CRANFIELD UNIVERSITY

QUINTAIN MCENTEGGART

ON THE GENERATION OF ENVIRONMENTALLY EFFICIENT FLIGHT TRAJECTORIES

SCHOOL OF AEROSPACE, TRANSPORT and MANUFACTURING (SATM)

PhD THESIS

CRANFIELD UNIVERSITY

School of Aerospace, Transport and Manufacturing (SATM)

Ph.D. Thesis

Academic Year 2015-16

Quintain McEnteggart

On the Generation of Environmentally Efficient Flight Trajectories

Supervisor: Dr. James F. Whidborne

May 2015

Abstract

To achieve a sustainable future for air transport, the International Civil Aviation Organization has proposed goals for reductions in community noise impact, local air quality and climate impacting emissions. The goals are intended to be achieved through advances in engine design, aircraft design and through improvements in aircraft operational procedures.

This thesis focuses on operational procedures, and considers how trajectory generation methods can be used to support flight and airspace planners in the planning and delivery of environmentally efficient flight operations.

The problem of planning environmentally efficient trajectories is treated as an optimal control problem that is solved through the application of a direct method of trajectory optimisation combined with a stochastic Non Linear Programming (NLP) solver. Solving the problem in this manner allows decision makers to explore the relationships between how aircraft are operated and the consequent environmental impacts of the flights.

In particular, this thesis describes a multi-objective optimisation methodology intended to support the planning of environmentally efficient climb and descent procedures. The method combines environmental, trajectory and NLP methods to generate Pareto fronts between several competing objectives. It is shown how Pareto front information can then be used to allow decision makers to make informed decisions about potential trade-offs between different environmental goals. The method is demonstrated through its application to a number of real world, many objective procedure optimisation studies. The method is shown to support in depth analysis of the case study problems and was used to identify best balance procedure characteristics and procedures in an objective, data driven approach not achievable through existing methods.

Driven by operator specific goals to reduce CO2 emissions, work in this thesis also looks at trajectory based flight planning of CO2 efficient trajectories. The results are used to better understand the impacts of ATM constraints and recommended procedures on both the energy management and fuel efficiency of flights. Further to this, it is shown how trajectory optimisation methods can be applied to the analysis of conventional assumptions on fuel efficient aircraft operations.

While the work within is intended to be directly relevant to the current air traffic management system, both consideration and discussion is given over to the evolution and continued relevance of the work to the Single European Sky trajectory based concept of operation.

Acknowledgements

I would like to thank my supervisor James Whidborne, whose generosity of knowledge, good humour and quite possibly infinite patience have seen me all the way through to the end of this work. Our early talks were the first time I felt I was having a conversation on trajectory optimisation, and they opened the door to me finally understanding so many new things. Later on, you were always supportive in letting me find my own particular path to the end. I shall be forever grateful for all your kindnesses.

To Ken Lai, if James opened the door to trajectory optimisation for me, then you helped me take the first tentative steps through it. In those early days, whenever I got stuck you had a wonderful habit of giving me just the boost I needed to get going again. You once told me you'd like to stay in academia, to be a lecturer, a supervisor and a teacher in your own right. I really hope you get a chance to do this, as, speaking from experience, I think you're going to be pretty good at it.

To Rick Drury, thank you for letting me build upon so many of your good ideas. Not only letting me, but actively encouraging me to run with them. A great deal of the work in this thesis was shaped by how generous you were with both your time and your thoughts.

Thanks to Sam Prince, whose kind words have opened up so many opportunities to me.

To Robert Christie, Klementina Gerova and Olivier Bouzigues, who have been along on this journey with me from the very beginning to the very end, thanks for the good company and all the laughter. To all my other friends, who havent heard from me in a while, but who never seem to give up on me anyway. I hope these words find you well and I hope I get a chance to see you soon. Thanks to Roy Chamberlain, Craig Wilde and my extended family in the fitness centre. Thanks also to my housemates, my office mates and to the various characters I got to know along the way.

To my parents Oliver and Rosealeen. I know I dont say this often, actually, I know I dont say this ever, but given that this is sort of a special occasion, I'd like to say that I love you both very much. Thank you for all the support with my studies, for always being there, and for all the things you've done through the years to make me the person I am today.

Contents

List of Figures

List of Tables

List of Algorithms

Chapter 1

Introduction

1.1 Motivation

Over the decade 2003 to 2013, global passenger air traffic increased by more than 70% [1]. Even factoring in the global economic downturn, continued high levels of traffic growth are projected for the coming decades [1, 2]. This growth however has come with an environmental cost. Increasing traffic levels, relying on greater consumption of fossil fuels, have led to increased levels of aircraft emissions, impacting climate change and local air quality [3]. Aviation's continued and rapid growth has seen it become the mode of transport with the fastest growing climate change impact [4, 5, 6]. Increasing traffic levels have also, despite an increasingly quiet aircraft fleet, led to an increase in the number of people exposed to significant levels of aircraft noise [7].

