concept “dynamic airspeed adjustment” or DASA.

The proposed scheme is consistent with the concept of making small tactical speed changes⁵ in order to increase predictability. This idea has been explored as a potential future ATM system concept; for example, in [250] and related work. Under the “subliminal control” paradigm examined there, pilots or automated systems would implement tactical speed changes in the interval $[-6\%, 3\%]$ (under future technology) in order to reduce the risk of conflicts; this modification would be small enough that it would fall within the uncertainty observed by the air traffic controller (thus the “subliminal” label). In the present work, a speed change of a similar magnitude would be individually triggered by delays or leads following a pre-computed rule instead of a periodic and centralized sector-wide calculation. Therefore, the impact on conflicts would be indirect, through increased adherence to the scheduled flyby times (which would allow for earlier deconfliction).

⁵Note that, in this work, the rule implementing these speed changes is still computed at the pre-tactical planning stage.

The DASA law

We consider a simple control law⁶:

$$v(s) - \bar{v}(s) = K \cdot (t(s) - \bar{t}(s)), \quad (5.8)$$

where $\bar{v}(s)$ is a fixed airspeed schedule, $\bar{t}(s)$ is the expected flyby time at position s , and K is a constant. This scheme constitutes a control law analogous to a proportional regulator, where the airspeed increments or decrements are proportional to the accumulated time lead or lag with respect to the expected trajectory. Under this law, the pilot or the FMS would change the planned airspeed according to the delay or lead compared to the scheduled times. We will optimize this law jointly with the flight plan at the planning stage (i.e., the gain K will be a result of the optimization process).

DASA formulation

In Section 4.6, we described the implementation of feedback policies of the form $\mathbf{u} = \mathbf{K}_p(\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{p}, t)$ among other potential alternative specifications. We will employ one of such variations, as the variable that is set by a feedback law in this concept is the airspeed, which is a state and not a control. Note that the airspeed dynamics will still be taken into account in the optimization problem.

In order to cast the DASA law into this framework, we define a parameter vector composed by the gain K , i.e., $\mathbf{p} = [K]$; and the reference variables \bar{v} and \bar{t} (i.e. $\mathbf{y}(s) = [\bar{v} \ \bar{t}]^T$). Both K and \bar{v} are left for the NLP solver to find; however, in order to ensure that \bar{t} does indeed represent the average flyby time, we will also introduce a constraint relating it to the state variables:

$$\bar{t} = \frac{1}{N} \sum_{k=1}^N t_k \quad (5.9)$$

We can now integrate the DASA law into the framework described in Section 5.1.4. Since the airspeed will now be specific to each ensemble member, we do not collapse it into a single variable; instead, we let v_1, \dots, v_N be the member-specific airspeeds with the associated dynamic function:

$$\frac{dv_k}{ds} = (T_k - D_k)/(m_k \cdot v_{G,k}), \quad k \in \{1, \dots, N\} \quad (5.10)$$

We replace the tracking condition on the airspeed by the algebraic condition implementing the control law:

$$v_k - \bar{v} = K \cdot (t_k - \bar{t}), \quad k \in \{1, \dots, N\} \quad (5.11)$$

⁶Note that v is not a control variable, but the dynamics of airspeed tracking would again happen at a timescale of much smaller characteristic time than that of the optimal control problem.

The ensemble dynamical equation is now:

$$\frac{d\mathbf{x}_E}{ds} = \frac{d}{ds} \begin{bmatrix} \phi \\ \lambda \\ v_1 \\ \vdots \\ v_N \\ t_1 \\ \vdots \\ t_N \\ m_1 \\ \vdots \\ m_N \end{bmatrix} = \begin{bmatrix} \cos(\psi)/(R_N + h) \\ \sin(\psi)/(R_M + h)/\cos\phi \\ (T_1 - D_1)/(m_1 \cdot v_{G,1}) \\ \vdots \\ (T_N - D_N)/(m_N \cdot v_{G,N}) \\ 1/v_{G,1} \\ \vdots \\ 1/v_{G,N} \\ f_c T_1/v_{G,1} \\ \vdots \\ f_c T_N/v_{G,N} \end{bmatrix} \quad (5.12)$$

The equality constraints now include the control law:

$$\begin{bmatrix} \bar{t} \\ v_{G,1} \cos(\psi) \\ \vdots \\ v_{G,N} \cos(\psi) \\ v_{G,1} \sin(\psi) \\ \vdots \\ v_{G,N} \sin(\psi) \\ v_1 - \bar{v} \\ \vdots \\ v_1 - \bar{v} \end{bmatrix} = \begin{bmatrix} N^{-1} \sum_{k=1}^N t_k \\ v \cos(\chi_i) + w_{y,1}(\phi, \lambda, t) \\ \vdots \\ v \cos(\chi_i) + w_{y,N}(\phi, \lambda, t) \\ v \sin(\chi_i) + w_{x,1}(\phi, \lambda, t) \\ \vdots \\ v \sin(\chi_i) + w_{x,N}(\phi, \lambda, t) \\ K \cdot (t_1 - \bar{t}) \\ \vdots \\ K \cdot (t_N - \bar{t}) \end{bmatrix} \quad (5.13)$$

The inequality constraints are now all scenario-specific:

$$\left. \begin{aligned} v_{CAS, stall} &\leq v_{CAS}(v_k) \leq v_{CAS, max} \\ M(v_k) &\leq M_{max} \\ T_{idle}(v_k) &\leq T_k \leq T_{max} \\ 0 &\leq v_{G,k} \end{aligned} \right\} \forall k \in \{1, \dots, N\} \quad (5.14)$$

and the boundary conditions now depend on \bar{v} instead of v (note that the initial airspeeds are already equated to $\bar{v}(0)$ by the DASA law):

$$\begin{aligned} (\phi, \lambda, \bar{v})(0) &= (\phi_0, \lambda_0, v_0), \\ (\phi, \lambda, \bar{v})(s_f) &= (\phi_f, \lambda_f, v_f), \\ \left. \begin{aligned} t_k(0) &= 0 \\ m_k(0) &= m_0 \\ t_{f,\min} &\leq t_k(s_f) \leq t_{f,\max} \end{aligned} \right\} \forall k \in \{1, \dots, N\}. \end{aligned} \quad (5.15)$$

Therefore, the uncertainty-discretized robust optimal control with tracking and feedback associated to the DASA problem can be defined as:

$$\left. \begin{aligned} \text{minimize} \quad & J_{RFP} \quad (5.6) \\ \text{subject to} \quad & \text{dynamical equation (5.12)} \\ & \text{equality constraints (5.13)} \\ & \text{inequality constraints (5.14)} \\ & \text{boundary conditions (5.15)} \end{aligned} \right\} \quad (\text{DASA})$$

5.2 Results

We will now proceed to illustrate the method with a flight planning scenario. We consider a cruise flight from the vertical of New York to the vertical of Lisbon at FL380. The aircraft is a wide-body twin-engine jet, modeled according to the BADA 4 specification and with an initial mass of 200 tons. Initial speed is set to 240 m/s and final speed is set to 220 m/s; for the constant airspeed case, the airspeed is fixed to 240 m/s. In order to facilitate visualization and understanding, we employ a static weather picture; nevertheless, the proposed methodology allows for the usage of dynamic weather forecasts. We employ the 6-hours-lead-time forecast from the Météo France PEARP ensemble for the 20th of January 2016 at the pressure level of 200 hPa. In this forecast, an area of high uncertainty can be identified, centered around 5 longitude degrees East of the Azores archipelago. Under these wind conditions, the optimal deterministic flight plan takes the aircraft straight through this region of high uncertainty; thus, it represents a useful scenario for our purposes.

We will consider three different variants of the analysis. In Section 5.2.1, we consider a fixed airspeed scenario (equivalent to a routing algorithm). Then, we will consider the full horizontal dynamics and allow the algorithm to optimize the airspeed profile as well in Section 5.2.2. Finally, in Section 5.2.3 we study the dynamic airspeed management concept proposed in Section 5.1.5.

We employ the Python bindings for the CasADi [251] library for NLP modeling and interfacing and the NLP solver IPOPT [77]. Table 5.1 shows the size of the resulting

nonlinear optimization problems after direct transcription with 80 time nodes. NLP CPU time varies depending on the type of problem and the quality of the initial guess, ranging from a few seconds in the simplest fixed-airspeed cases to a few minutes in the more complex DASA cases.

#	Constant TAS	Variable TAS	DASA
NLP variables	11144	14104	16912
Equality constraints	11039	14016	16782
Inequality constraints	7565	11437	16877
Number of nonzeros in ...			
Eq. constr. Jacobian	57271	88189	115945
Ineq. constr. Jacobian	7849	11673	17113
Lagrangian Hessian	18616	44496	60656

Table 5.1: Problem size.

For the remainder of the Chapter, the units of the DP and CP parameters are understood to be in the corresponding International System units (i.e. kg/s and kg/m) whenever they are not mentioned.

5.2.1 Fixed Airspeed

Figure 5.2 displays the geographical routes for different values of the dispersion penalty DP. We also plot regions of higher uncertainty, which we have defined as $\sqrt{\sigma_u^2 + \sigma_v^2}$, with σ_u and σ_v being the standard deviation of the u and v components of wind across different members. It can be seen that routes computed with higher DP tend to avoid the high uncertainty zone in the Atlantic in order to increase predictability, at the cost of taking a more indirect route that is longer on average.

Figure 5.3 shows the evolution of the groundspeed and the time lead or lag for different DP settings. Note that each individual line represents the trajectory that corresponds to an ensemble member and the time lead or lag is measured with respect to the average flyover time at each point in the route. The groundspeed is different for each member despite sharing the same airspeed schedule; naturally, this is due to the spread in wind speeds and directions. It can be seen that the spread in the ensemble times and ensemble groundspeeds increases markedly when the aircraft crosses the area of high uncertainty near the destination airport (see Figure 5.2). However, as the DP setting is increased and the aircraft avoids the areas with highest uncertainty, the spread in groundspeeds is reduced and, therefore, the time leads or lags are closer together on arrival.

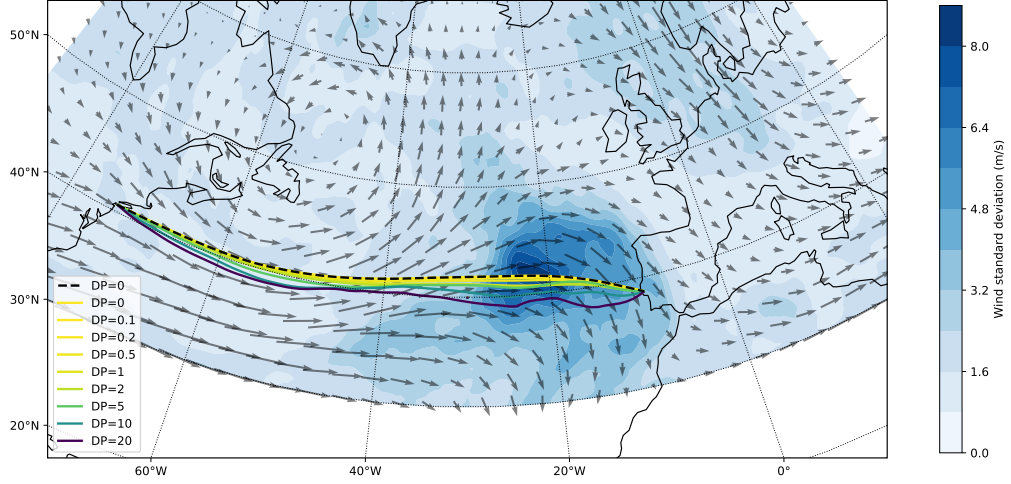


Figure 5.2: [Fixed airspeed] Optimal trajectories from New York to Lisbon, for values of DP from 0 kg/s to 20 kg/s. Higher brightness in the trajectory color indicates lower values of DP.




Figure 5.3: Groundspeeds and time leads or lags for different DP settings. Time leads and lags are defined with respect to the average trajectory.

Finally, Figure 5.4 shows the Pareto frontier of the problem, obtained by solving problems with different penalties (from $DP = 0$ kg/s to $DP = 20$ kg/s). For the maximum average efficiency case ($DP = 0$), the time dispersion at the final fix is just above 6 minutes, which is almost halved by the usage of a higher predictability flight plan ($DP = 20$ kg/s). This comes, however, at the cost of flying a longer route that takes 10 additional minutes of flight time (corresponding to more than 700 kg of extra fuel burnt). Nevertheless, a more modest improvement in uncertainty, from 6 minutes to 4.5 minutes, can be achieved at a tenth of the cost (around a minute of extra flight time).




Figure 5.4: Pareto frontier of the problem.

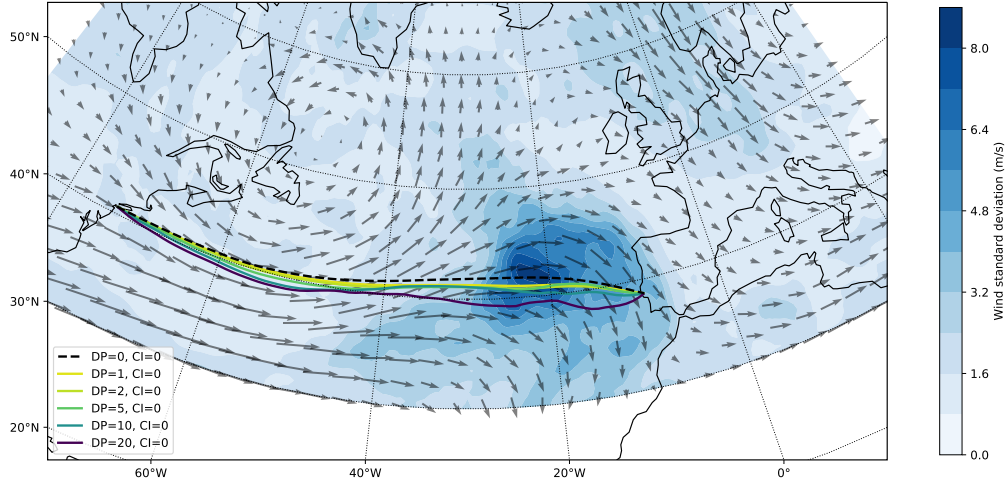


Figure 5.5: [Variable airspeed] Optimal trajectories from New York to Lisbon, for values of DP from 0 kg/s to 20 kg/s and $CI = 0$ kg/s. Higher brightness in the trajectory color indicates lower values of DP.

5.2.2 Variable Airspeed

We now proceed to consider variable airspeed profiles. In this case, the cost index (CI) becomes relevant, as it will have a first-order influence on the average airspeed (with higher CI settings leading to higher airspeed profiles).

Figures 5.5 and 5.6 show the geographical path of the optimized flight plans for $CI = 0$ kg/s and $CI = 5$ kg/s. As in the constant airspeed case, the horizontal profiles of the trajectories that place a higher weight on reducing uncertainty are farther from the high uncertainty zone; however, at a higher cost index setting the trajectories will deviate by a smaller amount due to the higher relative importance of the average flight duration objective. Figures 5.7 and 5.8 show the geographical path of the optimized flight plans for varying CI values at fixed values of the dispersion penalty. In the $DP = 0$ case, it can be observed that there is a negligible influence of the cost index parameter in the route: it's only influencing the airspeed profile, as we shall see shortly. However, at a high DP value there is an interaction between the two objectives, and we can see again that higher CI values will lead to solutions that are closer to the deterministic flight plan.




Figure 5.6: [Variable airspeed] Optimal trajectories from New York to Lisbon, for values of DP from 0 kg/s to 20 kg/s and CI = 5 kg/s. Higher brightness in the trajectory color indicates lower values of DP.




Figure 5.7: [Variable airspeed] Optimal trajectories from New York to Lisbon, for values of CI from 0 kg/s to 5 kg/s and DP = 0 kg/s. Higher brightness in the trajectory color indicates lower values of DP.

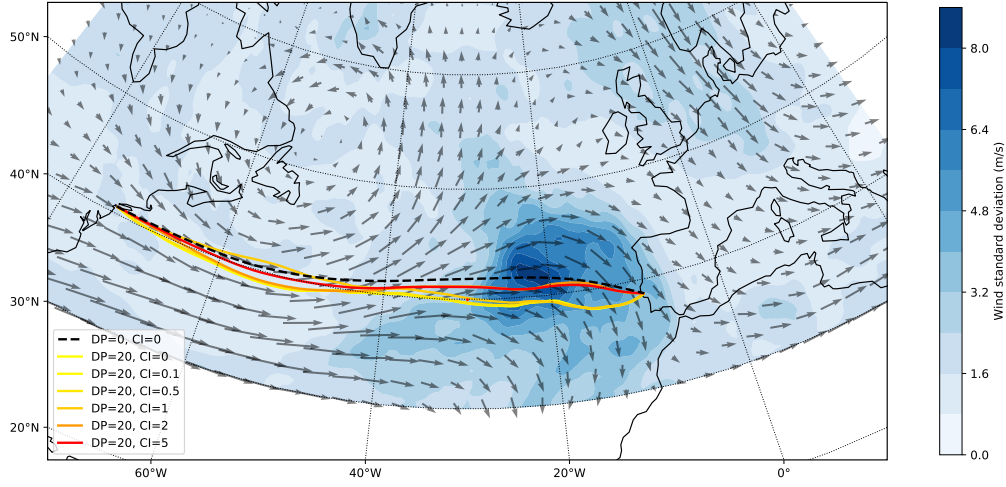


Figure 5.8: [Variable airspeed] Optimal trajectories from New York to Lisbon, for values of CI from 0 kg/s to 5 kg/s and DP = 20 kg/s. Higher brightness in the trajectory color indicates lower values of DP.

Figure 5.9 shows the speed profiles for different values of the CI and DP parameters. As expected, trajectories optimized for a higher CI tend to feature higher speeds. However, there is another relevant difference between trajectories: as the penalty in uncertainty grows, the optimal airspeed *when the aircraft crosses uncertain regions* increases. We attribute this result to the idea that, by flying at a higher speed, the relative importance of the wind on the groundspeed is reduced⁷. We show the groundspeeds and time leads or delays of three scenarios in Figure 5.10; the pattern is similar to the constant airspeed one (see Figure 5.3), with the addition of the airspeed profile.

These airspeed changes reduce uncertainty more efficiently than direct avoidance, as it can be observed in Figure 5.11. For a given CI, there are two regions in the locus of the solutions for different DP values. Reducing arrival time dispersion in the left branch is cheaper (around 100 kg when reducing the arrival window size from 5 minutes to 4.5 minutes) than in the right branch (around 300 kg when reducing from 4 minutes to 3.5 minutes).

⁷Consider the following illustrative example: the along-track wind can take one of two constant values, w_1 or w_2 , on a region where an aircraft flies a distance x at the airspeed v . The difference in arrival times is $\Delta t = x \frac{|w_1 - w_2|}{v^2 + (w_1 + w_2)v + w_1 w_2}$, which is decreasing on v .

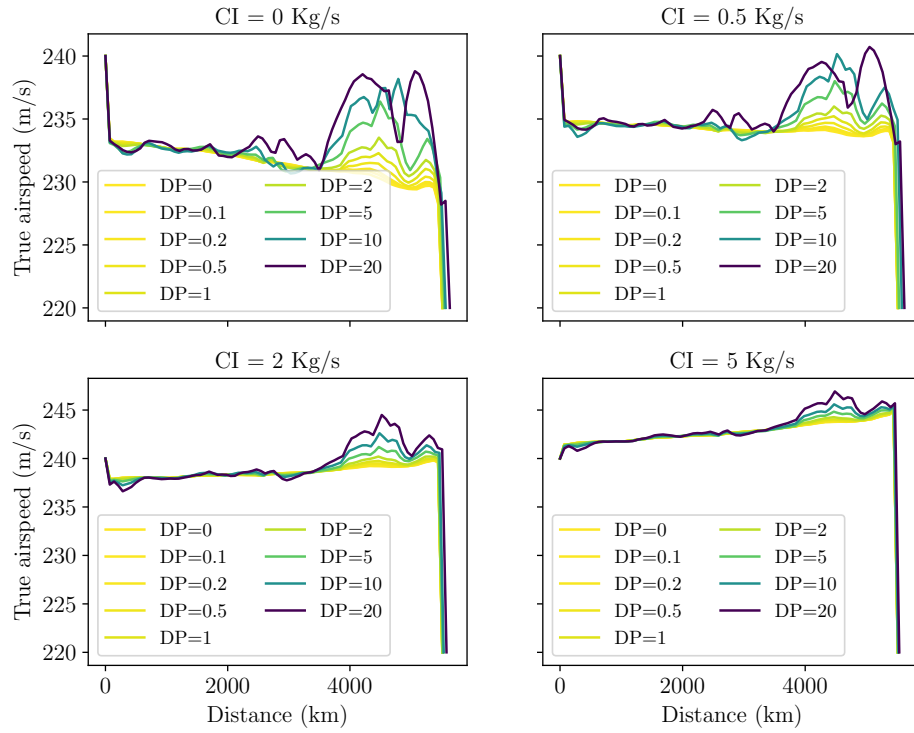


Figure 5.9: True Airspeed profiles for different DP and CI settings.




Figure 5.10: Groundspeeds and time leads or lags for different DP settings and $CI = 0$ kg/s. Time leads and lags are defined with respect to the average trajectory.