To achieve a sustainable future for air transport, the International Civil Aviation Organization (ICAO) has proposed 3 high level environmental goals for international aviation [7]:

- to limit or reduce the number of people affected by significant aircraft noise,
- to limit or reduce the impact of aircraft engine emissions on local air quality,
- to limit or reduce the impact of aviation greenhouse gas emissions on the global climate.

In Europe, ICAOs goal of limiting or reducing aviation related greenhouse gas emissions is complemented by the European Union Emissions Trading Scheme (EU ETS)[8, 9]. The EU ETS is a market-based cap and trade system, where a cap is set on the total level of CO2 emissions and emitters are allocated permits to emit CO2 within a progressively reducing cap [10]. The system regulates CO2 emissions from energy intensive industries, which as of 2012, includes the aviation industry. The system aims to create a market-based incentive for efficient operations by allowing efficient airline operators with a surplus

of permits to sell excess permits to inefficient operators with a deficit of permits [10]. A worldwide, CO2-based, aviation emissions trading scheme has also been proposed by ICAO for implementation in 2020 [11, 12].

More long term strategic goals for CO2 reduction have been proposed by the Advisory Council for Aeronautics Research in Europe (ACARE) and the Single European Sky ATM Research (SESAR) programme for the years 2020 through to 2050. ACARE, a group of leading aviation stakeholders from industry, academia and the European Commission, have created the Strategic Research and Innovation Agenda (SRIA) [13, 14]. The SRIA is a high level roadmap for employing technology to meet the societal, economic, environmental and safety challenges facing the aviation industry in the coming decades. The ACARE goals are to be achieved through changes to aircraft airframes, engines and operational procedures. ACARE have proposed, from a 2000 baseline, a target of 10% improvement in the operational CO2 efficiency of flights [15].

Improving Air Traffic Management (ATM) related aircraft operations is the aim of SESAR, which is the research and development initiative of the Single European Sky (SES). The SES targets are based around a new trajectory centric concept of operations that is enabled through the development and adoption of new air traffic controller and flight crew support tools, along with developments in communication, navigation and surveillance technologies. The high level SESAR targets are from a 2005 baseline [2],

- to enable a 3-fold increase in capacity which will also reduce delays both on the ground and in the air,
- to improve safety by a factor of 10,
- to enable a 10 % reduction in the effects flights have on the environment,
- to provide ATM services to the airspace users at a cost of at least 50% less.

Therefore, aligning with ACARE, the most specific environmental goal of SESAR is a 10% reduction in carbon dioxide emissions per flight (from a 2005 baseline) [16]. SESAR does not have noise and air-quality targets. Note that this is not because they are not considered to be significant and important environmental impacts [17, 18, 19]. The lack of targets reflects that there can be significant trade-offs between aircraft related environmental impacts. This is particularly true within the terminal control area, where trade-offs between environmental impacts currently need to be managed on case by case basis [17]. Therefore, there is a difficulty in developing complementary European wide environmental performance goals [17]. While there may not currently be SESAR related targets, EU legislation on aircraft noise mitigation [20, 21], local air quality targets and limit values [22], will continue to dictate that aircraft operations be conducted in a manner that supports the mitigation of different forms of environmental impact.

Therefore, through the aforementioned goals, targets and legislation, there is a clear mandate for Air Navigation Service Providers (ANSPs) and aircraft operators to mitigate environmental impact through desirable and achievable improvements in aircraft operations.

1.2 Goals

This thesis is concerned with the generation of environmentally efficient flight trajectories. The goal of the work is to use trajectory generation methods based on optimal control to support air navigation service providers and aircraft operators in the planning and delivery of environmentally efficient aircraft flight operations. The problem of planning environmentally efficient trajectories is treated as an optimal control problem that is solved through the application of a direct method of trajectory optimisation combined with a stochastic Non Linear Programming (NLP) solver. Solving the problem in this manner allows decision makers to explore the relationships between how aircraft are operated and the consequent environmental impacts of the flights. This information can then be used to support

- airspace designers in environmentally optimising the ATM system constraints (see Section 2.6.3),
- airline operators in planning environmentally efficient flight trajectories within the ATM system constraints (see Section 2.6.4).

A particular focus of work in this thesis is the use of multi-objective trajectory generation to manage necessary trade-offs between conflicting environmental impacts within the terminal area. In this work, environmentally optimised arrival and departure procedures are developed, where multi-objective trajectory optimisation is used to generate Pareto fronts. Pareto fronts, as used in this work, are multi-dimensional plots that allow for the identification of trajectory solutions that provide the best trade-offs between competing environmental goals.