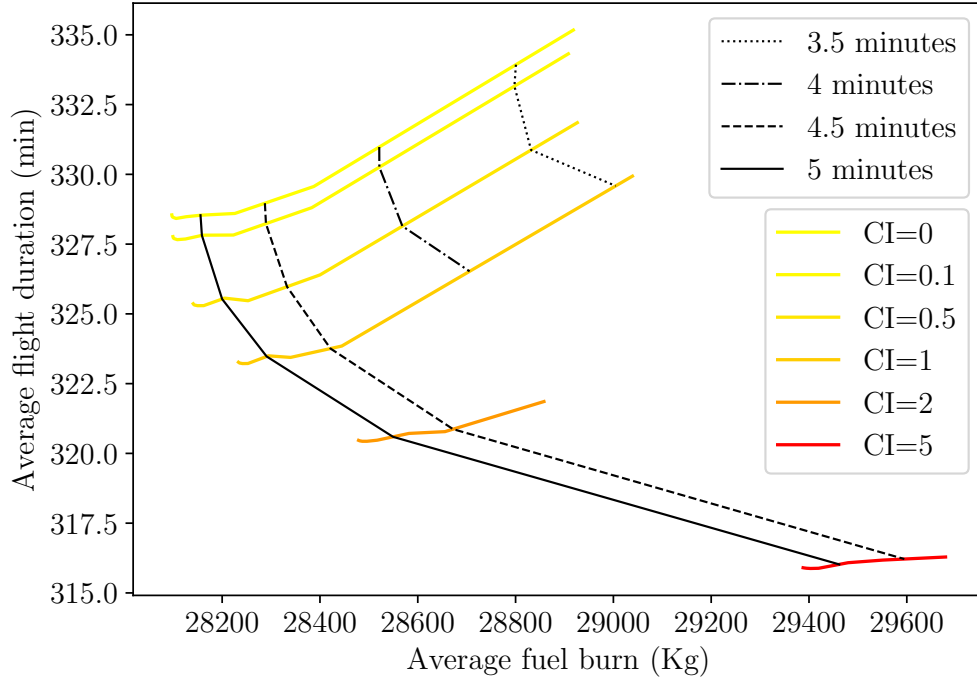


Figure 5.11: Pareto frontier and locus of the solutions of the variable airspeed profile case.

5.2.3 Dynamic Airspeed Adjustment (DASA)

Finally, we discuss the results of the Dynamic Airspeed Adjustment (DASA) concept introduced in Section 5.1.5. The results of the optimization suggest that, in this case, it is cheaper to reduce uncertainty by increasing the feedback strength K (and, therefore, absorbing some uncertainty in airspeed instead of groundspeed) rather than modifying the geographical path: Figure 5.12 shows that geographical paths are very similar.

Figure 5.13 shows the evolution of the state variables for different values of DP. The lowest predictability scenario has a similar profile to the analogous case in Figure 5.10; however, as DP increases, some of the wind uncertainty starts being compensated by airspeed variations instead of being fully passed onto the groundspeed. As a consequence of the TAS dispersion, the fuel burn shows more dispersion as arrival time uncertainty is decreased; we illustrate this result in Figure 5.14.




Figure 5.12: [DASA] Optimal trajectories from New York to Lisbon, for values of DP from 0 kg/s to 20 kg/s and $CI = 0$ kg/s. Higher brightness in the trajectory color indicates lower values of DP.

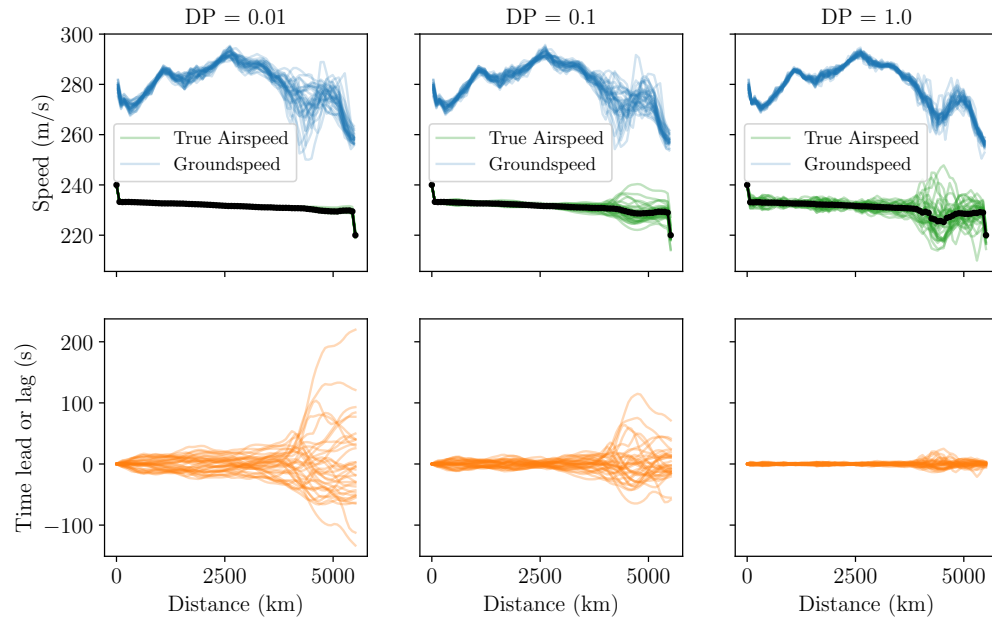


Figure 5.13: [DASA] Groundspeeds and time leads or lags for different DP settings and $CI = 0$ kg/s. Time leads and lags are defined with respect to the average trajectory.




Figure 5.14: Time uncertainty - fuel uncertainty tradeoff. The range of a variable is defined as the difference between its maximum value over all scenario and its minimum value.

Figure 5.15 shows the scheduled TAS profiles for different values of CI and DP. As it can be expected, the airspeed increases as CI increases, therefore putting more importance on reducing flight time than fuel burn. Another apparent feature of these profiles is that the reference airspeed decreases when crossing a region of high uncertainty instead of increasing, as in the variable airspeed case. This can be attributed to the result that higher predictability trajectories will now feature higher gains K (see Figure 5.17) and thus the airspeeds will be more spread when crossing a region of high uncertainty (recall that, under DASA, $\Delta v = K\Delta t$). Therefore, in order to prevent the airspeeds in the most unfavourable members from rising to a level that is either too inefficient or even outside the flight envelope, the reference airspeed needs to be lower.




Figure 5.15: True Airspeed profiles for different DP and CI settings.

Figure 5.16 shows that the Pareto frontiers for different values of the arrival time dispersion are very close to one another (compare with Figure 5.11): as discussed earlier, the reduction of uncertainty through dynamic airspeed adjustment is cheap enough (in average terms) that the impact on average efficiency is minimal. The main cost is now an increase in fuel burn dispersion (see again Figure 5.14).

Finally, we show the magnitude of the gains K in Figure 5.17. As expected, increasing DP has the effect of increasing the strength of the feedback in order to reduce the temporal divergence with respect to the reference schedule (otherwise, the gain will be low as deviating from the most efficient airspeed is inefficient in general). We can observe that an increase of the CI setting results in a lower gain: as average flight time is prioritized, uncertainty reduction is less important in relative terms.




Figure 5.16: Pareto frontier and loci of the solutions of the dynamic airspeed management case.




Figure 5.17: Gain K for varying DP and CI parameters.

Flight Planning under Convection

Convection represents an important meteorological source of uncertainty in flight plans, as discussed in Chapters 1 and 2. At the planning horizon, it is often impossible to forecast the specific location, timing and evolution of individual convective events with enough accuracy to produce highly predictable flight plans in convective scenarios. However, the necessary conditions for convection can be identified in probabilistic fashion at the planning stage, and this information can be integrated into the Robust Flight Planning framework presented in Chapter 5. Building and demonstrating such an extension is the purpose of the current chapter.

This chapter is divided in two parts:

- Section 6.1 explains the extension of the Robust Flight Planning framework for including convection as an uncertainty source.
- Section 6.2 demonstrates the described approach in a planning scenario under both constant airspeed and variable airspeed paradigms.

6.1 Concept

In the current chapter, we will develop an extension of the Robust Flight Planning framework introduced in Chapter 5. In a flight planning scenario, we design the planned aircraft trajectory around 1 to 3 hours before departure. Depending on the duration of the flight, the elapsed time between the moment that the flight plan is produced and the instant where the aircraft would encounter convective weather may range between 2-3 hours to more than 10 hours. At these lead times, forecasting and nowcasting systems are not able to produce estimates of the future position and shape of convective cells with enough accuracy to design reliable flight plans, as explained in Section 3.3. Thus, it is assumed that, if the aircraft encounters adverse weather, it will need to deviate from the planned trajectory in order to avoid it. This deviation leads to higher cost trajectories in average (as the avoidance routes are usually longer) and more unpredictable trajectories. Uncertainty is amplified when taking into account multiple trajectories, as an aircraft may receive additional ATC advisories due to conflicts with other aircraft that are also trying to avoid the same convective cells.

However, even if we cannot determine with precision the future location and timing of convective triggers and individual cells, it is possible to identify regions where the necessary conditions for convection are forecasted to be met. We can employ this information to generate more efficient and predictable trajectories at the planning horizon.

We will model these conditions with the aid of the function $c_{t_p}(\phi, \lambda, t)$, introduced in Section 3.3.1. Given the latest EPS forecast available at planning time t_p , this function represents the fraction of the EPS members forecasting that two indicators of convective conditions (Total Totals and Convective Precipitation) will exceed their thresholds at point (ϕ, λ) and at time t . For convenience, we assume that the planning time t_p is unambiguous and denote the function $c(\phi, \lambda, t)$.

Assume now that an aircraft has a trajectory described by the time domain $\mathcal{T} = [t_0, t_f]$, the groundspeed profile $v_g(t)$, and a lateral profile described by $(\phi(t), \lambda(t))$. We define the exposition to convection EC as

$$EC = \int_{t_0}^{t_f} c(\phi(t), \lambda(t), t) v_g(t) dt \quad (6.1)$$

Using the position variable s , defined as in Chapter 5, we can write this definition more concisely as:

$$EC = \int_{s_0}^{s_f} c(\phi(s), \lambda(s), t(s)) ds \quad (6.2)$$

We now assume that the expected cost of crossing convective zones, in terms of additional fuel burn (Δm_c), flight time (Δt_c) and unpredictability ($\Delta \sigma_c$) due to avoidance maneuvers, is linear on EC . In other words, if one possible route overflies a region with

high likelihood of convection for 200 km and another possible route overflies it only for 100 km, we assume that the expected costs of the former route will be twice the costs of the latter. These relationships will be modeled by the parameters β_m , β_t and β_σ , i.e.

$$\mathbb{E}[\Delta m_c] = \beta_m EC \quad (6.3)$$

$$\mathbb{E}[\Delta t_c] = \beta_t EC \quad (6.4)$$

$$\Delta \sigma_c = \beta_\sigma EC \quad (6.5)$$

Assume now that the flight planner has time-fuel preferences described by a Cost Index (CI) and average cost - uncertainty cost preferences described by a Dispersion Penalty (DP) parameter, i.e., the cost functional (without the effect of convection) is given by

$$J_{RFP} = -\mathbb{E}[m(t_f)] + CI \cdot \mathbb{E}[t_f] + DP \cdot \sigma[t_f] \quad (6.6)$$

where $\sigma[t_f]$ is a measure of the uncertainty in the arrival time (such as the range or a multiple of the standard deviation). The increase in total cost due to the average effect of convection ΔJ_c can now be modeled as:

$$\Delta J_c = \mathbb{E}[\Delta m_c] + CI \cdot \mathbb{E}[\Delta t_c] + DP \cdot \Delta \sigma_c = \underbrace{(\beta_m + \beta_t CI + \beta_\sigma DP)}_{\text{Convective penalty}} EC \quad (6.7)$$

If we now define the *convective penalty* parameter $CP(CI, DP) = \beta_m + \beta_t CI + \beta_\sigma DP$, we can extend the cost functional of the robust flight planning problem to include the effect of convection. The new cost functional is given by:

$$J_{RFP_c} = -\mathbb{E}[m(t_f)] + CI \cdot \mathbb{E}[t_f] + DP \cdot \sigma[t_f] + CP(CI, DP) \cdot EC \quad (6.8)$$

with the rest of the problem formulation remaining identical to the one in Chapter 5.

Estimating the statistical relationship between the exposition to convection and the deterioration of the objectives of interest (i.e. determining the values of the coefficients β_m , β_t and β_σ) is a challenging research problem that falls outside of the scope of this thesis, as it requires either a complex statistical analysis of historical traffic patterns and deviations (which may not generalize to an evolving concept of operations in the future) or a simulation-based analysis that demands realistic models for the stochastic generation of convective thunderstorms, as well as a realistic model of the processes of replanning avoidance routes and traffic deconfliction. Instead, for the remainder of the chapter we will assume that the convective penalty CP is an independent parameter chosen by the flight planner, in the same way as the cost index or the dispersion penalty.

6.2 Results

In order to understand how the planned trajectories produced by this framework manage to improve the desired objectives, we now proceed to apply the proposed variation of the robust flight planning methodology to a flight planning problem for different settings of the optimization parameters.


We consider a wide-body twinjet aircraft, modeled according to the BADA4 specification, flying from the vertical of New York (-73.8 deg, 40.6 deg) to the vertical of Argel (3.2 deg, 36.7 deg) at constant barometric altitude 200hPa and Mach 0.82. Initial mass has been set to 200 tons. We use a forecast for a pressure of 200 hPa 9 hours in advance for the 19th of December 2016 from the ECMWF ensemble, elaborated by the European Center for Medium-Range Weather Forecasts (ECMWF)¹ with 51 members. As in Chapter 5, we rely on a constant weather picture for ease of exposition and analysis. We rely on the CasADi library[251] as NLP interface [252] and IPOPT [77] as NLP solver.

We will perform two analyses. In first place, we study the constant airspeed (with $TAS = 240$ m/s) routing problem (Section 6.2.1) by considering the 2Dt model (see Section 3.1.4) and, afterwards, we demonstrate the method for the variable airspeed problem (Section 6.2.2), based on the 2Dt-m-e model.


6.2.1 Constant Airspeed

Figure 6.1 displays the geographical routes for different values of DP and CP . It can be seen that routes computed with higher DP (setting $CP = 0$) tend to avoid the high uncertainty zone in the North Atlantic in order to increase predictability, at the cost of taking a more indirect route that is longer on average. It can be also observed that routes computed with higher CP (setting $DP = 0$) tend to reduce the exposure to convective risk zones, again at the cost of taking a more indirect route. Four selected flight plans produced with different CP - DP pairs are shown in Figure 6.2.

¹<http://apps.ecmwf.int/datasets/>



(a) Optimal paths with different DP values (CP=0).



(b) Optimal paths with different CP values (DP=0).

Figure 6.1: Optimal trajectories for DP/CP values. Higher brightness in the trajectory color indicates higher values of the penalty. Color contour scale indicates wind uncertainty characterized as $\sqrt{\sigma_u^2 + \sigma_v^2}$, with σ_u being the standard deviation of the u component of wind across different members and σ_v analogous for the v -component. Dashed regions indicate regions of convective exposure.




Figure 6.2: Optimal paths for different DP-CP pairs.

The evolution of time spreads and convective exposure over the flown distance is presented in Figure 6.3. Comparing the maximum average efficiency trajectory (corresponding to $DP = 0$ and $CP = 0$) with a more predictable trajectory, e.g., that of $DP = 6$ and $CP = 0$, it can be seen how the spread in the ensemble times increases slightly when the aircraft crosses the area of high uncertainty in the middle of the North Atlantic. Similarly, comparing the maximum average efficiency trajectory with another with less exposure to convection, e.g., that of $DP = 0$ and $CP = 0.01$, it can be readily seen in Figure 6.3 that the exposure (the integral, or area below the curve) is reduced. For the minimum average fuel case ($DP = 0$), the time dispersion at the final fix is around 200 seconds, whereas for the $DP = 6$ case, the time dispersion at the final fix is around 100 sec. In other words, around one and a half minutes reduction in *time uncertainty* could be achieved by flying a more predictable trajectory ($DP = 6$). This improvement comes at the cost of extra flight time and associated fuel burnt on average.

Figure 6.4 presents different Pareto frontiers (for different pre-set average flight times) and shows different possible solutions with trade-offs between time spread, accumulated convection, and arrival time (directly related to consumption). In other words, and in order to display the trade-off effects of dispersion and convection, we have solved the problem with an additional constraint that enforces the average arrival time to be a pre-set one. For the sake of illustration, we present result for average final times set equal to 395 minutes, 400 minutes, and 405 min.

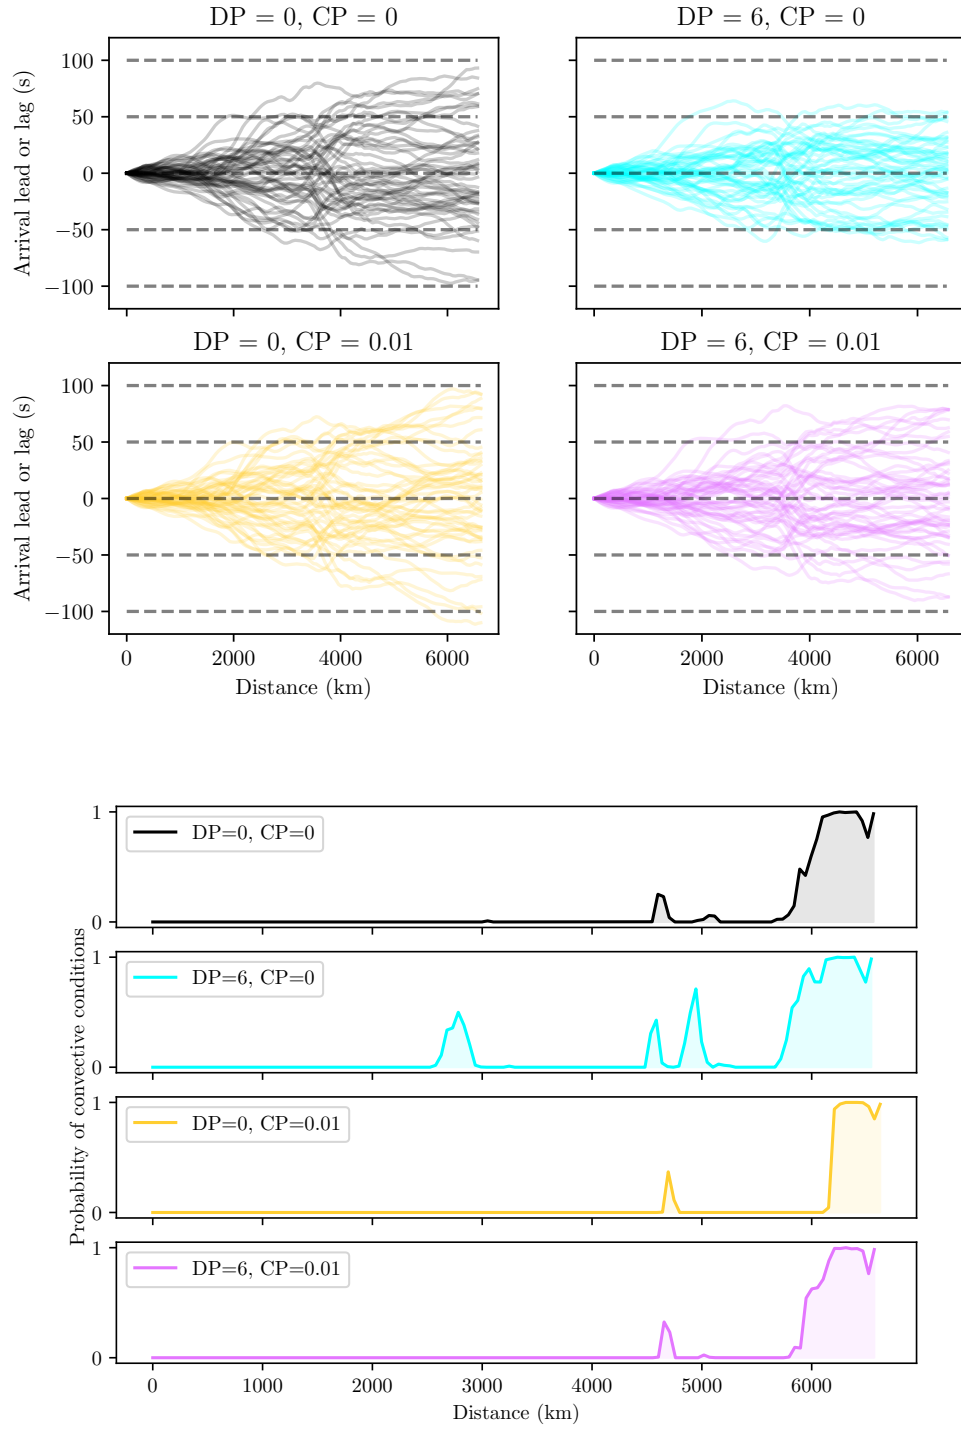




Figure 6.3: Evolution of time spreads and convective exposure.



(a) Solution loci of the problem.



(b) Time spreads and accumulated convection. Each segment represents the arrival time window

Figure 6.4: Trade-offs between variables (spread and convection exposure) for three different arrival times; ekm denotes "equivalent kilometer"


Quantitate values of flight dispersion and exposure to convection can be readily extracted out of Figure 6.4. Consider the flight constraint to reach the final fix at an average flight time of 400 minutes (orange line in Figure 6.4.a): for a dispersion of roughly 1.5 minutes the exposure to convective areas would be of roughly 1000 e-km (equivalent kilometers, i.e., kilometers flown at probability of convection equal to one). If ones want to reduce this exposure to convection to roughly 500 e-km, it comes at a cost of an extra 1.5 minutes (3 in total) of flight dispersion. The same information, however differently represented, has been included in Figure 6.4.b. Nevertheless, notice that these numbers correspond to this particular case (route, day, and weather forecasts) and may or may not be representative of characteristic costs and benefits. Further studies should assess these quantities in a more systematic fashion.

6.2.2 Variable Airspeed

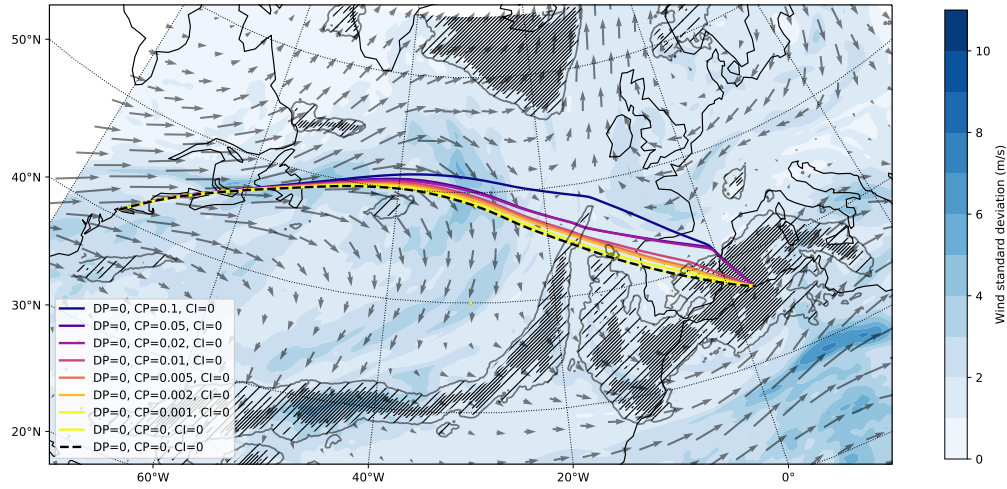
We now proceed to relax the assumption of constant airspeed and allow the algorithm to optimize the airspeed schedule too. In this case, fuel consumption is no longer tightly related to flight time, and we employ the Cost Index (CI) parameter along the DP and CP parameter in order to represent the preferences of the airline.

Figures 6.5 and 6.7.a) represent the geographical paths of the variable airspeed profile case by changing each of the preference parameters. The behaviour of the algorithm for varying DP and CP is very similar to the one in the constant airspeed case (cf. Figure 6.1), while the geographical path does not change significantly when modifying the cost index.

The airspeed profiles are shown in Figure 6.6 and Figure 6.7.b). It can be observed that the main influence in the airspeed profile is the variation in the CI, which increases, as expected, the airspeed values along the whole trajectory. The CP parameter does not have a big influence on the shape of the airspeed profile, which changes as the lateral path shifts to a different location. Finally, the flight plans with high DP exhibit the same phenomena as the results in Chapter 5, where the airspeed changes locally depending on the uncertainty of the winds at each point in the trajectory, with higher airspeeds when the aircraft crosses higher uncertainty zones. This influence is moderated by higher CI settings (compare the dashed lines with the solid lines in Figure 6.6.a).




(a) Optimal paths with different DP values (CP = 0).




(b) Optimal paths with different CP values (DP = 0).

Figure 6.5: Optimal trajectories for different CP/DP/CI values. Color contour scale indicates wind uncertainty characterized as $\sqrt{\sigma_u^2 + \sigma_v^2}$, with σ_u being the standard deviation of the u component of wind across different members and σ_v analogous for the v -component. Dashed regions indicate regions of convective exposure.

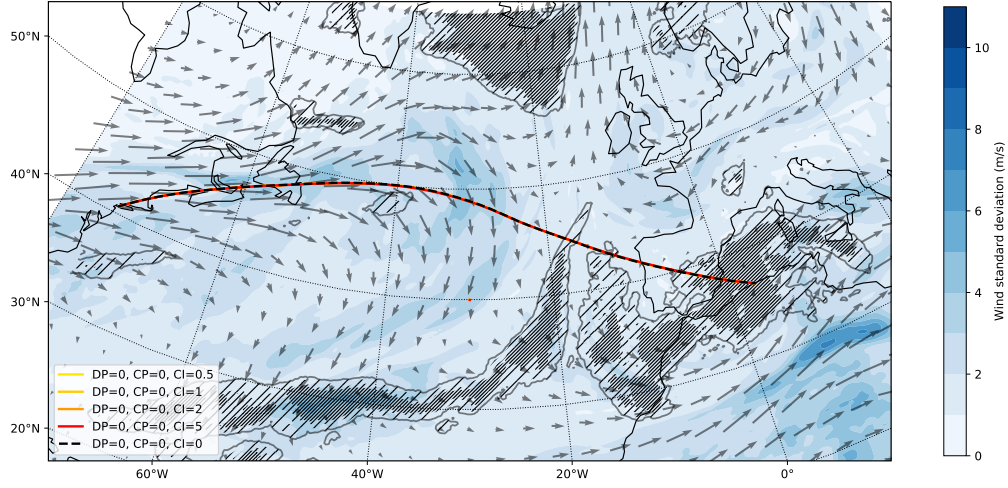


(a) Optimal airspeed schedules with different DP values (CP=0).




(b) Optimal airspeed schedules with different CP values (DP=0).

Figure 6.6: Airspeed profiles for varying cost functional parameters.



(a) Optimal paths with different CI values (DP=CP=0).



(b) Optimal airspeed schedules with different CI values (DP=CP=0).

Figure 6.7: Geographical path and airspeed profile for varying cost index.

In order to compare the joint effect of changes in multiple parameters, we now highlight 8 trajectories obtained with different sets of parameters, resulting from the combination of two values for each of the parameters. Figure 6.8 illustrates the optimal paths of these highlighted trajectories. The clearest interaction takes place when CI increases, since it “pulls back” the trajectories with higher CP and DP values towards the optimal trajectory on average; this can be explained by the fact that the dispersion and convection terms lose relative weight as CI increases. The other interaction that can be observed is that, when DP is increased at high CP, the optimal trajectory shifts even more towards the North instead of being pulled South towards the “high DP” solution, which illustrates the nonlinear effects caused by the nonconvexity of the uncertainty and convection fields.

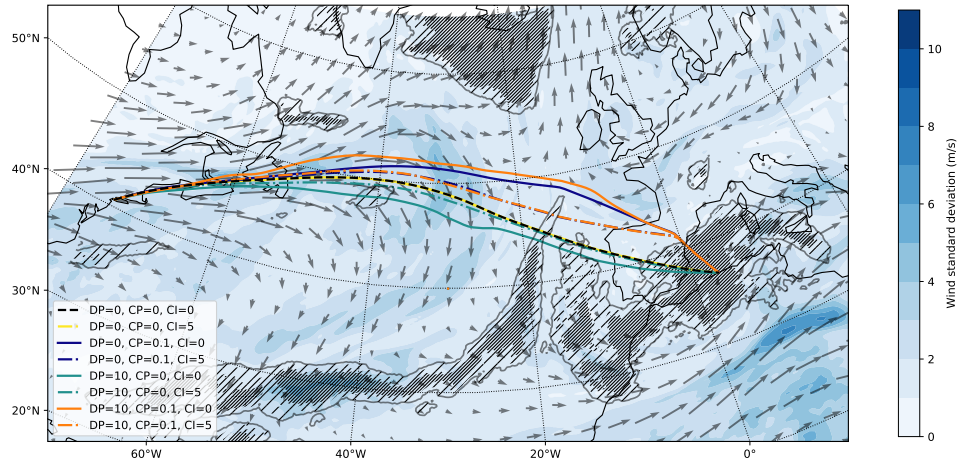


Figure 6.8: Optimal path for the featured trajectories.

Figure 6.9 shows the respective airspeed profiles. Clearly, at high CI values the airspeed profiles also become high and very similar. The airspeed profiles for the low CI trajectories show more variability between them, which should be expected as they fly different lateral paths and thus face different winds. Finally, the “speed-up on higher wind uncertainty” effect is also clearly present for the trajectories with high DP but low CI. Again, a higher CI setting has a regularizing effect, leading to smoother airspeed profiles as the relative influence of the wind and its uncertainty is reduced.






Figure 6.9: Airspeed profiles for the featured trajectories.

We show the evolution of the time dispersion along the trajectory in Figure 6.10. In this chart, we plot the flyby times at each point in the trajectory for each of the 50 ensemble members. It can be observed that higher DP settings lead to tighter uncertainty profiles, just as expected; however, this trend is moderated by the increase in the other parameters (particularly the cost index), since the relative importance of reducing dispersion due to wind uncertainty decreases.

Figure 6.11 illustrates the probability of convection $c(\phi, \lambda, t)$ along the route (compare with Figure 6.8 to see how the segments of high probability of convective conditions correspond to the crossing of the highlighted regions). Again, it is clear that a higher CP parameter emphasizes avoidance of areas of likely convective conditions; the trajectories generated with CP=0.1 only encounter (unavoidable) risk of convection at the end of the trajectory, near the destination airport.




(a) CI=0




(b) CI=5

Figure 6.10: Evolution of time spreads for the different ensemble members.



(a) $CI=0$



(b) $CI=5$

Figure 6.11: Evolution of the convective exposure.

Finally, we illustrate the trade-offs that are available to the flight planner in Figure 6.12. In the top figure, we illustrate the trajectory options, in terms of the average performance, that can be obtained for selected levels of uncertainty (both from wind and convection). In first place, it can be observed that the cost of reducing uncertainty due to wind has a nonlinear shape, with the level curves being relatively closer between 200 s and 180 s of arrival time range than between 180 s and 140 s. This is consistent with the results of Section 5.2: using airspeed variations allows the algorithm to cheaply improve predictability up to a point where it saturates or becomes inefficient; afterwards, the more expensive lateral path modifications need to be employed and the cost of reducing the arrival time window by 10 seconds is 100 to 200 kg in average or 3 to 6 minutes. Reducing exposure to convection appears to have a more homogeneous penalty in terms of cost, with the cost of encountering 50 km less of potential convective conditions being on the order of several minutes or several hundred kilograms.

In the bottom panel, we can observe the uncertainty that a trajectory will face for a given average performance. We can again observe the mentioned nonlinearity: at higher widths of the arrival time window, it is relatively cheap to decrease it (now in terms of exposition to convection at fixed $\mathbb{E}[t_f]$ and $\mathbb{E}[m_f]$), with the level curves being flatter at the right of the kink point. However, they become steeper at the left of this point. The mentioned trade-offs between exposition to convection and average performance can also be observed more clearly in this graph: at around 180 seconds of arrival window size, 250 kg of average fuel burn can reduce exposition to convection by 15 to 20 equivalent km in faster flights and by 30 to 40 equivalent km in later arrivals. In other words, at high Cost Index settings (faster flights) where flight time is more relatively important, predictability improvements are more expensive in terms of average fuel burn than they are at slower flights, as the algorithm has more “room” to tune the flight plan.

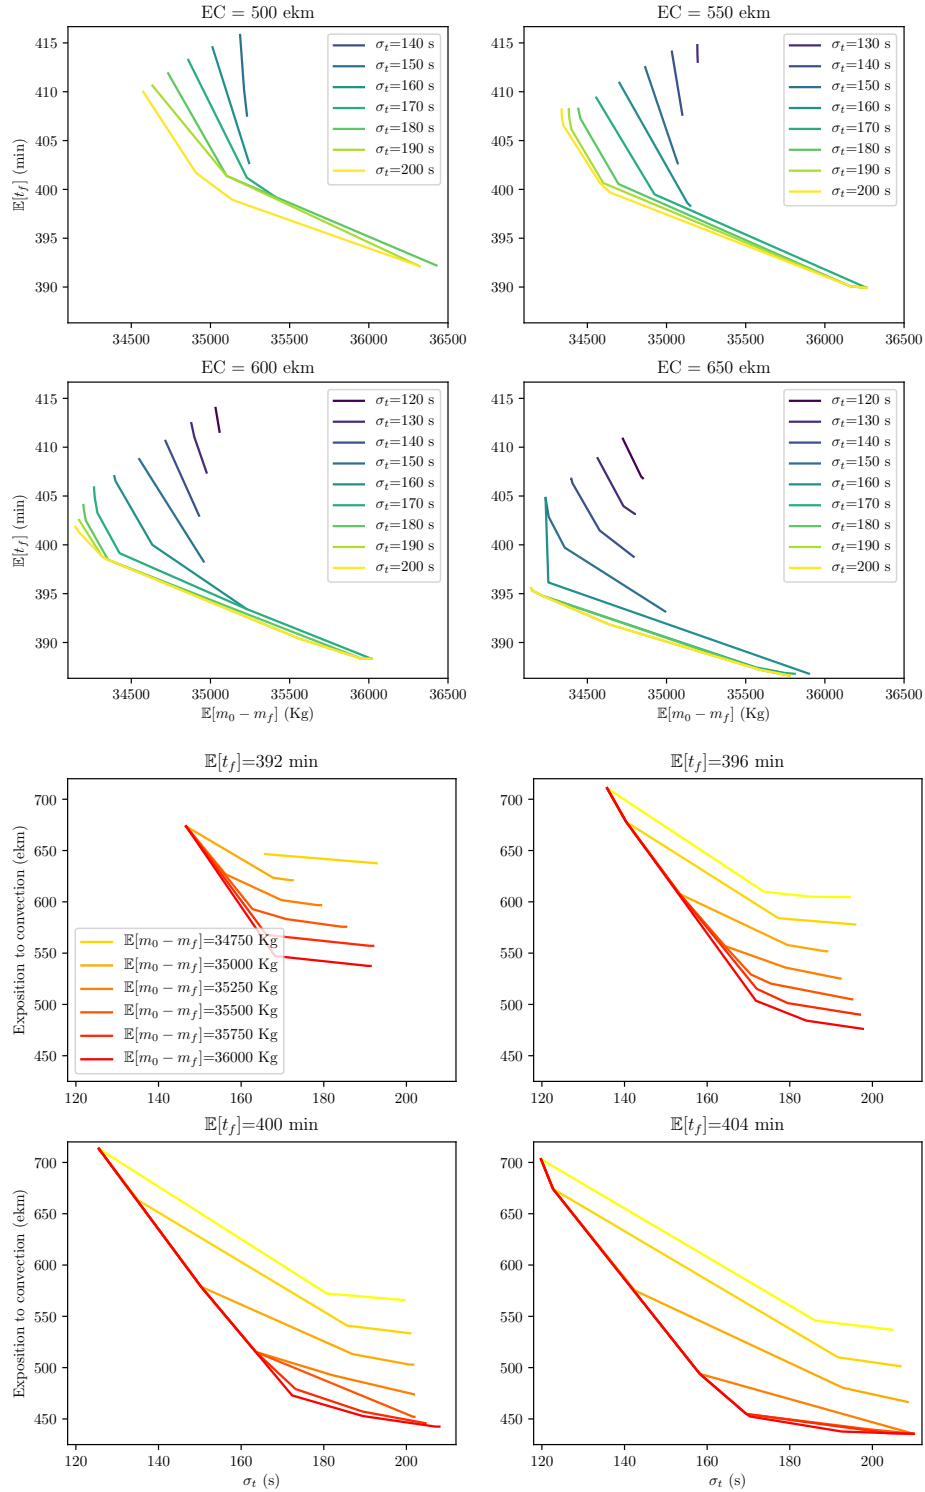


Figure 6.12: Trade-offs (Pareto fronts) between variables

Modeling Uncertainty in Convective Nowcasts

In Chapter 6, we introduced a framework for the consideration of convection at a planning horizon based on numerical weather forecasts. As discussed in 1.4.2, radar and satellite nowcasts are more accurate than NWP forecasts at a tactical horizon. However, these nowcasts often rely on linear extrapolation and rarely consider uncertainty in thunderstorm evolution. In this chapter, we present a model of the uncertainty in the deterministic RDT nowcasts, which we can use to generate probabilistic forecasts in short time horizons; the resulting framework will be employed in Chapter 8 for the purpose of automated trajectory planning in convective environments.

7.1 Deterministic Input

We will now proceed to build a probabilistic forecast based on the deterministic RDT forecasts. The end goal is to obtain a function of the form $p_{t_r}(\mathbf{r}, t)$ that represents the probability that the point $\mathbf{r} = (\phi, \lambda) \in \mathbb{R}^2$ (where ϕ and λ denote geodetic latitude and longitude, in degrees) lies inside a convective storm at time t , as estimated using the forecasts available up to time t_r .

Let t_a be the time at which we build the probabilistic forecast (not the deterministic RDT forecast). We define the reference time $t_{ref} \leq t_a$ as the time at which the latest RDT analysis was released at time t_a ¹. See Figure 7.1 for reference.

As input data, we will make use of the latest N_{steps} forecasts that are available before t_a . We will denote the corresponding forecasts times as follows (see Figure 7.1):

$$t_{-i}^f = t_{ref} - i \cdot 15 \text{ min}, i \in \mathcal{I} := \{0, 1, \dots, N_{steps} - 1\} \quad (7.1)$$




Figure 7.1: Construction of the time-lagged ensemble

For example, if $N_{steps} = 3$ (as in our implementation), we will employ the forecasts made at t_{ref} , $t_{ref} - 15$ min and $t_{ref} - 30$ min.

¹While each analysis contains forecasts for multiple discrete time horizons (15, 30, 45, and 60 minutes), we can extrapolate them to any instant between them by making use of the speed and direction information.

For each of these forecasts, the projected position of each storm can be extrapolated to any posterior time t in a linear and deterministic manner by employing the velocity information provided by the forecast and the temporal distance; in other words, we add the quantity $\mathbf{v}_s \cdot |t - t_{-i}^f|$ to each point in the polygon that defines the contour of the storm, where \mathbf{v}_s is the estimated velocity of the storm in latitude-longitude coordinates and can be expressed as:

$$\mathbf{v}_s = \frac{180}{\pi} R_E^{-1} v_f \begin{bmatrix} \cos \chi_f \\ (\cos \phi)^{-1} \sin \chi_f \end{bmatrix} \quad (7.2)$$

where R_E is the mean radius of the Earth, v_f is the velocity of the storm as estimated by the RDT forecast and χ_f is the estimated direction of movement of the storm (again provided by the RDT extrapolation).

Let $\mathcal{J} = \{\text{Triggering, Triggering from split, Growing, Mature, Decaying}\}$ be the set of possible phases of the storm. For a given reference time, we now compute the indicator functions $I_{-i,j}(\mathbf{r}, t)$ with $i \in \mathcal{I}, j \in \mathcal{J}$. These functions take the value 1 if, according to the forecast made at t_{-i}^f and extrapolated to time t in the aforementioned fashion, the point \mathbf{r} lies inside a storm with phase j ; they take the value 0 otherwise.




Figure 7.2: Example realization of Equation 7.3

7.2 Stochastic Motion

We will now proceed to incorporate stochasticity into the movement of the convective cell in order to represent the forecast uncertainty². We model this evolution as a Brownian motion with drift. This process is described by the following SDE (see Section 2.5.3):

$$d\mathbf{c}_t = \mathbf{v}_s dt + \Sigma_w d\mathbf{W} \quad (7.3)$$

In this equation, \mathbf{c}_t denotes the position of a reference point in the storm, such as the center; \mathbf{v}_s denotes the average speed of the storm. Σ_w is a matrix describing the magnitude and shape of the random perturbations; we model them as isotropic, i.e., $\Sigma_w = \sigma_w \mathbb{I}_{2 \times 2}$ with σ_w denoting the amplitude of the perturbations. Finally, \mathbf{W} denotes a standard two-dimensional Wiener process.

We define the stochastic displacement \mathbf{D} as the difference between the stochastic position \mathbf{c}_t and the deterministic forecast $\mathbf{c}_{t_{-i}^f} + \mathbf{v}_s(t - t_{-i}^f)$ (see Figure 7.2). For a fixed t , this displacement is a Gaussian random variable with mean 0 and standard deviation $\sigma_w \sqrt{t - t_{-i}^f}$; its probability density function is, therefore:

$$p_{\mathbf{D}}^t(\mathbf{y}) = \frac{1}{\sqrt{2\pi}\sigma_w|t - t_{-i}^f|} \exp\left(-\frac{\|\mathbf{y}\|^2}{2\sigma_w^2|t - t_{-i}^f|}\right) \quad (7.4)$$

Because the position of the storm is now a stochastic process, the indicator functions $I_{-i,j}(\mathbf{r}, t)$ become stochastic processes too. We can, however, compute their expected value $\bar{I}_{-i,j}(\mathbf{r}, t)$ with a straightforward convolution operation, as shown by equation (7.5). See Figure 7.3 for a schematic illustration of the convolution process.

$$\bar{I}_{-i,j}(\mathbf{r}, t) = \mathbb{E}[I_{-i,j}(\mathbf{r}, t)] = \int_{\mathbb{R}^2} p_{\mathbf{D}}^t(\mathbf{y}) I_{-i,j}(\mathbf{r} - \mathbf{y}, t) d\mathbf{y} = p_{\mathbf{D}}^t * I_{-i,j}(\cdot, t) \quad (7.5)$$

Performing this approximation for multiple storms in the same manner requires the assumption that the stochastic displacement is common to all storms with the same phase j ; in other words, the movement of the storms is assumed to be perfectly correlated and linearity allows us to take the convolution after aggregation. While it would be desirable to model the correlation of the movements of the storms, we consider this approximation good enough for our purposes, as it approximates a “correct” $\bar{I}_{-i,j}$ by an upper bound.³

²In this work, we will not consider changes in size or shape

³Assume that there are two storms, and the events that the point \mathbf{r} is in each one of them at time t are denoted by A and B . Then

$$P(A \cup B) = P(A) + P(B) - P(A \cap B) = \bar{I}_{-i,j}(\mathbf{r}, t) - P(A \cap B).$$

Therefore, if $P(A \cap B)$ (which represents the probability of the point being inside both extrapolated storms) is small enough, the approximation is close. For more storms, the argument is similar, as the difference is again only intersection terms.




Figure 7.3: Deterministic and stochastic (averaged) indicator functions in one dimension at different, successive times.

7.3 Statistical Model

At this stage, we have $|\mathcal{I}| \times |\mathcal{J}|$ smoothed, extrapolated indicator functions. In order to integrate the information from all of them in a data-driven fashion, we will use them as predictors in a statistical model. We select a logistic regression model, as it is a simple specification that models the output variable as the probability that the event of interest takes place (in this case, the existence of convective conditions in a future step). Therefore, by comparing the predicted indicators with the actual outcomes (according to the analysis in the next time steps), a logistical regression procedure can be employed to predict convective conditions in a probabilistic manner.