Historically, ANSPs and operators have planned environmentally efficient climbs and descents separately, with ANSPs planning the horizontal routing and operators defining the aircraft operating steps along the routing [23]. This fragmented approach has prevented the full realisation of potential environmental benefits from the resulting combined operating procedure [23]. The multi-objective method proposed in this thesis offers a harmonised approach, simultaneously optimising both routing and aircraft operating steps in determining the most environmentally efficient climb or descent operating procedures. Combining the trajectory generation and environmental methods with full Pareto front analysis has also allowed trade-offs in environmental objectives to be assessed in a far more objective, data driven, manner than was previously possible.

Driven by goals to reduce CO2 emissions, work in this thesis also considers the flight planning problem and how operators can plan CO2 efficient flight trajectories. It further considers what impact ATM restrictions have on the fuel efficiency of a flight trajectory. Although the work is intended to be relevant for the current day ATM system, close consideration is given to the continued relevance of the work within the SES trajectory based concept of operation.

1.3 Thesis Structure

By chapter, the thesis it arranged as follows:

Chapter 2 discusses aircraft flight trajectories within the context of the air traffic system. It examines how constraints come to exist within the system and what impact those constraints have on the environmental efficiency of flights. The role that trajectory optimisation methods have in improving the environmental efficiency of flights is then discussed, particularly in terms of meeting a need for improved approaches to constraint definition and flight planning.

Chapters 3 & 4 discusses the different environmental metrics and presents the environmental modelling methods used in the calculation of aircraft emissions and noise impact. It discusses the use of the Annoyance Score single event noise impact measure. This recently developed metric [24] consolidates aircraft community noise impact into a single value, making it well-suited to trajectory optimisation case studies. However, the application of the metric for this purpose has not been explored prior to work in this thesis.

Chapter 5 provides an overview of trajectory generation techniques and how they are used to convert the optimal control problem into a non-linear programming problem. It discusses the properties of different direct methods and describes the differences between direct and indirect methods. The chapter highlights the Inverse Dynamics in the Virtual Domain method (IDVD), and discusses its adoption as an approach suited to the generation of environmentally efficient trajectories.

Chapter 6 provides an overview of Non Linear Programming (NLP) techniques. It highlights the stochastic Differential Evolution (DE) algorithm and its prior success when combined with the IDVD method (IDVD-DE). The chapter considers the further use of DE with the inverse method and the extensions required to apply the IDVD-DE approach to the many-objective trajectory optimisation problems considered in this thesis.

Chapter 7 defines the Pareto front analysis approach used in the work. It also assess the performance of the IDVD-DE method on a simple multi-objective environmental trajectory optimisation problem. The IDVD-DE method is then applied to a noise abatement optimisation problem with a known solution. The known solution was developed by the Sourdine project, a leading project in the field of noise abatement trajectory operations.

Chapter 8 sees the IDVD-DE method applied to two real world environmental procedure optimisation case studies. Previous multi-environmental-objective trajectory optimisation studies in the literature only considered very idealised scenarios. The use of real world case studies in this thesis set demanding requirements on the number and type of environmental objectives that needed to be considered in the trade-off analysis. By applying the IDVD-DE method to real world case studies it can be determined whether the proposed data driven approach could identify procedures that provide better trade-offs between the environmental impacts than those proposed by current best practice approaches.

Chapter 9 examines the applicability of the IDVD-DE method to problems beyond climb and descent trajectory optimisation. It investigates the origin-destination planning of CO2 efficient short-haul trajectories through the perspective of the 3Di flight efficiency measure. The results highlight how inefficiencies introduced by ATM constraints can be quantified and show that the most fuel/CO2 efficient trajectory may not be the airlines preferred trajectory. This is due to trade-offs in fuel, operating and maintenance costs. It is proposed, for future work, that multi-objective trajectory generation methods be further applied to better understand the trade-offs in airline environmental and operating cost performance.

The conclusions summarise how trajectory generation has been used in this work to propose a method useful to air traffic route designers and to airline flight planners in predicting and optimising the environmental impact of commercial aircraft trajectory operations. It summarises some of the non-intuitive results presented in the thesis and discusses the implications of these results with regard to how ATM constraints and flight operations are typically planned.

1.4 Contributions

The contributions to knowledge which have been made as part of this work are summarized below:

- By defining the use of a direct multi-objective trajectory generation method, which supports the calculation of Pareto fronts between several competing environmental objectives, environmental trajectory optimisation research has been advanced.

- A new data driven approach that supports airspace designers and flight planners in determining the most environmentally efficient routing and aircraft operating steps within the TMA has been developed.

- A harmonised approach to procedure design that simultaneously considers both route planning and aircraft operating steps has been developed.