In a Generalized Linear Model (GLM) such as linear regression, an outcome variable y (that, in a binary setting like ours, takes values in $\{0, 1\}$) is modeled as a random variable whose expected value depends on the values of a vector \mathbf{z} of predictors or “features”:

$$\ell(\mathbb{E}[y]) = \mathbf{w}^T \mathbf{z} \quad (7.6)$$

where ℓ represents a *link function* while \mathbf{w} is a vector of coefficients whose values are determined, in GLM routines, by maximizing their log-likelihood on a training set. The quantity $\mathbf{w}^T \mathbf{z}$ is called the *linear predictor*.⁴ In the case of logistic regression, the link

⁴Note that, just like in the linear regression case, nonlinear predictors can be included as long as

function is the *logit* function:

$$\text{logit}(p) = \log\left(\frac{p}{1-p}\right) \quad (7.7)$$

In our case, the predictor variables will be the extrapolated indicator functions as well as a constant offset w_0 , so the linear predictor is given by:

$$\text{LP}(\mathbf{r}, t) = w_0 + \sum_{\substack{i \in \mathcal{I} \\ j \in \mathcal{J}}} w_{i,j} \bar{I}_{-i,j}(\mathbf{r}, t) \quad (7.8)$$

and the desired probability function can be written as:

$$p_{t_{ref}}(\mathbf{r}, t) = \text{logit}^{-1}(\text{LP}(\mathbf{r}, t)) = \frac{1}{1 + \exp(-\text{LP}(\mathbf{r}, t))} \quad (7.9)$$

7.4 Model Fitting

The complete model is defined by the parameters N_{steps} and the coefficients σ_w, w_0 and $w_{i,j}$. We choose $N_{steps} = 3$ after our exploratory analysis concluded that the coefficients after $i = 2$ were too close to 0 to be relevant and the statistical skill of the model was not appreciably deteriorated by the exclusion of the $i > 2$ forecasts; in other words, the information contained in the forecasts before the third to last seems to be redundant for prediction purposes. In order to determine the remaining parameters, we will make use of three procedures:

- An “inner loop” that computes the values of w_0 and $w_{i,j}$ for a given value of σ_w , as well as a measure of the statistical skill of the model.
- An “outer loop” that calls the inner loop for different values of σ_w and then selects the value σ_w^* that leads to the best statistical skill.
- A “bootstrap-like” method that runs the inner loop with $\sigma_w = \sigma_w^*$ and different, randomly selected slices of the dataset in order to obtain more robust estimates of the w_0 and $w_{i,j}$ coefficients as well as to estimate their volatility.

We start by describing the inner loop procedure. For a given value of σ_w , we can compute the values of the coefficients w_0 and $w_{i,j}$ with the following method:

1. Select a number of dates and times t_a in random fashion from a training dataset and load the corresponding RDT analyses, as well as the N_{steps} preceding forecasts.

they are included in additive fashion.

2. For each analysis, compute the “areas of interest” (AoI), defined as rectangular regions containing any detected or forecasted storms while adding proportional margins (see Figure 7.5). This step reduces the subsequent computational requirements of the process.
3. At each point \mathbf{r} in a rectangular grid inside an AoI, compute the values of $\bar{I}_{-i,j}(\mathbf{r}, t_a)$ and the actual value $I_a(\mathbf{r}, t_a)$ (i.e., 1 if the point lies inside a storm according to the analysis and 0 otherwise). Note that the extrapolated and smoothed indicator $\bar{I}_{-i,j}(\mathbf{r}, t_a)$ depends on the value of σ_w .
4. Launch a logistic regression routine employing the computed values of $\bar{I}_{-i,j}$ as predictors and the values of the points as data samples, returning the values of the w_0 and $w_{i,j}$ coefficients.

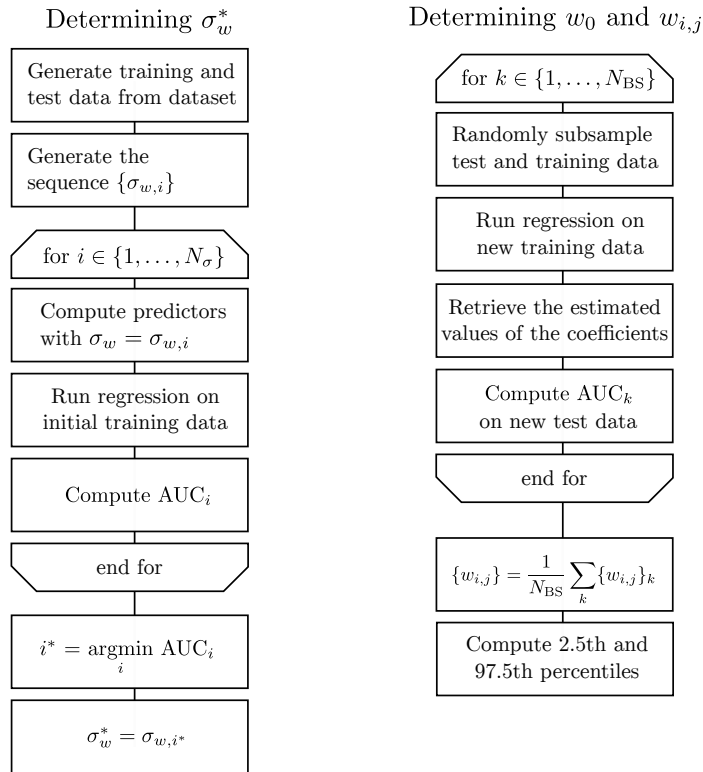


Figure 7.4: Schematic representation of the model fitting procedures. N_{BS} denotes the number of bootstrap analyses.

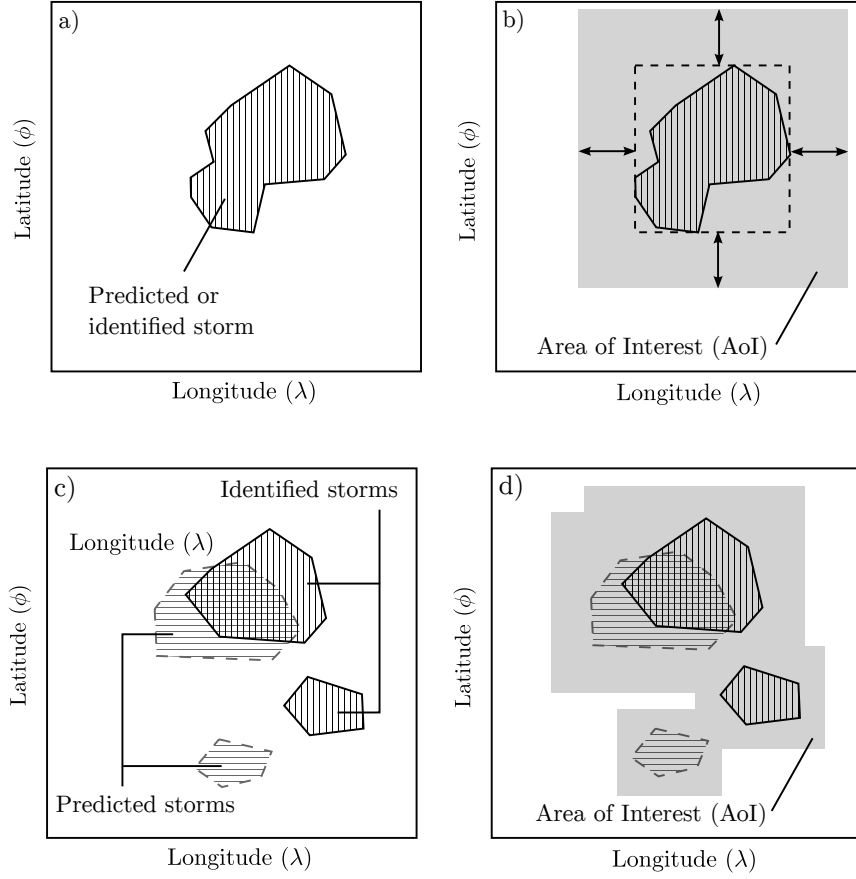


Figure 7.5: Area of Interest determination. For a given storm (as in Figure 7.5.a), the corresponding AoI is the rectangle containing it, with proportional margins (Figure 7.5.b). For each analysis, comprised of multiple identified storms plus forecasted storms from earlier analysis (Figure 7.5.c), the AoIs from each storm are merged.

The performance of the regression model can be evaluated with several metrics. Suppose that we consider the test set (composed by data that have not been employed in training) and predict that every data point where the forecasted probability is above a threshold to have convective conditions and every other data point to have clear air conditions. If we plot the true positive rate against the false positive rate as the threshold changes from 0 to 1, we obtain the Receiver Operating Characteristic (ROC) curve. The integral of the area below the ROC curve is called the Area Under Curve (AUC) metric; AUC values of around 0.5 indicate that the model is no more skilled than random classification on average, while values closer to 1 indicate high skill.⁵

⁵Values closer to 0 indicate “reverse skill”, i.e., the model is worse than chance at classifying outcomes




Figure 7.6: Values of the AUC (Area Under Curve) metric. The dashed line represents the polynomial approximation employed to determine the value of σ_w that optimizes the AUC.

We can now employ this inner loop process in order to choose the value of σ_w in the outer loop and the values of w_0 and $w_{i,j}$ in the bootstrap-like procedure. Both methods are illustrated in Figure 7.4.

For the outer loop, we can employ the AUC criterion to choose the “best” value of σ_w . We compute the logistic regression on a training set for different values of the σ_w parameter (using the same randomly chosen data). Then, we compute the AUC score on a test set. Finally, we select the value of σ_w that lies on the maximum of the σ_w - AUC curve. This process is performed twice: first, with a logarithmic sweep of σ_w in order to identify the scale of interest (upper plot in Figure 7.6); then, with a linear sweep and polynomial approximation in order to identify the optimum with more precision (though we will note that somewhat lower or higher values lead to very similar AUC scores). With our dataset, comprising RDT analysis and forecasts from the 20th to the 29th of March, 2017, the chosen value is $\sigma_w = 0.73 \text{ km} \cdot \text{s}^{-0.5}$.

Finally, with σ_w^* fixed, we can now estimate the mean values and confidence intervals of the coefficients $\{w_{i,j}\}$ with the bootstrap-like procedure. We randomly select 250 dates and times from the available data set and subsample 25% of the resulting areas of interest at random; these settings allow us to hold the resulting sample in-memory. Then, we obtain the coefficient values by regression and evaluate the performance by

so the reverse prediction has skill.

computing the AUC score on a test dataset composed by a random 25% of a random 150 dates and times from the remaining files. We repeat this process 1000 times and take the average, as well as the 2.5th and the 97.5th percentiles which allow us to define the 95% confidence intervals.

The computed values of the coefficients $\{w_{i,j}\}$ are presented in Table 7.1 and Figure 7.7. It can be observed that the step of the forecast is more relevant in the prediction than the analyzed storm phase, with the latest forecast being 2 to 4 times more influential than the previous one, while the earliest forecast is statistically significant only partially. Additionally, we can see that the estimated values of the coefficients for storms that are triggering are more volatile; we attribute this outcome to the fact that they are less numerous within the data set, so the corresponding coefficient estimates have higher variance.

The computed AUC values range from 0.949 to 0.958. This seems like an excellent value, but we will note that our dataset features a large amount of points where no storm is extrapolated and no storm occurs; this inflates the AUC score by providing a lot of “easy predictions”. In practice, our model will often miss newly created storms that are hard to detect using extrapolation procedures, unless they are close enough to other storms that the probability dispersion effect assigns some non-zero probability.

Figure 7.8 compares the deterministic extrapolation of an RDT nowcast to the probability field generated by the fitted model. It can be observed how the probability field diffuses as the forecast horizon increases.

	0 min	-15 min	-30 min
Triggering	5.71 (4.95, 6.79)	2.56 (1.47, 3.54)	0.05 (-0.87, 1.23)
Tr. (split)	6.45 (6.11, 6.81)	1.76 (1.20, 2.34)	-0.05 (-0.68, 0.57)
Growing	5.65 (5.40, 5.89)	1.24 (0.87, 1.61)	-0.08 (-0.43, 0.28)
Mature	5.76 (5.52, 6.03)	1.75 (1.38, 2.10)	0.70 (0.35, 1.03)
Decaying	5.70 (5.49, 5.92)	1.55 (1.23, 1.87)	0.58 (0.25, 0.91)
w_0	-4.15 (-4.19, -4.12)		

Table 7.1: Coefficient values $\{w_{i,j}\}$ and 95% confidence intervals computed as described in the text.

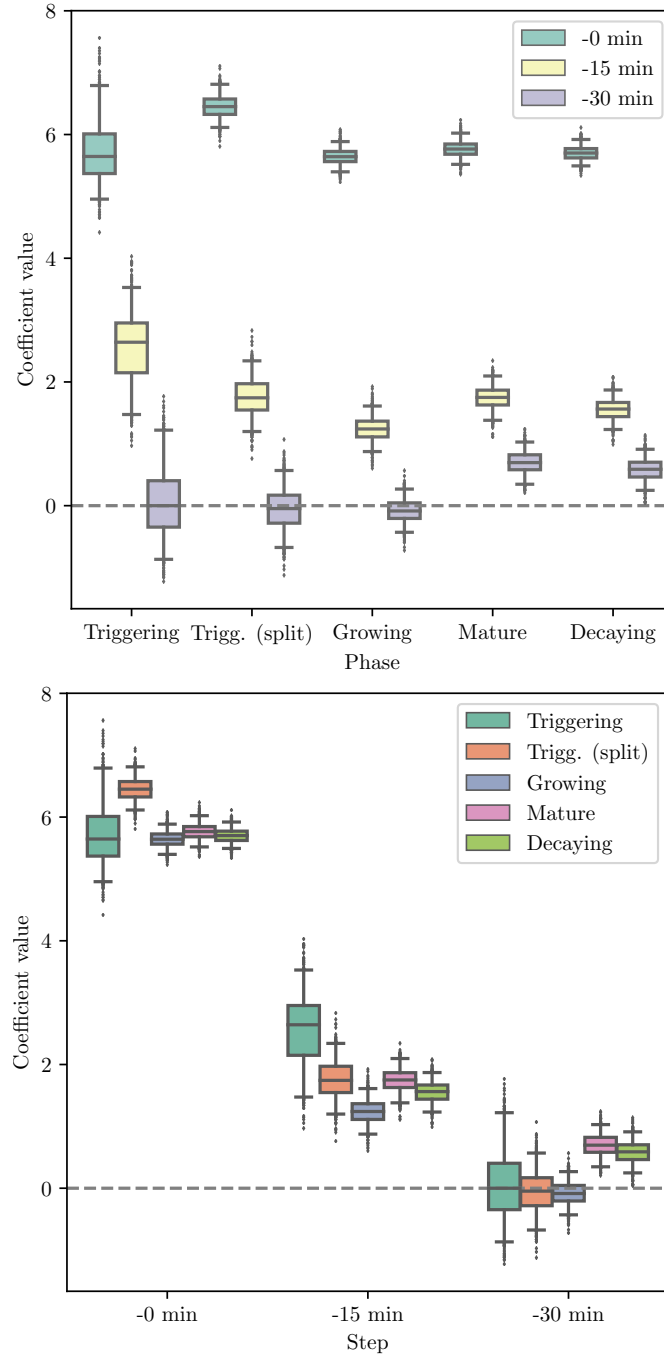


Figure 7.7: Coefficient estimates w_0 and $\{w_{i,j}\}$. Each boxplot represents the distribution among different bootstrap-like fits, with the box representing the 25th, 50th and 75th percentiles and the whiskers representing the 10th and 90th percentiles.

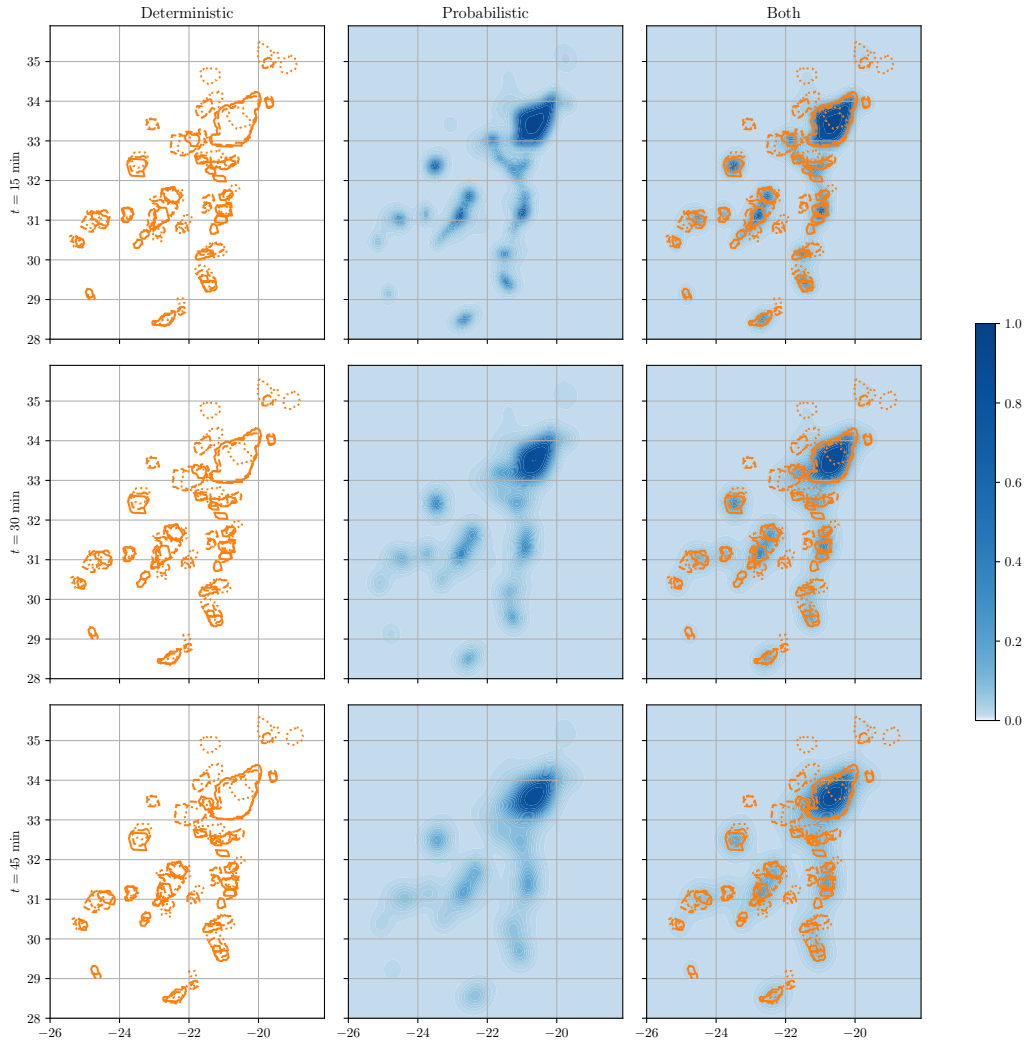


Figure 7.8: Comparison between the deterministic RDT extrapolation and the probabilistic model at different forecast horizons. The solid, dashed, and dotted lines represent the extrapolation from the forecast at t_{-0}^f , t_{-1}^f , and t_{-2}^f , respectively. Color bar denotes $p_t(\mathbf{r}, t)$.

We will now turn our attention to the issue of how to employ these generated probabilistic forecasts for the purpose of flight planning under uncertainty in Chapter 8.

Optimal Stochastic Storm Avoidance

In the preceeding chapters, we have studied the problem of robust flight planning at a “planning horizon” of 2 to 3 hours before departure, optimizing an entire cruise operation. In the current chapter, we turn our attention to a narrower “tactical horizon” of 5 to 30 minutes before a potential encounter with meteorological hazards, optimizing only a leg or segment of the flight. At this stage, convection represents the most salient weather phenomenon in terms of risks to the aircraft, and thus pilots are led to avoid convective areas with maneuvers and trajectory modifications that are hard to predict in advance, introducing uncertainty into ATM operations.

The chapter describes a trajectory planning framework under convective activity, based on the probabilistic thunderstorm evolution model introduced in Section 3.3.2. We describe the modeling strategy and the problem formulation in Section 8.1. The resulting optimal control problem is initialized with a randomized heuristic, described in Section 8.2, with the goal of exploring the multiple local optima of the problem. Finally, Section 8.3 contains a case study that showcases the proposed approach in a convective scenario, as well as a discussion of the results.

8.1 Problem Modeling

8.1.1 Dynamical Model

We assume that the uncertainty in the movement of the storm is substantially greater than the uncertainty in aircraft dynamics. Therefore, we model the motion of the aircraft as a deterministic system, in contrast to Chapter 5. We assume that the aircraft moves in a constant flight level and that turning dynamics are important; thus, we choose to employ the **2Dt-m μ -e** model from Section 3.1.4. The corresponding differential-algebraic equations are:

2Dt-m μ -e model

State variables: $\phi, \lambda, v, \chi, m$

Control variables: C_L, C_T, μ

Differential equations:

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

Algebraic constraints:

$$L(C_L) \cos \mu = mg$$

In addition, we add the following constraint:

$$v \leq v_{max} \tag{8.1}$$

where v_{max} represents the potential speed limit when flying through turbulent air near convective zones.

8.1.2 Optimal Control Model

We turn again to optimal control to solve the flight planning problem at hand. The objective function is defined as follows:

$$J = \left[-m(t_f) + \text{CI} \cdot t_f + r_p \int_{t_0}^{t_f} p(\mathbf{r}(t), t) dt \right], \quad (8.2)$$

where CI is the “Cost Index” parameter, representing the preference for reduced flight time instead of fuel burn, and r_p is the “risk penalty”, which we define as the preference for reduced exposition to regions of high storm probability instead of the other objectives. A higher r_p setting generates a trajectory that is less likely to be exposed to convective activity.

The optimal control problem is formulated as:

$$\min J \quad (8.3)$$

subject to:

$$\frac{d}{dt} \begin{bmatrix} \phi \\ \lambda \\ v \\ \chi \\ m \end{bmatrix} = \begin{bmatrix} (R_N + h)^{-1} (v \cos(\chi) + w_x(\phi, \lambda, t)) \\ (R_M + h)^{-1} \cos^{-1}(\phi) (v \sin(\chi) + w_y(\phi, \lambda, t)) \\ (\text{T}(C_T) - \text{D}(C_L, v))/m \\ (g/v) \tan \mu \\ -\eta(v) \text{T}(C_T, v) \end{bmatrix} \quad (8.4)$$

$$\text{L}(C_L, v) \cos \mu = mg \quad (8.5)$$

$$|\mu| \leq 45^\circ, \quad (8.6)$$

$$C_L \leq C_{Lmax} \quad (8.7)$$

$$M \leq M_{MO} \quad (8.8)$$

$$C_{Tmin} \leq C_T \leq C_{Tmax} \quad (8.9)$$

and the boundary conditions:

$$(\phi, \lambda, v, \chi, m)(t_0) = (\phi_0, \lambda_0, v_0, \chi_0, m_0) \quad (8.10)$$

and

$$(\phi, \lambda, v)(t_f) = (\phi_f, \lambda_f, v_f) \quad (8.11)$$

If a required time of arrival (RTA) to the end waypoint is specified, we set $t_f = t_{RTA}$ instead of leaving it free. Naturally, if an RTA is set in this fashion, the CI setting becomes irrelevant.

As in previous chapters, we work with direct transcription methods that transform the optimal control problem into a nonlinear optimization problem. Therefore, a user-provided initial guess for the decision variables (the values of the trajectory state and control variables at every node) is required. We describe the generation of this initial guess in Section 8.2.

Note that the presented optimal control scheme requires a probability field of function of the form $p(\mathbf{r}(t), t) \in [0, 1]$ (in either analytical or gridded representation), such as the one we generate in Chapter 7. Therefore, a different weather forecast or product could be employed as input, as long as it is provided in the required shape and represents the same concept (the probability of an aircraft at \mathbf{r} being within a convective cell at time t) or a similar one.

8.2 Initialization

The selection of the initial guess is important for two reasons. First, the quality of the initial guess is, in general, an important factor in the performance of NLP solvers. The number of iterations, the computational time (particularly relevant for our purposes) and the likelihood that the algorithm converges to an optimal solution are all dependant on the initial guess.

Second, we expect the optimization problem that we formulate to have multiple local minima; therefore, the initial guess will determine which of these multiple solutions will be found by the algorithm. We expect different local minima because of the non-convex nature of the problem, which is derived to a high extent from the existence of multiple potential avoidance routes that are separated (in trajectory space) by “worse” routes that cross the obstacles. This implies the existence of multiple basins of attraction around the locally optimal routes.

For these reasons, we have designed a randomized heuristic initialization procedure that produces different initial guess trajectories with the goal of exploring the solution space and find the different local minima. Generating multiple solutions also has operational benefits, as pilots and controllers can then choose one of them according to a different criterion (for example, facilitating deconfliction or sequencing); additionally, it is possible for them to employ their trained intuition and experience to judge the feasibility and complexity of the trajectories, or any factor not included in the model, to take the best decision in practice.

The generated trajectories will be generated assuming no wind and constant airspeed v_{IG} . Under these assumptions, any trajectory from \mathbf{r}_0 to \mathbf{r}_f can be represented by the

arrival time t_f and the heading history $\chi(t)$. However, it is not true in general that any randomly chosen arrival time and heading history will lead to a satisfactory solution, as there is no guarantee that the arrival point will be \mathbf{r}_f . Therefore, we will generate a “normalized” heading history $\hat{\chi} : [0, 1] \rightarrow \mathbb{R}$ such that

$$\chi(t) = \chi_0 + \hat{\chi} \left(\frac{t}{t_f - t_0} \right), \quad (8.12)$$

where χ_0 and t_f are adjustable constants that allow us to generate a valid heading history from any continuous normalized heading history $\hat{\chi}$ through a straightforward rotation (adjustment of χ_0) and stretch (adjustment of t_f).

In order to generate $\hat{\chi}$, we consider the following basis expansion:

$$\hat{\chi}(\tau) = \sum_{k=1}^{n_L} a_k Z_k P_k(\tau) \quad (8.13)$$

where n_L represents the degree of the expansion (we use 4), $\{Z_k\}$ is a sequence of independent standard normal variables, $\{a_k\}$ represents a weighting sequence (in our implementation, $a_k = (3/4)^{1-k}$) and $\{P_k\}$ denotes the shifted Legendre polynomials (the Legendre polynomials on $[0, 1]$). Note that the zeroth term is omitted, as it corresponds to a constant offset and χ_0 already fulfills that role. Table 8.1 lists the shifted Legendre polynomials up to degree 5.

k	$P_k(\tau)$
0	1
1	$2\tau - 1$
2	$6\tau^2 - 6\tau + 1$
3	$20\tau^3 - 30\tau^2 + 12\tau - 1$
4	$70\tau^4 - 140\tau^3 + 90\tau^2 - 20\tau + 1$
5	$252\tau^5 - 630\tau^4 + 560\tau^3 - 210\tau^2 + 30\tau - 1$

Table 8.1: Legendre polynomials on $[0, 1]$.

By sampling $\{Z_k\}$ in random fashion, we can generate different functions $\hat{\chi}(t)$, which determine different heading histories $\chi(t)$. We generate the rest of the state and control trajectory by integrating Equation (8.4), thus completing the initial guess. Finally, by solving the problem with different starting points obtained in this fashion and collecting the corresponding solutions, we have a higher chance of obtaining a global optimum instead of a single local optimum.

We note that the relationship between a given starting point and the optimal solution found by the solver is not fixed, but may depend on the choice of NLP solver and its settings.

8.3 Results

Our test scenario is based on a storm group that is detected by the RDT algorithm on November 16th, 2017. The storm group is evolving towards North-East. The scenario starts at 6:00Z and ends at 07:30Z. A twinjet narrow-body airliner modeled according to the BADA 4 specification flies at FL330 from $(\phi_0, \lambda_0) = (34^\circ, -24^\circ)$ to $(\phi_f, \lambda_f) = (28^\circ, -19^\circ)$, therefore crossing the group of storms. The remaining initial conditions are given by:

$$\chi_0 = 115^\circ \quad (8.14)$$

$$m_0 = 71495 \text{ kg} \quad (8.15)$$

$$v_0 = 220 \text{ m/s} \quad (8.16)$$

We select $r_p = 10 \text{ kg/s}$, a required time of arrival at the end waypoint of 07:12Z, a v_{max} of 230 m/s, and a final airspeed of 200 m/s. We solve the problem with the initialization procedure described in Section 8.2, starting from 3400 randomly chosen starting points. We will discuss the best solution found in Section 8.3.2 and the other local minima afterwards, in Section 8.3.3.

8.3.1 Computational Setup

We employ the Python-based CasADi [251] library for NLP modeling. The probability field is computed in a $0.1^\circ \times 0.1^\circ$ grid, sampled every 5 minutes, and then interpolated through 3-dimensional B-splines. We employ the interior-point NLP solver IPOPT [77] running with the MA27 sparse symmetric linear solver from the HSL Mathematical Software Library [253] and initial barrier parameter $\mu = 10^{-3.8}$. The computations are performed in a workstation equipped with an Intel Xeon E3-1240 v5 CPU running at 3.5 GHz.

We employ a trapezoidal transcription scheme [75, Chapter 4] with piecewise-constant controls as direct collocation method, with defect constraints of the form

$$\mathbf{x}_{n+1} - \mathbf{x}_n = \frac{\Delta t}{2} (\mathbf{f}(\mathbf{x}_{n+1}, \mathbf{u}_n, t_n) + \mathbf{f}(\mathbf{x}_n, \mathbf{u}_n, t_{n+1})), \quad \forall n \in \{0, \dots, N_{\text{nodes}} - 1\},$$

with the subscript n denoting the value of the directized variables at the n -th node or interval and $\Delta t = \frac{t_f - t_0}{N_{\text{nodes}} - 1} \equiv t_{n+1} - t_n$ denoting the node spacing. By employing an homogeneous node spanning h , this transcription scheme samples the probability field in a homogeneous fashion, unlike higher-order pseudospectral methods [81]; furthermore, the potential existence of constrained arcs (such as legs at maximum speed) does not allow us to assume that pseudospectral methods will provide spectral accuracy unless coupled with complex adaptive methods. In any case, the choice of discretization method is not critical for our work.

8.3.2 Global Optimum

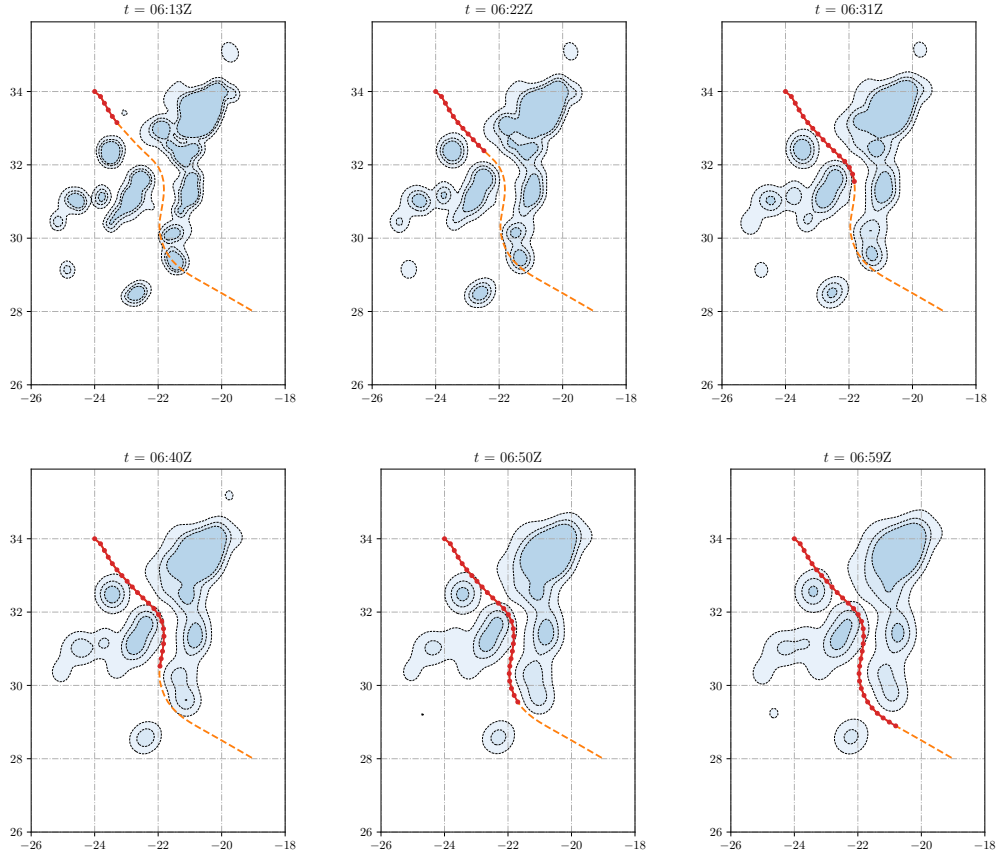


Figure 8.1: Computed trajectory and storm probability $p(\mathbf{r}, t)$ at different points in the trajectory. The solid line with dots represents the flown part of the trajectory while the dashed one represents the remaining part. The three contour levels represent the 2.5%, 5% and 10% probability levels.

Figure 8.1 illustrates the evolving probability field $p(\mathbf{r}, t)$ at different instants in time, as well as the planned trajectory of the aircraft, passing through a corridor between the storms. It can be observed that the planned trajectory takes the future potential evolution of the storms into account, and not just the present location of the storms.

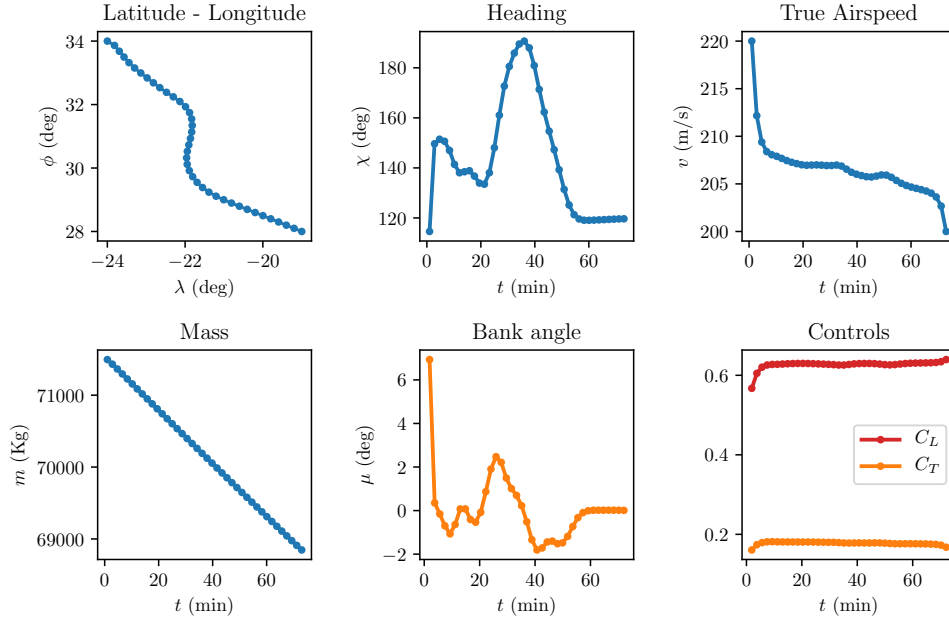


Figure 8.2: State space and control trajectories for the global solution.

State and control trajectories are illustrated in Figure 8.2. Since the resulting airspeed profile is relatively smooth, the longitudinal controls (C_T and C_L) are stable, while the lateral-directional control (the bank angle, μ , with C_L playing a compensation role) changes in order to steer the aircraft. Beyond the initial and final transitions towards the initial and final conditions, the speed slowly decreases as the mass decreases. This could be attributed to multiple factors: it might be optimal to go faster at the beginning in order to avoid the storms incoming from starboard and to go slower at the final leg in order to wait for the storms at the port side to go away. Alternatively, the TAS profile might be similar to the one in a simple, one-dimensional cruise problem [94]; in this kind of setting, the optimal solution requires that the airspeed slowly goes down as the mass of the aircraft decreases.

8.3.3 Local Optima

We illustrate all the locally optimal paths found by the method in Figure 8.3. We highlight the five solutions that are present in more than 1% of the runs and have objective values closer than 100 kg to the global optimum.




Figure 8.3: Trajectories from different randomized starting points. For the most likely trajectories with close-to-optimal objective function values, we annotate the difference in cost with respect to the global optimum and the percentage of the time that the randomized initial guess produces the trajectory. The storm probability field is represented at 06:30Z

It can be observed that about two thirds of the starting points lead to the most common trajectory, which is very close in cost (12 kg) to the global optimum. Another 4.4% of the starting points lead to the global optimum, as discussed in Section 8.3.2. Together, the five highlighted trajectories account for nearly 85% of the initialization points; the remaining 15% of the starting points lead to trajectories that have higher costs or are found infrequently. Some of these solutions try to route around the whole storm group, but the RTA requirement prevents them from fully avoiding the bigger storm clusters.

These results confirm our hypothesis about the multiplicity of local optima formulated in Section 8.2. Therefore, the usage of a randomized initial guess procedure is justified.

Figure 8.4 displays the airspeed profiles of these trajectories. The most common solution features the slowest airspeed profile and, thus, the lowest expected fuel burn (as long as the trajectory does not have to reroute). It does so, however, at the expense of increasing the potential exposition to convective hazards when compared to the global optimum (see Figure 8.5 and Figure 8.6).




Figure 8.4: TAS profiles of the randomly initialized trajectories

Another feature present in the airspeed profiles shown in Figure 8.4 is the appearance of peaks of increased airspeed, which correspond to the crossing of zones of higher storm probability (compare with Figure 8.3 and Figure 8.5). This phenomenon can be explained by the shape of our cost functional, where the cost of flying through risky zones depends not only on the probability level but on the time spent in risky zones, and thus flying faster is recommended. This might seem to contradict recommendations to not fly at high speed through turbulent air, but we address this issue with the v_{max} setting in Equation (8.1). The best solution, which does not cross zones where the probability field is above 2.5%, is the only one that does not feature these temporary airspeed increases.




Figure 8.5: Probability field along the trajectory. The horizontal dashed lines represent the 2.5%, 5% and 10% probability levels.

Figure 8.6 shows the two components of the objective value that is achieved by each trajectory. As discussed earlier, approximately 85% of the starting points lead to one of the five highlighted trajectories, which have close-to-optimal cost. A few of the remaining trajectories also have relatively small cost increases with respect to the optimum (50 - 200 kg), manifesting mainly in the form of increased fuel burn. The second cluster has

similar fuel consumption patterns, but the potential exposition to convective weather rises by about 50% (from 1.4 equivalent minutes to 2.1 equivalent minutes).




Figure 8.6: Solution cost. The dashed lines represent level curves of equal total cost.

The information provided in these figures could be presented to the decision-maker in order to assess the safety of the proposed trajectory, as Figure 8.5 represents the risk of encountering thunderstorms at each instant in the trajectory while Figure 8.6 illustrates the accumulated risk. Note that, in our approach, the event to avoid (because it would likely lead to an unplanned route deviation) is an encounter with a convective cell, not an instance of loss of separation with the convective cell, as we deal with uncertainty in an explicit form instead of employing pre-determined safety margins. Nevertheless, it is certainly possible to add additional safety margins to our approach by enlarging the deterministically nowcasted storms by the desired amount before the application of the model.

The distribution of the computational cost associated with the solution of the optimization problem (not including preprocessing) is shown in Figure 8.7. It can be observed that the NLP solver takes less than 30 seconds of CPU time to find the corresponding solution around 90% of the time. The computational cost of the algorithm

can be thus considered as moderate. While the 10-40 seconds of optimization time is useful for analysis and close to real-time performance, the multiple runs needed to obtain the alternative trajectories with the initialization scheme represent a limit in terms of practical implementation. Nevertheless, because these runs are independent processes, it is possible to run them in parallel.




Figure 8.7: Distribution of the computational cost.

Conclusions and future work

In this dissertation, we have introduced methods based on optimal control to design efficient aircraft trajectories in the presence of meteorological uncertainty. In first place, an appropriate modeling framework for aircraft motion in ATM contexts was proposed. Meteorological uncertainty was modeled at multiple timescales, using ensemble forecasts to model wind and convection at the planning horizon and satellite nowcasts to model convective activity at the tactical horizon.

In order to select a computational framework, an examination of the literature on dynamic optimization and flight planning was performed. Because of its ability to model complex aircraft dynamics and handle dynamical systems of moderate size, numerical optimal control techniques based on direct methods are adequate for the solution of free-routing aircraft trajectory planning problems.

The first contribution of the thesis represents a framework for the solution of nonlinear optimal control problems under parametric uncertainty. It extends previous concepts from the literature, based on restrictive “open-loop” operational paradigms, in order to solve a much wider class of practical optimization problems under the “tracking” concept, including aircraft trajectory planning problems. In addition, this approach has been shown to handle policy-based problems, although this idea has not been explored in-depth within this work. Since the size of the resulting NLP problem is linear on the number of scenarios, the computational load is manageable.

Building on these pillars, we have designed a framework for flight planning under uncertainty at the planning horizon. With this method, it is possible to design trajectories with optimal average efficiency (depending on the relative cost of time and fuel); in addition, the capability to generate flight plans that are more predictable at the cost of some efficiency is also enabled by the proposed method. These predictability improvements are found by the algorithm by shifting routes to zones where the winds are predicted with more certainty, which is expensive in terms of efficiency, and by choosing higher airspeeds when overflying these regions, which is cheaper. Finally, a dynamic airspeed management paradigm where the aircraft increases or decreases its true airspeed in proportion to the accumulated delay is studied within this framework. It is shown that the predictability can be improved at almost no cost in terms of average efficiency;

however, the price is a significant increase on the uncertainty in fuel burn.

A transatlantic flight scenario was studied to illustrate these findings. While these numbers are likely not representative of general costs or benefits, they might give an indication of the shape and order of magnitude of the benefits. If the airspeed is fixed, the arrival window for the maximum average efficiency flight plan of this flight has a size of 6 minutes; a 1-minute reduction of the size of this window can be achieved with a 30-second increment in average flight duration. However, additional reductions in uncertainty become substantially more expensive: for example, reducing this range by another minute demands a 3-minute increase in average flight time. A similar pattern of diminishing returns is present in the variable airspeed case, though the relationship becomes more complex due to the interplay between fuel burn and flight duration, mediated through the airspeed schedule. Finally, we found that the proposed dynamic airspeed management can improve predictability arbitrarily, but even a halving of the time uncertainty can result on a two-fold or even three-fold increase in the fuel burn uncertainty at a high cost index.

While convective activity cannot be forecasted with complete precision at planning horizons, the presence of preconditions for the development of thunderstorms can be computed in a probabilistic manner. We have integrated this indicator into the proposed trajectory planning framework in order to consider the predictability impact of flying through areas where convection may develop. The test scenario illustrates the emergent tradeoffs, showing that exposition to potential convective phenomena can be reduced through flight planning; however, this again comes at a cost of either average efficiency or additional exposition to wind-related uncertainty.

At a tactical horizon, we have presented a framework for the generation of trajectories under uncertain convective environments. We build a model of the stochastic evolution of convective cells based on deterministic nowcasts derived from satellite data. Then, we employ the information from this model in an optimal control framework that is initialized through a randomized heuristic procedure, with the goal of exploring the local minima of the problem. By employing this method, a set of efficient and safe (and, thus, predictable) trajectory options can be generated automatically, allowing pilots and controllers to choose the one that best fits into the traffic situation. This development could be employed in the future to develop applications in both onboard and ground-based Decision Support Tool (DST). At the airborne side, pilots could employ such a DST to quickly obtain and evaluate flight path alternatives; in the ground, it could help air traffic controllers to assign time-based metering slots to aircraft flying through weather impacted regions and/or group aircraft trajectories into weather-avoiding flows that are easier to manage. Finally, it could be employed as a basis for trajectory synchronization and negotiation between air traffic control and the aircraft.

We employ this approach to solve a scenario where an aircraft has to arrive to a

given waypoint at a specific time, crossing a group of convective cells. The method successfully identifies potential deviation trajectories, employing different corridors between the thunderstorms: they are classified in terms of trajectory cost and exposition to thunderstorm encounters. As these encounters lead to trajectory modifications that would reduce downstream predictability, it is in the interest of the whole ATM system to plan trajectories that reduce the chance that these events take place.

To summarize, this thesis develops methods based on optimal controls for trajectory planning under uncertainty. We have studied the trade-offs involved and shown how these methods can generate the relevant information and insights to assess the costs of enhancing predictability through improved flight planning. This improvement in predictability can then be leveraged by the ATM system to increase capacity and reduce delays and emissions. From the results presented, we can therefore conclude that the proposed optimal control techniques could serve as a step towards systems and tools for automated trajectory planning under uncertainty, leading to more efficient and predictable ATM operations.

9.1 Future Work

The work carried out within this thesis opens up numerous avenues of research. In this section, we will proceed to explore the most prominent ones:

9.1.1 Validation

Validation exercises are also necessary in order to demonstrate the benefits of our methods in practical terms. On one hand, it must be shown that the solution provided by the method is valid in statistical terms; on the other hand, the potential benefits must be quantified for different routes and dates in a certain airspace region. Further analysis is also necessary in order to improve our understanding of the relationship between trajectory predictability and ATM performance and the costs of improving predictability through flight planning should be quantified in a systematic way. These studies would lay the necessary groundwork for the design of policies that provide the right incentives to flight planners to produce more predictable flight plans and, therefore, increase the efficiency of the system as a whole.

9.1.2 Extended Aircraft Dynamics

A clear field of potential future inquiry would be an extension of the aircraft dynamics to a full 4D concept with the addition of a vertical profile (changing flight levels) and, perhaps, multiple aerodynamic configurations. Because we work within an optimal

control framework, incorporating a vertical dimension in a free routing framework¹ does not fundamentally change the core problem and the same concepts can be applied. This is particularly interesting in the context of tactical trajectory planning, as overflying a convective cell at a higher altitude is sometimes another alternative.

A more challenging modification would be the consideration of structured airspaces, as most aircraft operations will take place under this paradigm in the short and medium term. The selection of discrete flight levels could be modeled with an appropriately shaped penalty function and a multi-step initialization procedure; however, the selection of the horizontal route (as a sequence of selected airways and waypoints) represents a greater challenge, likely requiring the employment of MINLP optimization techniques or hybrid optimization methods.

9.1.3 Multiple Uncertainty Sources

Another relevant addition to the planning model would be the incorporation of additional uncertainty sources. Temperature is the next obvious choice, as it influences engine performance and air density (at a given pressure level) and, thus, aerodynamic performance. Furthermore, its inclusion is straightforward as uncertainty in temperature can be modeled using the same ensemble forecasts as for wind.

Other relevant sources of uncertainty can be included, such as aircraft mass, aircraft performance and takeoff time. The inclusion of these factors would be somewhat more complex, as the new stochastic quadrature rule would need to integrate them “orthogonally” with respect to weather uncertainty. Two possibilities can be mentioned: modeling meteorological deviations from the mean as significant “patterns” weighed by continuous random variables, in order to employ a multidimensional SQR, or generating the SQR as a tensor product of the ensemble SQR and another SQR.

The consideration of these additional uncertainties would be useful to model shorter flight legs, where there is not enough time for wind-related uncertainties to accumulate. In particular, climb and descent operations represent a popular research topic in the field of trajectory optimization, and studying them under existing uncertainties would be valuable.

9.1.4 Fast Quasi-Optimal Heuristics

While numerical optimal control with direct methods is an effective framework for the calculation of optimal trajectories, it is not guaranteed to converge to a solution in a given timeframe. For certain applications, requiring real-time or near real-time computational times, it may be necessary to sacrifice some of the advantages of optimal

¹Discrete flight levels present a bigger challenge for optimal control, but some options are still available.

control in favour of speed. One possible approach to this problem would involve fast heuristics or metaheuristics, leveraging high-performance computational platforms (such as the Graphics Processing Unit (GPU)) to simulate multiple scenarios simultaneously. We present some preliminary results on this research line in [8].

9.1.5 Estimation of the Cost of Convection

In order to consider convection at the planning stage, a value for the convective penalty parameter (CP) must be chosen. As shown in Section 6.1, this requires the estimation of a model of the costs of flying through regions where convection may be present. In this way, its impact can be translated into average costs and predictability losses, which allows the determination of CP from the CI and DP parameters.

9.1.6 Stochastic Convection Extensions

The stochastic model of the evolution of convective cells can be enhanced in multiple ways. In first place, the consideration of the vertical dimension by including the estimated echo tops of the convective objects would allow for the formulation of full 4D convective avoidance problems. In second place, additional meteorological and orographic variables could be included in the set of predictors to enhance accuracy, if they prove informative. Finally, modeling the growth, formation, decay and dissipation processes could also prove valuable.

Alternatively, other information such as NWP convective ensembles or pilot avoidance models could be included in the model in order to increase the realism of the proposed trajectories.

Bibliography

- [1] 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, pp. 673–688, mar 2018.
- [2] D. González-Arribas, M. Soler, M. Sanjurjo-Rivo, M. Kamgarpour, and J. Simarro, “Robust aircraft trajectory planning under uncertain convective environments with optimal control and rapidly developing thunderstorms,” *Aerospace Science and Technology*, vol. 89, pp. 445–459, jun 2019.
- [3] D. González Arribas, M. Soler, M. Sanjurjo Rivo, J. García-Heras Carretero, D. Sacher, U. Gelhardt, J. Lang, T. Hauf, and J. Simarro, “Robust optimal trajectory planning under uncertain winds and convective risk,” in *Air Traffic Management and Systems III*, Springer Nature Singapore Pte Ltd., 2019.
- [4] D. González Arribas, M. Soler, M. Sanjurjo Rivo, J. García-Heras Carretero, D. Sacher, U. Gelhardt, J. Lang, T. Hauf, and J. Simarro, “Robust Optimal Trajectory Planning under Uncertain Winds and Convective Risk,” in *EIWAC 2017 - The 5th ENRI International Workshop on ATM/CNS*, oct 2017.
- [5] V. Courchelle, M. Soler, D. González-Arribas, and D. Delahaye, “A simulated annealing approach to 3d strategic aircraft deconfliction based on en-route speed changes under wind and temperature uncertainties,” *Transportation Research Part C: Emerging Technologies*, vol. 103, pp. 194–210, jun 2019.
- [6] D. Hentzen, M. Kamgarpour, M. Soler, and D. González-Arribas, “On maximizing safety in stochastic aircraft trajectory planning with uncertain thunderstorm development,” *Aerospace Science and Technology*, vol. 79, pp. 543–553, Aug. 2018.
- [7] D. González-Arribas, M. Soler, J. López-Leonés, E. Casado, and M. Sanjurjo-Rivo, “Automated optimal flight planning based on the aircraft intent description language,” *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, vol. 233, pp. 928–948, jan 2018.
- [8] 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*, pp. 453–466, Springer International Publishing, Sept. 2018.
- [9] V. Courchelle, D. Delahaye, D. González-Arribas, and M. Soler, “Simulated annealing for strategic traffic deconfliction by subliminal speed control under wind

- uncertainties,” in *SIDs 2017 - Proceedings of the Seventh SESAR Innovation Days*, SESAR, nov 2017.
- [10] D. González-Arribas, D. Hentzen, M. Sanjurjo-Rivo, M. Soler, and M. Kamgar-pour, “Optimal aircraft trajectory planning in the presence of stochastic convective weather cells,” in *17th AIAA Aviation Technology, Integration, and Operations Conference*, p. 3431, American Institute of Aeronautics and Astronautics, June 2017.
- [11] P. Machuca, D. González-Arribas, D. Morante-González, M. Sanjurjo-Rivo, and M. Soler, “Robust optimization of descent trajectories on irregular-shaped bodies in the presence of uncertainty,” in *Proceedings of the AAS/AIAA Astrodynamics Specialist Conference*, SESAR, aug 2017.
- [12] D. González Arribas, M. Soler, and M. Sanjurjo Rivo, “Wind-based robust trajectory optimization using meteorological ensemble probabilistic forecasts,” in *SIDs 2016 - Proceedings of the Sixth SESAR Innovation Days*, SESAR, dec 2016.
- [13] D. Gonzalez Arribas, M. Soler, and M. S. Rivo, “Wind-optimal cruise trajectories using pseudospectral methods and ensemble probabilistic forecasts,” in *Proceedings of the 5th International Conference on Application and Theory of Automation in Command and Control Systems - ATACCS '15*, pp. 160–167, ACM, ACM Press, 2015.
- [14] D. Gonzalez-Arribas, M. Sanjurjo-Rivo, and M. Soler, “Optimization of path-constrained systems using pseudospectral methods applied to aircraft trajectory planning,” *IFAC-PapersOnLine*, vol. 48, no. 9, pp. 192–197, 2015.
- [15] Eurocontrol, “European Aviation in 2040 - Challenges of Growth,” 2018.
- [16] Sims R., R. Schaeffer, F. Creutzig, X. Cruz-Núñez, M. D’Agosto, D. Dimitriu, M.J. Figueroa Meza, L. Fulton, S. Kobayashi, O. Lah, A. McKinnon, P. Newman, M. Ouyang, J.J. Schauer, D. Sperling, G. Tiwari, “Transport,” in *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge University Press, 2014.
- [17] Myhre G., D. Shindell, F.-M. Bréon, W. Collins, J. Fuglestedt, J. Huang, D. Koch, J.-F. Lamarque, D. Lee, B. Mendoza, T. Nakajima, A. Robock, G. Stephens, T. Takemura and H. Zhang, “Anthropogenic and natural radiative forcing,” in *Climate Change 2013 - The Physical Science Basis: Working Group I Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, pp. 659–740, Cambridge University Press, 2014.
- [18] A. Cook and D. Rivas, *Complexity Science in Air Traffic Management*. Routledge, June 2016.

- [19] R. Vazquez and D. Rivas, "Propagation of initial mass uncertainty in aircraft cruise flight," *Journal of Guidance, Control, and Dynamics*, vol. 36, pp. 415–429, Mar. 2013.
- [20] E. Casado, M. Vilaplana, and C. Goodchild, "Sensitivity of trajectory prediction accuracy to aircraft performance uncertainty," in *AIAA Infotech@Aerospace (I@A) Conference*, American Institute of Aeronautics and Astronautics, Aug. 2013.
- [21] E. Casado, M. Vilaplana, and C. Goodchild, "Sensitivity of trajectory prediction accuracy to aircraft performance uncertainty," in *AIAA Infotech@Aerospace (I@A) Conference*, American Institute of Aeronautics and Astronautics, Aug. 2013.
- [22] Q. Zheng and Y. Zhao, "Modeling wind uncertainties for stochastic trajectory synthesis," in *11th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, American Institute of Aeronautics and Astronautics, Sept. 2011.
- [23] A. Motodani and N. Takeichi, "Feasibility study on modeling of flight time prediction uncertainty using flight and meteorological conditions," *JOURNAL OF THE JAPAN SOCIETY FOR AERONAUTICAL AND SPACE SCIENCES*, vol. 66, no. 5, pp. 123–127, 2018.
- [24] N. Takeichi, "An adaptive model of flight time uncertainty and its application to time-based air traffic operations," in *2018 AIAA Modeling and Simulation Technologies Conference*, American Institute of Aeronautics and Astronautics, Jan. 2018.
- [25] M. B. Rhudy, "Predicting the parameters of stochastic wind models for time-varying wind estimation techniques," *Journal of Aerospace Information Systems*, vol. 16, pp. 71–76, Feb. 2019.
- [26] J. Sun, H. Vũ, J. Ellerbroek, and J. M. Hoekstra, "Weather field reconstruction using aircraft surveillance data and a novel meteo-particle model," *PLOS ONE*, vol. 13, p. e0205029, Oct. 2018.
- [27] G. Chaloulos and J. Lygeros, "Effect of wind correlation on aircraft conflict probability," *Journal of Guidance, Control, and Dynamics*, vol. 30, pp. 1742–1752, Nov. 2007.
- [28] I. Lymperopoulos and J. Lygeros, "Improved multi-aircraft ground trajectory prediction for air traffic control," *Journal of Guidance, Control, and Dynamics*, vol. 33, pp. 347–362, Mar. 2010.
- [29] S. Choi, Y. J. Kim, S. Briceno, and D. Mavris, "Prediction of weather-induced airline delays based on machine learning algorithms," in *2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC)*, IEEE, Sep. 2016.
- [30] H. Arneson, A. Bombelli, A. Segarra-Torne, and E. Tse, "Analysis of convective-weather impact on pre-departure routing decisions for flights traveling between Fort Worth Center and New York Air Center," in *17th AIAA Aviation Technology*,

- Integration, and Operations Conference*, American Institute of Aeronautics and Astronautics, jun 2017.
- [31] 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.
 - [32] B. Sridhar and N. Chen, “Short-term national airspace system delay prediction using weather impacted traffic index,” *Journal of Guidance, Control, and Dynamics*, vol. 32, pp. 657–662, mar 2009.
 - [33] E. Casado Magaña, *Trajectory prediction uncertainty modelling for Air Traffic Management*. PhD thesis, University of Glasgow, 2016.
 - [34] T. Zeh, J. Rosenow, and H. Fricke, “Interdependent uncertainty handling in trajectory prediction,” *Aerospace*, vol. 6, p. 15, feb 2019.
 - [35] M. Leutbecher and T. Palmer, “Ensemble forecasting,” *Journal of Computational Physics*, vol. 227, pp. 3515–3539, mar 2008.
 - [36] J.-P. B. Clarke, S. Solak, Y. Chang, L. Ren, and A. Vela, “Air traffic flow management in the presence of uncertainty,” in *8th USA - Europe ATM Seminar*, 2009.
 - [37] M. Steiner, R. Bateman, D. Megenhardt, Y. Liu, M. Xu, M. Pocernich, and J. Krozel, “Translation of ensemble weather forecasts into probabilistic air traffic capacity impact,” *Air Traffic Control Quarterly*, vol. 18, pp. 229–254, July 2010.
 - [38] R. Kikuchi, T. Misaka, S. Obayashi, H. Inokuchi, H. Oikawa, and A. Misumi, “Nowcasting algorithm for wind fields using ensemble forecasting and aircraft flight data,” *Meteorological Applications*, vol. 25, pp. 365–375, dec 2017.
 - [39] C. Ichard, *Random media and processes estimation using non-linear filtering techniques: application to ensemble weather forecast and aircraft trajectories*. PhD thesis, Université de Toulouse, Université Toulouse III-Paul Sabatier, 2015.
 - [40] 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.
 - [41] 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.
 - [42] W. R. Knecht, “Current nexrad cannot reliably enable safe flight around heavy weather,” *The International Journal of Aviation Psychology*, vol. 26, no. 1-2, pp. 46–61, 2016.

- [43] A. Franco, D. Rivas, and A. Valenzuela, “Probabilistic aircraft trajectory prediction in cruise flight considering ensemble wind forecasts,” *Aerospace Science and Technology*, vol. 82-83, pp. 350–362, nov 2018.
- [44] E. Hernández-Romero, A. Valenzuela, and D. Rivas, “A probabilistic approach to measure aircraft conflict severity considering wind forecast uncertainty,” *Aerospace Science and Technology*, vol. 86, pp. 401–414, mar 2019.
- [45] R. Buizza, M. Milleer, and T. N. Palmer, “Stochastic representation of model uncertainties in the ECMWF ensemble prediction system,” *Q.J.R. Meteorol. Soc.*, vol. 125, pp. 2887–2908, Aug. 2007.
- [46] P. Bauer, A. Thorpe, and G. Brunet, “The quiet revolution of numerical weather prediction,” *Nature*, vol. 525, pp. 47–55, Sept. 2015.
- [47] A. Gardi, R. Sabatini, and S. Ramasamy, “Multi-objective optimisation of aircraft flight trajectories in the ATM and avionics context,” *Progress in Aerospace Sciences*, vol. 83, pp. 1–36, may 2016.
- [48] J. Fisher and R. Bhattacharya, “Optimal trajectory generation with probabilistic system uncertainty using polynomial chaos,” *J. Dyn. Sys., Meas., Control*, vol. 133, no. 1, p. 014501, 2011.
- [49] T. Flanzer, G. Bower, and I. Kroo, “Robust trajectory optimization for dynamic soaring,” in *AIAA Guidance, Navigation, and Control Conference*, p. 4603, American Institute of Aeronautics and Astronautics, Aug. 2012.
- [50] X. Li, P. B. Nair, Z. Zhang, L. Gao, and C. Gao, “Aircraft robust trajectory optimization using nonintrusive polynomial chaos,” *Journal of Aircraft*, vol. 51, pp. 1592–1603, Sept. 2014.
- [51] I. M. Ross, R. J. Proulx, M. Karpenko, and Q. Gong, “Riemann–stieltjes optimal control problems for uncertain dynamic systems,” *Journal of Guidance, Control, and Dynamics*, vol. 38, pp. 1251–1263, July 2015.
- [52] W. B. Powell, *Approximate Dynamic Programming*, vol. 703. John Wiley & Sons, Inc., Aug. 2011.
- [53] D. P. Bertsekas, *Dynamic programming and optimal control*, vol. 1. Athena Scientific, 4th ed., 2017.
- [54] E. W. Dijkstra, “A note on two problems in connexion with graphs,” *Numerische Mathematik*, vol. 1, pp. 269–271, dec 1959.
- [55] P. Hart, N. Nilsson, and B. Raphael, “A formal basis for the heuristic determination of minimum cost paths,” *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.

- [56] A. Stentz, "Optimal and efficient path planning for partially-known environments," in *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, IEEE Comput. Soc. Press, 1994.
- [57] S. Koenig and M. Likhachev, "Fast replanning for navigation in unknown terrain," *IEEE Transactions on Robotics*, vol. 21, pp. 354–363, jun 2005.
- [58] A. Nash and S. Koenig, "Any-angle path planning," *AI Magazine*, vol. 34, p. 85, sep 2013.
- [59] M. Elbanhawi and M. Simic, "Sampling-based robot motion planning: A review," *IEEE Access*, vol. 2, pp. 56–77, 2014.
- [60] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The International Journal of Robotics Research*, vol. 30, pp. 846–894, jun 2011.
- [61] I. Noreen, A. Khan, and Z. Habib, "Optimal path planning using RRT* based approaches: A survey and future directions," *International Journal of Advanced Computer Science and Applications*, vol. 7, no. 11, 2016.
- [62] V. Gabrel, C. Murat, and L. Wu, "New models for the robust shortest path problem: complexity, resolution and generalization," *Annals of Operations Research*, vol. 207, pp. 97–120, oct 2011.
- [63] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, pp. 671–680, may 1983.
- [64] J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proceedings of ICNN'95 - International Conference on Neural Networks*, IEEE, 1995.
- [65] Y. Shi and R. Eberhart, "A modified particle swarm optimizer," in *1998 IEEE International Conference on Evolutionary Computation Proceedings. IEEE World Congress on Computational Intelligence (Cat. No.98TH8360)*, IEEE, 1998.
- [66] H.-G. Beyer and H.-P. Schwefel, "Evolution strategies - a comprehensive introduction," *Natural Computing*, vol. 1, no. 1, pp. 3–52, 2002.
- [67] R. Storn and K. Price, "Differential evolution - a simple and efficient heuristic for global optimization over continuous spaces," *Journal of Global Optimization*, vol. 11, no. 4, pp. 341–359, 1997.
- [68] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan, "A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II," in *Parallel Problem Solving from Nature PPSN VI*, pp. 849–858, Springer Berlin Heidelberg, 2000.
- [69] N. Mladenović and P. Hansen, "Variable neighborhood search," *Computers & Operations Research*, vol. 24, pp. 1097–1100, nov 1997.

- [70] J. Hart and A. W. Shogan, “Semi-greedy heuristics: An empirical study,” *Operations Research Letters*, vol. 6, pp. 107–114, jul 1987.
- [71] L. Pontryagin, *Mathematical Theory of Optimal Processes*. Classics of Soviet Mathematics, Routledge, May 2018.
- [72] R. F. Hartl, S. P. Sethi, and R. G. Vickson, “A survey of the maximum principles for optimal control problems with state constraints,” *SIAM Rev.*, vol. 37, pp. 181–218, June 1995.
- [73] J. A. Sethian and A. Vladimirovsky, “Ordered upwind methods for static Hamilton–Jacobi equations: Theory and algorithms,” *SIAM J. Numer. Anal.*, vol. 41, pp. 325–363, Jan. 2003.
- [74] U. M. Ascher, R. M. M. Mattheij, and R. D. Russell, *Numerical Solution of Boundary Value Problems for Ordinary Differential Equations*, vol. 13. Society for Industrial and Applied Mathematics, Jan. 1995.
- [75] J. T. Betts, *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*, vol. 19. Society for Industrial and Applied Mathematics, Jan. 2010.
- [76] P. E. Gill, W. Murray, and M. A. Saunders, “SNOPT: An SQP algorithm for large-scale constrained optimization,” *SIAM Rev.*, vol. 47, pp. 99–131, Jan. 2005.
- [77] A. Wächter and L. T. Biegler, “On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming,” *Math. Program.*, vol. 106, pp. 25–57, Apr. 2005.
- [78] D. Kraft, “On converting optimal control problems into nonlinear programming problems,” in *Computational Mathematical Programming*, pp. 261–280, Springer Berlin Heidelberg, 1985.
- [79] C. R. Hargraves and S. W. Paris, “Direct trajectory optimization using nonlinear programming and collocation,” *Journal of Guidance, Control, and Dynamics*, vol. 10, pp. 338–342, July 1987.
- [80] A. L. Herman and B. A. Conway, “Direct optimization using collocation based on high-order Gauss-lobatto quadrature rules,” *Journal of Guidance, Control, and Dynamics*, vol. 19, pp. 592–599, May 1996.
- [81] I. M. Ross and M. Karpenko, “A review of pseudospectral optimal control: From theory to flight,” *Annual Reviews in Control*, vol. 36, pp. 182–197, Dec. 2012.
- [82] A. V. Rao, “A survey of numerical methods for optimal control,” *Advances in the Astronautical Sciences*, vol. 135, no. 1, pp. 497–528, 2009.
- [83] G. Elnagar, M. Kazemi, and M. Razzaghi, “The pseudospectral Legendre method for discretizing optimal control problems,” *IEEE Trans. Automat. Contr.*, vol. 40, no. 10, pp. 1793–1796, 1995.

- [84] G. N. Elnagar and M. A. Kazemi, "Pseudospectral Legendre-based optimal computation of nonlinear constrained variational problems," *J. Comput. Appl. Math.*, vol. 88, pp. 363–375, Mar. 1998.
- [85] F. Fahroo and I. M. Ross, "Direct trajectory optimization by a chebyshev pseudospectral method," *Journal of Guidance, Control, and Dynamics*, vol. 25, pp. 160–166, Jan. 2002.
- [86] D. Benson, *A Gauss pseudospectral transcription for optimal control*. PhD thesis, Massachusetts Institute of Technology, 2005.
- [87] S. Kameswaran and L. T. Biegler, "Convergence rates for direct transcription of optimal control problems using collocation at radau points," *Comput Optim Appl*, vol. 41, pp. 81–126, Nov. 2007.
- [88] D. Garg, M. Patterson, C. Darby, C. Francolin, G. Huntington, W. Hager, and A. Rao, "Direct trajectory optimization and costate estimation of general optimal control problems using a radau pseudospectral method," in *AIAA Guidance, Navigation, and Control Conference*, American Institute of Aeronautics and Astronautics, Aug. 2009.
- [89] C. Canuto, M. Y. Hussaini, A. Quarteroni, and T. A. Zang, *Spectral Methods in Fluid Dynamics*. Scientific Computation, Springer Berlin Heidelberg, 1988.
- [90] F. Fahroo and I. M. Ross, "Costate estimation by a Legendre pseudospectral method," *Journal of Guidance, Control, and Dynamics*, vol. 24, pp. 270–277, Mar. 2001.
- [91] Q. Gong, I. M. Ross, W. Kang, and F. Fahroo, "Connections between the covector mapping theorem and convergence of pseudospectral methods for optimal control," *Comput Optim Appl*, vol. 41, pp. 307–335, Oct. 2007.
- [92] F. Fahroo and I. M. Ross, "Advances in pseudospectral methods for optimal control," in *AIAA Guidance, Navigation and Control Conference and Exhibit*, pp. 18–21, American Institute of Aeronautics and Astronautics, Aug. 2008.
- [93] I. M. Ross and F. Fahroo, "Pseudospectral knotting methods for solving nonsmooth optimal control problems," *Journal of Guidance, Control, and Dynamics*, vol. 27, pp. 397–405, May 2004.
- [94] A. Franco, D. Rivas, and A. Valenzuela, "Minimum-fuel cruise at constant altitude with fixed arrival time," *Journal of Guidance, Control, and Dynamics*, vol. 33, pp. 280–285, Jan. 2010.
- [95] P. Tsiotras, E. Bakolas, and Y. Zhao, "Initial guess generation for aircraft landing trajectory optimization," in *AIAA Guidance, Navigation, and Control Conference*, American Institute of Aeronautics and Astronautics, Aug. 2011.

- [96] J. T. Betts and W. P. Huffman, "Mesh refinement in direct transcription methods for optimal control," *Optimal Control Applications and Methods*, vol. 19, no. 1, pp. 1–21, 1998.
- [97] J. T. Betts, N. Biehn, S. L. Campbell, and W. P. Huffman, "Compensating for order variation in mesh refinement for direct transcription methods," *J. Comput. Appl. Math.*, vol. 125, pp. 147–158, Dec. 2000.
- [98] S. Jain and P. Tsiotras, "Trajectory optimization using multiresolution techniques," *Journal of Guidance, Control, and Dynamics*, vol. 31, pp. 1424–1436, Sept. 2008.
- [99] Y. Zhao and P. Tsiotras, "Density functions for mesh refinement in numerical optimal control," *Journal of Guidance, Control, and Dynamics*, vol. 34, pp. 271–277, Jan. 2011.
- [100] Q. Gong, F. Fahroo, and I. M. Ross, "Spectral algorithm for pseudospectral methods in optimal control," *Journal of Guidance, Control, and Dynamics*, vol. 31, pp. 460–471, May 2008.
- [101] C. L. Darby, W. W. Hager, and A. V. Rao, "An hp-adaptive pseudospectral method for solving optimal control problems," *Optim. Control Appl. Meth.*, vol. 32, pp. 476–502, Aug. 2010.
- [102] M. Schlegel and W. Marquardt, "Detection and exploitation of the control switching structure in the solution of dynamic optimization problems," *J. Process Control*, vol. 16, pp. 275–290, Mar. 2006.
- [103] S. Winderl and C. Büskens, "Exploiting the special structure for real-time control of optimal control problems with linear controls: Numerical experience," *Proc. Appl. Math. Mech.*, vol. 1, p. 484, Mar. 2002.
- [104] P. Wang, C. Yang, and Z. Yuan, "The combination of adaptive pseudospectral method and structure detection procedure for solving dynamic optimization problems with discontinuous control profiles," *Ind. Eng. Chem. Res.*, vol. 53, pp. 7066–7078, Apr. 2014.
- [105] S. Malan, M. Milanese, and M. Taragna, "Robust analysis and design of control systems using interval arithmetics," *IFAC Proceedings Volumes*, vol. 29, pp. 3422–3427, June 1996.
- [106] D. Limon, T. Alamo, E. Camacho, and J. Bravo, "Robust MPC of constrained nonlinear systems based on interval arithmetic," *IEEE Proceedings - Control Theory and Applications*, vol. 152, pp. 325–332, May 2005.
- [107] B. Øksendal, *Stochastic Differential Equations*. Springer Berlin Heidelberg, 2003.
- [108] D. J. Higham, "An algorithmic introduction to numerical simulation of stochastic differential equations," *SIAM Rev.*, vol. 43, pp. 525–546, Jan. 2001.

- [109] G. Maruyama, “Continuous Markov processes and stochastic equations,” *Rend. Circ. Mat. Palermo*, vol. 4, pp. 48–90, Jan. 1955.
- [110] G. N. Mil’shtejn, “Approximate integration of stochastic differential equations,” *Theory Probab. Appl.*, vol. 19, pp. 557–562, June 1975.
- [111] P. E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations*, vol. 23. Springer Berlin Heidelberg, 1992.
- [112] S. Brooks, A. Gelman, G. Jones, and X. Meng, *Handbook of Markov Chain Monte Carlo*. Chapman & Hall/CRC Handbooks of Modern Statistical Methods, Chapman and Hall/CRC, May 2011.
- [113] H. Niederreiter and A. Winterhof, *Quasi-Monte Carlo Methods*, ch. Quasi-Monte Carlo Methods, pp. 185–306. Springer International Publishing, 2015.
- [114] S. H. Paskov and J. F. Traub, “Faster valuation of financial derivatives,” *JPM*, vol. 22, pp. 113–123, Oct. 1995.
- [115] I. H. Sloan and H. Woźniakowski, “When are quasi-Monte Carlo algorithms efficient for high dimensional integrals?,” *Journal of Complexity*, vol. 14, pp. 1–33, Mar. 1998.
- [116] R. E. Caflisch, W. J. Morokoff, and A. B. Owen, *Valuation of mortgage backed securities using Brownian bridges to reduce effective dimension*. Department of Mathematics, University of California, Los Angeles.
- [117] S. Tezuka, “On the necessity of low-effective dimension,” *Journal of Complexity*, vol. 21, pp. 710–721, Oct. 2005.
- [118] D. Xiu, *Numerical Methods for Stochastic Computations*. Princeton University Press, July 2010.
- [119] B. J. Deusschere, H. N. Najm, P. P. Pébay, O. M. Knio, R. G. Ghanem, and O. P. Le Maître, “Numerical challenges in the use of polynomial chaos representations for stochastic processes,” *SIAM J. Sci. Comput.*, vol. 26, pp. 698–719, Jan. 2004.
- [120] Y. Matsuno, T. Tsuchiya, and N. Matayoshi, “Near-optimal control for aircraft conflict resolution in the presence of uncertainty,” *Journal of Guidance, Control, and Dynamics*, vol. 39, pp. 326–338, Feb. 2016.
- [121] D. Xiu and G. E. Karniadakis, “Modeling uncertainty in flow simulations via generalized polynomial chaos,” *J. Comput. Phys.*, vol. 187, pp. 137–167, May 2003.
- [122] S. A. Smolyak, “Quadrature and interpolation formulas for tensor products of certain classes of functions,” in *Dokl. Akad. Nauk SSSR*, vol. 4, p. 123, 1963.
- [123] A. O’Hagan, “Polynomial Chaos: A Tutorial and Critique from a Statistician’s Perspective,” 2013.

- [124] E. Novak and K. Ritter, “Simple cubature formulas with high polynomial exactness,” *Constr. Approx.*, vol. 15, pp. 499–522, Oct. 1999.
- [125] R. Cools and P. Rabinowitz, “Monomial cubature rules since “Stroud”: A compilation,” *J. Comput. Appl. Math.*, vol. 48, pp. 309–326, Nov. 1993.
- [126] R. Cools, “Monomial cubature rules since “Stroud”: A compilation — part 2,” *J. Comput. Appl. Math.*, vol. 112, pp. 21–27, Nov. 1999.
- [127] C. Litterer and T. Lyons, “Cubature on wiener space continued,” in *Stochastic Processes and Applications to Mathematical Finance*, vol. 460, pp. 169–198, The Royal Society, WORLD SCIENTIFIC, Apr. 2007.
- [128] J. Tsitsiklis, “Asynchronous stochastic approximation and q-learning,” in *Proceedings of 32nd IEEE Conference on Decision and Control*, pp. 395–400 vol.1, IEEE, dec 1993.
- [129] R. S. Sutton, “Learning to predict by the methods of temporal differences,” *Mach Learn*, vol. 3, pp. 9–44, Aug. 1988.
- [130] B. Recht, “A tour of reinforcement learning: The view from continuous control,” *Annu. Rev. Control Robot. Auton. Syst.*, vol. 2, Dec. 2018.
- [131] R. F. Stengel, *Optimal control and estimation*. Courier Corporation, 1994.
- [132] W. Fleming and R. Rishel, *Deterministic and Stochastic Optimal Control*, vol. 1. Springer New York, 1975.
- [133] H. J. Kushner and F. C. Schweppe, “A maximum principle for stochastic control systems,” *Journal of Mathematical Analysis and Applications*, vol. 8, pp. 287–302, Apr. 1964.
- [134] W. H. Fleming and H. M. Soner, *Controlled Markov Processes and Viscosity Solutions*, vol. 25. Springer-Verlag, 2006.
- [135] H. J. Kushner and P. Dupuis, *Numerical Methods for Stochastic Control Problems in Continuous Time*, vol. 24. Springer New York, 2001.
- [136] W. H. Fleming, “Exit probabilities and optimal stochastic control,” *Appl Math Optim*, vol. 4, pp. 329–346, Mar. 1977.
- [137] H. J. Kappen, “An introduction to stochastic control theory, path integrals and reinforcement learning,” *Cooperative Behavior in Neural Systems*, vol. 887, pp. 149–181.
- [138] B. Van den Broek, W. Wiegerinck, and B. Kappen, “Graphical model inference in optimal control of stochastic multi-agent systems,” *jair*, vol. 32, pp. 95–122, May 2008.

- [139] E. Theodorou, J. Buchli, and S. Schaal, “A generalized path integral control approach to reinforcement learning,” *The Journal of Machine Learning Research*, vol. 11, pp. 3137–3181.
- [140] E. Theodorou, J. Buchli, and S. Schaal, “Reinforcement learning of motor skills in high dimensions: A path integral approach,” in *2010 IEEE International Conference on Robotics and Automation*, pp. 2397–2403, IEEE, IEEE, May 2010.
- [141] M. B. Horowitz, A. Damle, and J. W. Burdick, “Linear Hamilton Jacobi bellman equations in high dimensions,” in *53rd IEEE Conference on Decision and Control*, pp. 5880–5887, IEEE, IEEE, Dec. 2014.
- [142] T. Huschto and S. Sager, “Solving stochastic optimal control problems by a wiener chaos approach,” *Viet J Math*, vol. 42, pp. 83–113, Jan. 2014.
- [143] C. Phelps, J. O. Royset, and Q. Gong, “Sample average approximations in optimal control of uncertain systems,” in *52nd IEEE Conference on Decision and Control*, pp. 1958–1965, IEEE, IEEE, Dec. 2013.
- [144] S. S. Altus, “Flight Planning – the Forgotten Field in Airline Operations,” tech. rep., Jeppesen.
- [145] “European ATM Master Plan Edition 2,” tech. rep., SESAR Joint Undertaking.
- [146] “NextGen Implementation Plan,” tech. rep., Federal Aviation Administration.
- [147] R. Dalmau and X. Prats, “Fuel and time savings by flying continuous cruise climbs,” *Transportation Research Part D: Transport and Environment*, vol. 35, pp. 62–71, mar 2015.
- [148] A. W. Hammad, D. Rey, A. Bu-Qammaz, H. Grzybowska, and A. Akbarnezhad, “Mathematical optimization in enhancing the sustainability of aircraft trajectory: A review,” *International Journal of Sustainable Transportation*, pp. 1–24, apr 2019.
- [149] J. W. Burrows, “Fuel optimal trajectory computation,” *Journal of Aircraft*, vol. 19, pp. 324–329, Apr. 1982.
- [150] P. K. A. Menon, “Study of aircraft cruise,” *Journal of Guidance, Control, and Dynamics*, vol. 12, pp. 631–639, Sept. 1989.
- [151] F. Neuman and E. Kreindler, “Minimum-fuel, three-dimensional flight paths for jet transports,” *Journal of Guidance, Control, and Dynamics*, vol. 8, pp. 650–657, Sept. 1985.
- [152] D. M. Pargett and M. D. Ardema, “Flight path optimization at constant altitude,” *Journal of Guidance, Control, and Dynamics*, vol. 30, pp. 1197–1201, July 2007.
- [153] A. Franco and D. Rivas, “Optimization of multiphase aircraft trajectories using hybrid optimal control,” *Journal of Guidance, Control, and Dynamics*, vol. 38, pp. 452–467, Mar. 2015.

- [154] M. G. Wu and A. V. Sadosky, “Minimum-Cost Aircraft Descent Trajectories with a Constrained Altitude Profile,” tech. rep., NASA, 2015.
- [155] D. Wu and Y. Zhao, “Performances and sensitivities of optimal trajectory generation for air traffic control automation,” in *AIAA Guidance, Navigation, and Control Conference*, American Institute of Aeronautics and Astronautics, Aug. 2009.
- [156] A. Valenzuela and D. Rivas, “Optimization of aircraft cruise procedures using discrete trajectory patterns,” *Journal of Aircraft*, vol. 51, pp. 1632–1640, Sept. 2014.
- [157] P. Tsiotras, E. Bakolas, and Y. Zhao, “Initial guess generation for aircraft landing trajectory optimization,” in *AIAA Guidance, Navigation, and Control Conference*, pp. 8–11, AIAA Reston, VA, American Institute of Aeronautics and Astronautics, Aug. 2011.
- [158] M. Soler, A. Olivares, and E. Staffetti, “Hybrid optimal control approach to commercial aircraft trajectory planning,” *Journal of Guidance, Control, and Dynamics*, vol. 33, pp. 985–991, May 2010.
- [159] P. Bonami, A. Olivares, M. Soler, and E. Staffetti, “Multiphase mixed-integer optimal control approach to aircraft trajectory optimization,” *Journal of Guidance, Control, and Dynamics*, vol. 36, pp. 1267–1277, Sept. 2013.
- [160] M. Soler, B. Zou, and M. Hansen, “Flight trajectory design in the presence of contrails: Application of a multiphase mixed-integer optimal control approach,” *Transportation Research Part C: Emerging Technologies*, vol. 48, pp. 172–194, Nov. 2014.
- [161] Y. Xu and X. Prats, “Effects of linear holding for reducing additional flight delays without extra fuel consumption,” *Transportation Research Part D: Transport and Environment*, vol. 53, pp. 388–397, jun 2017.
- [162] A. Murrieta-Mendoza, R. M. Botez, and A. Bunel, “Four-dimensional aircraft en route optimization algorithm using the artificial bee colony,” *Journal of Aerospace Information Systems*, vol. 15, pp. 307–334, jun 2018.
- [163] A. Murrieta-Mendoza, A. Hamy, and R. M. Botez, “Four- and three-dimensional aircraft reference trajectory optimization inspired by ant colony optimization,” *Journal of Aerospace Information Systems*, vol. 14, pp. 597–616, nov 2017.
- [164] C. Bouttier, O. Babando, S. Gadat, S. Gerchinovitz, S. Laporte, and F. Nicol, “Adaptive simulated annealing with homogenization for aircraft trajectory optimization,” in *Operations Research Proceedings*, pp. 569–574, Springer International Publishing, 2017.
- [165] A. Murrieta-Mendoza, L. Ternisien, B. Beuze, and R. M. Botez, “Aircraft vertical route optimization by beam search and initial search space reduction,” *Journal of Aerospace Information Systems*, vol. 15, pp. 157–171, mar 2018.

- [166] R. Chai, A. Savvaris, A. Tsourdos, Y. Xia, and S. Chai, "Solving multiobjective constrained trajectory optimization problem by an extended evolutionary algorithm," *IEEE Transactions on Cybernetics*, pp. 1–14, 2018.
- [167] R. F. Patron, A. Kessaci, and R. M. Botez, "Flight trajectories optimization under the influence of winds using genetic algorithms," in *AIAA Guidance, Navigation, and Control (GNC) Conference*, American Institute of Aeronautics and Astronautics, aug 2013.
- [168] S. Conway and S. Mondoloni, "An airborne conflict resolution approach using a genetic algorithm," in *AIAA Guidance, Navigation, and Control Conference and Exhibit*, American Institute of Aeronautics and Astronautics, aug 2001.
- [169] R. Vivona, D. Karr, and D. Roscoe, "Pattern-based genetic algorithm for airborne conflict resolution," in *AIAA Guidance, Navigation, and Control Conference and Exhibit*, American Institute of Aeronautics and Astronautics, aug 2006.
- [170] E. Zermelo, "Über das navigationsproblem bei ruhender oder veränderlicher windverteilung," *Z. angew. Math. Mech.*, vol. 11, no. 2, pp. 114–124, 1931.
- [171] M. R. Jardin and A. E. Bryson, "Neighboring optimal aircraft guidance in winds," *Journal of Guidance, Control, and Dynamics*, vol. 24, pp. 710–715, July 2001.
- [172] M. R. Jardin and A. E. Bryson, "Methods for computing minimum-time paths in strong winds," *Journal of Guidance, Control, and Dynamics*, vol. 35, pp. 165–171, Jan. 2012.
- [173] S. J. Bijlsma, "Optimal aircraft routing in general wind fields," *Journal of Guidance, Control, and Dynamics*, vol. 32, pp. 1025–1029, May 2009.
- [174] B. Sridhar, H. Ng, and N. Chen, "Aircraft trajectory optimization and contrails avoidance in the presence of winds," *Journal of Guidance, Control, and Dynamics*, vol. 34, pp. 1577–1584, Sept. 2011.
- [175] H. K. Ng, B. Sridhar, and S. Grabbe, "Optimizing aircraft trajectories with multiple cruise altitudes in the presence of winds," *Journal of Aerospace Information Systems*, vol. 11, pp. 35–47, Jan. 2014.
- [176] B. Girardet, L. Lapasset, D. Delahaye, and C. Rabut, "Wind-optimal path planning: Application to aircraft trajectories," in *2014 13th International Conference on Control Automation Robotics & Vision (ICARCV)*, IEEE, Dec. 2014.
- [177] B. Girardet, L. Lapasset, D. Delahaye, and C. Rabut, "Wind-optimal path planning: Application to aircraft trajectories," in *2014 13th International Conference on Control Automation Robotics & Vision (ICARCV)*, pp. 1403–1408, IEEE, IEEE, Dec. 2014.
- [178] A. Murrieta-Mendoza and R. Botez, "Lateral navigation optimization considering winds and temperatures for fixed altitude cruise using dijkstra's algorithm," in *Volume 1: Advances in Aerospace Technology*, ASME, nov 2014.

- [179] C. Kiss-Toth and G. Takacs, “A dynamic programming approach for 4D flight route optimization,” in *2014 IEEE International Conference on Big Data (Big Data)*, pp. 24–28, IEEE, IEEE, Oct. 2014.
- [180] A. Franco and D. Rivas, “Minimum-cost cruise at constant altitude of commercial aircraft including wind effects,” *Journal of Guidance, Control, and Dynamics*, vol. 34, pp. 1253–1260, July 2011.
- [181] A. Franco and D. Rivas, “Analysis of optimal aircraft cruise with fixed arrival time including wind effects,” *Aerosp. Sci. Technol.*, vol. 32, pp. 212–222, Jan. 2014.
- [182] A. Valenzuela and D. Rivas, “Analysis of Wind-Shear Effects on Optimal Aircraft Cruise,” in *Proc. of International Conferences on Research in Air Transportation (ICRAT)*.
- [183] M. Kamgarpour, M. Soler, C. J. Tomlin, A. Olivares, and J. Lygeros, “Hybrid optimal control for aircraft trajectory design with a variable sequence of modes,” vol. 44, pp. 7238 – 7243, 2011. 18th IFAC World Congress.
- [184] R. S. F. Patrón, Y. Berrou, and R. M. Botez, “New methods of optimization of the flight profiles for performance database-modeled aircraft,” *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, vol. 229, pp. 1853–1867, Dec. 2014.
- [185] R. S. Félix Patrón and R. M. Botez, “Flight trajectory optimization through genetic algorithms coupling vertical and lateral profiles,” in *Volume 1: Advances in Aerospace Technology*, pp. V001T01A048–V001T01A048, American Society of Mechanical Engineers, ASME, Nov. 2014.
- [186] 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.
- [187] G. Cottrill and F. Harmon, “Hybrid Gauss pseudospectral and generalized polynomial chaos algorithm to solve stochastic trajectory optimization problems,” in *AIAA Guidance, Navigation, and Control Conference*, p. 6572, American Institute of Aeronautics and Astronautics, Aug. 2011.
- [188] Y. Matsuno and T. Tsuchiya, “4D trajectory optimization in the presence of uncertainty,” in *2013 Aviation Technology, Integration, and Operations Conference*, pp. 1–11, American Institute of Aeronautics and Astronautics, Aug. 2013.
- [189] 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, June 2015.
- [190] H. Erzberger, T. Nikoleris, R. A. Paielli, and Y.-C. Chu, “Algorithms for control of arrival and departure traffic in terminal airspace,” *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, vol. 230, pp. 1762–1779, Feb. 2016.

- [191] R. Windhorst, M. Refai, and S. Karahan, "Convective weather avoidance with uncertain weather forecasts," in *2009 IEEE/AIAA 28th Digital Avionics Systems Conference*, pp. 3–D, IEEE, IEEE, Oct. 2009.
- [192] D. McNally, K. Sheth, C. Gong, M. Sterenchuk, S. Sahlman, S. Hinton, C. H. Lee, and F.-T. Shih, "Dynamic weather routes: Two years of operational testing at american airlines," *Air Traffic Control Quarterly*, vol. 23, pp. 55–81, jan 2015.
- [193] H. Erzberger, T. A. Lauderdale, and Y.-C. Chu, "Automated conflict resolution, arrival management, and weather avoidance for air traffic management," *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, vol. 226, pp. 930–949, Oct. 2011.
- [194] K. Sheth, A. Clymer, A. Morando, and F.-T. Shih, "Analysis of multiple flight common route for traffic flow management," in *16th AIAA Aviation Technology, Integration, and Operations Conference*, American Institute of Aeronautics and Astronautics, jun 2016.
- [195] K. D. Bilimoria, M. Hayashi, and K. Sheth, "Human-in-the-loop evaluation of dynamic multi-flight common route advisories," in *2018 Aviation Technology, Integration, and Operations Conference*, American Institute of Aeronautics and Astronautics, jun 2018.
- [196] J.-T. Chen, A. Yousefi, S. Krishna, B. Sliney, and P. Smith, "Weather avoidance optimal routing for extended terminal airspace in support of dynamic airspace configuration," in *2012 IEEE/AIAA 31st Digital Avionics Systems Conference (DASC)*, IEEE, oct 2012.
- [197] C. Taylor, S. Liu, C. Wanke, and T. Stewart, "Generating diverse reroutes for tactical constraint avoidance," *Journal of Air Transportation*, vol. 26, pp. 49–59, Apr. 2018.
- [198] H. K. Ng, S. Grabbe, and A. Mukherjee, "Design and evaluation of a dynamic programming flight routing algorithm using the convective weather avoidance model," in *AIAA Guidance, Navigation, and Control Conference*, p. 5862, American Institute of Aeronautics and Astronautics, Aug. 2009.
- [199] B. Zhang, L. Tang, and M. Roemer, "Probabilistic planning and risk evaluation based on ensemble weather forecasting," *IEEE Trans. Automat. Sci. Eng.*, vol. 15, pp. 556–566, Apr. 2018.
- [200] A. Mukherjee, B. Sridhar, and S. Grabbe, "Optimizing flight departure delay and route selection under en route convective weather," *Air Traffic Control Quarterly*, vol. 19, pp. 63–82, apr 2011.
- [201] M. C. Bartholomew-Biggs, S. C. Parkhurst, and S. P. Wilson, "Using DIRECT to solve an aircraft routing problem," *Computational Optimization and Applications*, vol. 21, no. 3, pp. 311–323.

- [202] M. Bartholomew-Biggs, S. Parkhurst, and S. Wilson, “Global optimization approaches to an aircraft routing problem,” *Eur. J. Oper. Res.*, vol. 146, pp. 417–431, Apr. 2003.
- [203] M. Kamgarpour, V. Dadok, and C. Tomlin, “Trajectory generation for aircraft subject to dynamic weather uncertainty,” in *49th IEEE Conference on Decision and Control (CDC)*, pp. 2063–2068, IEEE, Dec. 2010.
- [204] S. Summers, M. Kamgarpour, J. Lygeros, and C. Tomlin, “A stochastic reach-avoid problem with random obstacles,” in *Proceedings of the 14th international conference on Hybrid systems: computation and control - HSCC '11*, pp. 251–260, ACM, ACM Press, 2011.
- [205] A. d'Aspremont, D. Sohler, A. Nilim, L. E. Ghaoui, and V. Duong, “Optimal path planning for air traffic flow management under stochastic weather and capacity constraints,” in *2006 International Conference on Research, Innovation and Vision for the Future*, IEEE, 2006.
- [206] W. Liu and I. Hwang, “Probabilistic aircraft midair conflict resolution using stochastic optimal control,” *IEEE Trans. Intell. Transport. Syst.*, vol. 15, pp. 37–46, Feb. 2014.
- [207] A. Nilim, L. El Ghaoui, and V. Duong, “Robust dynamic routing of aircraft under uncertainty,” in *Proceedings. The 21st Digital Avionics Systems Conference*, vol. 1, pp. 1A5–1, IEEE, 2002.
- [208] A. Nilim, L. El Ghaoui, and V. Duong, “Robust dynamic routing of aircraft under uncertainty,” in *Proceedings. The 21st Digital Avionics Systems Conference*, pp. 23–27, IEEE, 2003.
- [209] A. Nilim and L. El Ghaoui, “Optimal path planning for air traffic flow management under stochastic weather and capacity constraints,” in *2006 International Conference on Research, Innovation and Vision for the Future*, vol. 4, pp. 3429–3434, IEEE, 2006.
- [210] A. Nuic, “User Manual for the Base of Aircraft Data (BADA) rev 3.11,” Tech. Rep. 13/04/16-01, Eurocontrol Experimental Centre.
- [211] 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*, pp. 1–12, IEEE, Oct. 2006.
- [212] A. Arribas, K. B. Robertson, and K. R. Mylne, “Test of a poor man’s ensemble prediction system for short-range probability forecasting,” *Mon. Wea. Rev.*, vol. 133, pp. 1825–1839, July 2005.
- [213] C. Lu, H. Yuan, B. E. Schwartz, and S. G. Benjamin, “Short-range numerical weather prediction using time-lagged ensembles,” *Wea. Forecasting*, vol. 22, pp. 580–595, June 2007.

- [214] F. Pappenberger, J. Bartholmes, J. Thielen, H. L. Cloke, R. Buizza, and A. de Roo, “New dimensions in early flood warning across the globe using grand-ensemble weather predictions,” *Geophys. Res. Lett.*, vol. 35, May 2008.
- [215] A. R. Lawrence and J. A. Hansen, “A transformed lagged ensemble forecasting technique for increasing ensemble size,” *Mon. Wea. Rev.*, vol. 135, pp. 1424–1438, Apr. 2007.
- [216] R. Buizza, “The TIGGE global, medium-range ensembles,” tech. rep., ECMWF, nov 2014.
- [217] Y.-Y. Park, R. Buizza, and M. Leutbecher, “TIGGE: Preliminary results on comparing and combining ensembles,” *Q.J.R. Meteorol. Soc.*, vol. 134, pp. 2029–2050, Oct. 2008.
- [218] F. Molteni, R. Buizza, T. N. Palmer, and T. Petroliagis, “The ECMWF ensemble prediction system: Methodology and validation,” *Q.J. Royal Met. Soc.*, vol. 122, pp. 73–119, Jan. 1996.
- [219] G. Pellerin, L. Lefaivre, P. Houtekamer, and C. Girard, “Increasing the horizontal resolution of ensemble forecasts at CMC,” *Nonlin. Processes Geophys.*, vol. 10, no. 6, pp. 463–468, 2003.
- [220] Z. Toth and E. Kalnay, “Ensemble forecasting at NCEP and the breeding method,” *Mon. Wea. Rev.*, vol. 125, pp. 3297–3319, Dec. 1997.
- [221] N. E. Bowler, A. Arribas, K. R. Mylne, K. B. Robertson, and S. E. Beare, “The MOGREPS short-range ensemble prediction system,” *Q.J.R. Meteorol. Soc.*, vol. 134, no. 632, pp. 703–722, 2008.
- [222] J.-A. García-Moya, A. Callado, P. Escribà, C. Santos, D. Santos-Muñoz, and J. Simarro, “Predictability of short-range forecasting: A multimodel approach,” *TELLUSA*, vol. 63, pp. 550–563, May 2011.
- [223] I.-L. Frogner, H. Haakenstad, and T. Iversen, “Limited-area ensemble predictions at the Norwegian meteorological institute,” *Q.J.R. Meteorol. Soc.*, vol. 132, pp. 2785–2808, Oct. 2006.
- [224] C. Marsigli, A. Montani, and T. Paccagnella, “A spatial verification method applied to the evaluation of high-resolution ensemble forecasts,” *Met. Apps*, vol. 15, no. 1, pp. 125–143, 2008.
- [225] L. Descamps, C. Labadie, A. Joly, E. Bazile, P. Arbogast, and P. Cébron, “PEARP, the météo-France short-range ensemble prediction system,” *Q.J.R. Meteorol. Soc.*, vol. 141, pp. 1671–1685, Dec. 2014.
- [226] T. Iversen, A. Deckmyn, C. Santos, K. Sattler, J. B. Bremnes, H. Feddersen, and I.-L. Frogner, “Evaluation of ‘GLAMEPS’—a proposed multimodel EPS for short range forecasting,” *Tellus A: Dynamic Meteorology and Oceanography*, vol. 63, pp. 513–530, Jan. 2011.

- [227] A. Horányi, M. Mile, and M. Szücs, “Latest developments around the ALADIN operational short-range ensemble prediction system in Hungary,” *TELLUSA*, vol. 63, pp. 642–651, May 2011.
- [228] K. A. Emanuel, *Atmospheric convection*. Oxford University Press, 1994.
- [229] J. R. Holton and G. J. Hakim, *An Introduction to Dynamic Meteorology*, vol. 88. Elsevier, 2013.
- [230] Y.-L. Lin, *Mesoscale Dynamics*. Cambridge University Press, 2007.
- [231] F. Molteni, T. Stockdale, M. Balmaseda, G. Balsamo, R. Buizza, L. Ferranti, L. Magnusson, K. Mogensen, T. Palmer, and F. Vitart, *The new ECMWF seasonal forecast system (System 4)*. 2011.
- [232] Y. Wang, M. Bellus, C. Wittmann, M. Steinheimer, F. Weidle, A. Kann, S. Iivatek-Šahdan, W. Tian, X. Ma, S. Tascu, and E. Bazile, “The central European limited-area ensemble forecasting system: Aladin-laef,” *Q.J.R. Meteorol. Soc.*, vol. 137, pp. 483–502, Jan. 2011.
- [233] J. W. Wilson, N. A. Crook, C. K. Mueller, J. Sun, and M. Dixon, “Nowcasting thunderstorms: A status report,” *Bull. Amer. Meteor. Soc.*, vol. 79, pp. 2079–2099, Oct. 1998.
- [234] J. E. Evans and E. R. Ducot, “Corridor integrated weather system,” *Lincoln Laboratory Journal*, vol. 16, no. 1, p. 59, 2006.
- [235] J. Sun, M. Xue, J. W. Wilson, I. Zawadzki, S. P. Ballard, J. Onvlee-Hooimeyer, P. Joe, D. M. Barker, P.-W. Li, B. Golding, M. Xu, and J. Pinto, “Use of NWP for nowcasting convective precipitation: Recent progress and challenges,” *Bull. Amer. Meteor. Soc.*, vol. 95, pp. 409–426, Mar. 2014.
- [236] Y. Hwang, A. J. Clark, V. Lakshmanan, and S. E. Koch, “Improved nowcasts by blending extrapolation and model forecasts,” *Wea. Forecasting*, vol. 30, pp. 1201–1217, Oct. 2015.
- [237] J. Pinto, W. Dupree, S. Weygandt, M. Wolfson, S. Benjamin, and M. Steiner, “Advances in the consolidated storm prediction for aviation (CoSPA),” tech. rep., MASSACHUSETTS INST OF TECH LEXINGTON LINCOLN LAB, 2010.
- [238] R. Osinski and F. Bouttier, “Short-range probabilistic forecasting of convective risks for aviation based on a lagged-average-forecast ensemble approach,” *Met. Apps*, vol. 25, pp. 105–118, Nov. 2017.
- [239] M. Sauer, T. Hauf, and C. Forster, “Uncertainty Analysis of Thunderstorm Nowcasts for Utilization in Aircraft Routing,” *Fourth SESAR Innovation Days, Madrid, Spain*, 2014.

- [240] L. Han, J. Sun, W. Zhang, Y. Xiu, H. Feng, and Y. Lin, “A machine learning nowcasting method based on real-time reanalysis data,” *J. Geophys. Res. Atmos.*, vol. 122, pp. 4038–4051, Apr. 2017.
- [241] S. Xingjian, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-c. Woo, “Convolutional LSTM network: A machine learning approach for precipitation nowcasting,” in *Advances in neural information processing systems*, pp. 802–810, 2015.
- [242] R. DeLaura, M. Robinson, M. Pawlak, and J. Evans, “Modeling convective weather avoidance in enroute airspace,” Citeseer.
- [243] R. DeLaura, B. Crowe, R. Ferris, J. Love, and W. Chan, “Comparing convective weather avoidance models and aircraft-based data,” 2008.
- [244] M. Matthews and R. DeLaura, “Evaluation of enroute convective weather avoidance models based on planned and observed flights,” in *14th American Meteorological Society Conference on Aviation, Range, and Aerospace Meteorology*, Citeseer, 2010.
- [245] K. Sheth, T. Amis, S. Gutierrez-Nolasco, B. Sridhar, and D. Mulfinger, “Development of a probabilistic convective weather forecast threshold parameter for flight-routing decisions,” *Wea. Forecasting*, vol. 28, pp. 1175–1187, Oct. 2013.
- [246] S.-C. Wu, C. L. Gooding, A. E. Shelley, C. G. Duong, and W. W. Johnson, “Pilot convective weather decision making in en route airspace,” 2012.
- [247] T. Flanzer, G. Bower, and I. Kroo, “Robust trajectory optimization for dynamic soaring,” in *AIAA Guidance, Navigation, and Control Conference*, American Institute of Aeronautics and Astronautics, Aug. 2012.
- [248] D. Bertsekas and J. Tsitsiklis, “Neuro-dynamic programming: An overview,” in *Proceedings of 1995 34th IEEE Conference on Decision and Control*, vol. 1, pp. 560–564, IEEE Publ. Piscataway, NJ, IEEE, 1995.
- [249] D. Garrido-López, R. G. Ledesma, G. Gershzhohn, and S. Moore, “Analysis of aircraft descent predictability: Implications for continuous four-dimensional navigation,” in *AIAA Guidance, Navigation, and Control Conference*, American Institute of Aeronautics and Astronautics, aug 2011.
- [250] G. Chaloulos, E. Crück, and J. Lygeros, “A simulation based study of subliminal control for air traffic management,” *Transportation Research Part C: Emerging Technologies*, vol. 18, pp. 963–974, Dec. 2010.
- [251] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, “CasADi: A software framework for nonlinear optimization and optimal control,” *Math. Prog. Comp.*, vol. 11, pp. 1–36, July 2018.
- [252] W. E. Hart, J.-P. Watson, and D. L. Woodruff, “Pyomo: Modeling and solving mathematical programs in python,” *Math. Prog. Comp.*, vol. 3, pp. 219–260, Aug. 2011.

- [253] “HSL. A collection of Fortran codes for large scale scientific computation. <http://www.hsl.rl.ac.uk/>.”

Tracking Feasibility for Linear Systems

In this appendix, we discuss the conditions under which a tracking scheme is feasible for a linear system (if f and h are regular enough, similar arguments can be made in the nonlinear case replacing the matrices by the jacobians of the corresponding functions). We neglect here the inequality constraints, as they fulfill a similar role as in deterministic optimal control (with the additional implication that they must be fulfilled for every realization of ξ). Consider the linear system:

$$\dot{\mathbf{x}} = A\mathbf{x} + B_z\mathbf{z} + B_u\mathbf{u} + C\xi$$

$$0 = D\mathbf{x} + E_z\mathbf{z} + E_u\mathbf{u} + F\xi,$$

and partition the state, algebraic and control vectors into tracked and non-tracked parts (with q denoting the tracked part and r the untracked part):

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}_q \\ \mathbf{x}_r \end{bmatrix}, \mathbf{z} = \begin{bmatrix} \mathbf{z}_q \\ \mathbf{z}_r \end{bmatrix}, \mathbf{u} = \begin{bmatrix} \mathbf{u}_q \\ \mathbf{u}_r \end{bmatrix},$$

with $\mathbf{x}_q \in \mathbb{R}^{q_x}$, $\mathbf{x}_r \in \mathbb{R}^{n_x - q_x}$, $\mathbf{z}_q \in \mathbb{R}^{q_z}$, $\mathbf{z}_r \in \mathbb{R}^{n_z - q_z}$, $\mathbf{u}_q \in \mathbb{R}^{q_u}$, and $\mathbf{u}_r \in \mathbb{R}^{n_u - q_u}$

Let us define

$$\mathbf{y}_q = \begin{bmatrix} \mathbf{z}_q \\ \mathbf{u}_q \end{bmatrix}, \mathbf{y}_r = \begin{bmatrix} \mathbf{z}_r \\ \mathbf{u}_r \end{bmatrix}, \text{ and } \mathbf{y} = \begin{bmatrix} \mathbf{y}_q \\ \mathbf{y}_r \end{bmatrix}$$

and rearrange the original system as

$$\dot{\mathbf{x}} = A\mathbf{x} + B\mathbf{y} + C\xi,$$

$$0 = D\mathbf{x} + E\mathbf{y} + F\xi.$$

We partition this system in the same form, according to tracked and untracked part:

$$\begin{bmatrix} \dot{\mathbf{x}}_q \\ \dot{\mathbf{x}}_r \end{bmatrix} = \begin{bmatrix} A_{qq} & A_{qr} \\ A_{rq} & A_{rr} \end{bmatrix} \begin{bmatrix} \mathbf{x}_q \\ \mathbf{x}_r \end{bmatrix} + \begin{bmatrix} B_{qq} & B_{qr} \\ B_{rq} & B_{rr} \end{bmatrix} \begin{bmatrix} \mathbf{y}_q \\ \mathbf{y}_r \end{bmatrix} + \begin{bmatrix} C_q \\ C_r \end{bmatrix} \xi \quad (\text{A.1})$$

$$0 = \begin{bmatrix} D_q & D_r \end{bmatrix} \begin{bmatrix} \mathbf{x}_q \\ \mathbf{x}_r \end{bmatrix} + \begin{bmatrix} E_q & E_r \end{bmatrix} \begin{bmatrix} \mathbf{y}_q \\ \mathbf{y}_r \end{bmatrix} + F\xi. \quad (\text{A.2})$$

For the tracking scheme to be feasible, we need to be able to track a trajectory for different values of the uncertainty ξ ; since this system is linear, it is enough to determine if the trajectory $\mathbf{x}_q(t) = 0, \mathbf{y}_q(t) = 0$ can be tracked. By plugging this tracking trajectory in Equations (A.1)(A.2), we can observe that this is true if, for every pair of values that (ξ, \mathbf{x}_r) attains in the complete trajectory, there is a value of \mathbf{y}_r that solves the linear system:

$$\begin{bmatrix} B_{qr} \\ E_r \end{bmatrix} \mathbf{y}_r = - \begin{bmatrix} C_q \\ F \end{bmatrix} \xi - \begin{bmatrix} A_{qr} \\ D_r \end{bmatrix} \mathbf{x}_r.$$

The matrix $Q := \begin{bmatrix} B_{qr} \\ E_r \end{bmatrix} \in \mathbb{R}^{n_y - q_y \times n_y - q_y}$ is square; therefore, a sufficient (but not necessary) condition for the feasibility of a tracking scheme is Q being full-rank. Since B_{qr} represents the dependence of the tracked states on the untracked controls and algebraic states and E_r represents the dependence of the algebraic conditions on the untracked controls and algebraic states, an intuitive interpretation of this condition is the capacity of choosing the algebraic conditions and the derivatives of the tracked states with the untracked controls and algebraic states, which implies that the impact of the uncertainty on the tracked states can be fully compensated without violating the algebraic condition.