- The usability of the noise Annoyance Score metric for environmental trajectory optimisation studies has been investigated. The metric has been shown to be a useful and needed measure for assessing community noise impact.

- The Inverse method and the Differential Evolution solver have been applied to a new problem.

- How trajectory optimisation methods can be used to better evaluate the impact of ATM constraints on aircraft flight efficiency has been highlighted.

Chapter 2

Environmentally Efficient Flight Trajectories

2.1 Introduction

When considering the environmental efficiency of commercial aircraft flight trajectories, the inefficiencies introduced to trajectories by the ATM system must first be considered. The following sections first examine trajectory constraints from an air traffic control perspective and consider how ATM related flight constraints come to exist in the first place. The ATM imposed flight constraints are then looked at from an aircraft operators perspective, by considering how ATM constraints typically limit flown flight trajectories within European airspace.

Near and long term approaches for improving the environmental efficiency of trajectories are then considered, including planned developments in flight crew and controller support tools, along with improvements in communication, navigation and surveillance technologies. These are largely described through the perspective of their planned deployment as part the Single European Sky (SES) initiative.

It is then discussed how trajectory generation methods have a role to play, both in the current and SES ATM systems, towards improving the environmental operation of aircraft by enabling better ATM constraint definition and flight planning.

2.2 Flight Efficiency - The ATM Perspective

When considering how the European ATM system introduces environmental inefficiencies into flight trajectories, the purpose of the ATM system and how this factors into the

system design must first be considered. The purpose of Air Traffic Control (ATC) is to provide a navigation support service that prevents collisions between aircraft and allows the safe and expeditious flow of air traffic [25]. The ATM system is then the systemised management of people and resources in the employ of delivering this ATC service [25].

Airspace globally is divided into Flight Information Regions. FIRs are regions where information, such as meteorological and traffic information is available to aircraft to support the safe and efficient conduct of flights [25]. Within FIRs, the airspace is divided into classes that define how separation assurance is achieved and the level of equipment required for the aircraft to access each airspace class. Controlled airspace encompasses the classes of airspace where ATC services are available and therefore where ATC have responsibility for providing separation assurance [26]. Strategic flight efficiency goals are principally concerned with commercial air traffic operating under Instrument Flight Rules in European controlled airspace. Therefore, the work in this thesis will similarly be concerned with commercial air traffic operating in European controlled airspace. With this in mind, the airspace within the European ATM system is designed and managed around finding a balance between the three strategic goals of safety, capacity and flight efficiency/environmental impact [19]. How each of these goals individually impacts the shape and operation of the airspace and ultimately the constraints that an aircraft trajectory may encounter are now discussed.

Safety

The shortest achievable route between an origin and destination airport is constrained by the design of the route network. Historically, the route network was structured with reference to aircraft navigational limitations and to enable air traffic control to provide separation assurance with the tools that were available [18]. As technology for both aircraft and ATM has improved, the need for such a rigid en-route structure has diminished [18]. However, for busy airspace, to ensure safety, flights are still concentrated onto routes, which create corridors of air traffic that can be more easily managed by human air traffic controllers. The corridors create consistent repeatable flows of traffic where traffic deviations from the norm can quickly be identified and corrected for by air traffic controllers.

The airspace itself is divided into zones of responsibility, or sectors, such that individual controllers only monitor an aircraft on a portion of its route [27]. Beyond the subdivision of controlled airspace into routes and sectors, a safe airspace sector configuration is one that minimises the risk of a loss of minimum separation or a collision between aircraft. Key safety indicators related to risk include level busts, airspace infringements and civil military interactions [28].

Level busts occur when aircraft continue to climb or descend to flight levels beyond the flight level they have been cleared to by ATC. Airspace infringements occur when aircraft

enter into controlled airspace without the knowledge and clearance from the controlling authority [28]. Related to airspace infringements is the interaction between civil and military aircraft. Civil and military traffic are controlled independently by civil and military controlling authorities, with military flights often operating in a non-structured, not easily predicted, manner. Therefore, safety of flights can be compromised if civil and military aircraft infringe upon on each others airspace.

The airspace is then designed to mitigate risk in a number of ways. The airspace configuration is designed so that if aircraft deviate from their anticipated trajectory then recovery is possible before a serious safety incident can occur. This includes, minimising the potential for level busts and airspace infringements by carefully choosing the placement of routes relative to each other and relative to sector boundaries [28, 19].

Another key to maintaining safe airspace is the moderation of controller workload, where workload refers to the busyness of the controller. The goal is for sector controllers to be efficiently utilised, neither under nor over worked, as both these scenarios can result in errors in traffic management [28, 19]. Controller workload can be moderated by minimising route conflict points and by carefully systemising the operation of a busy sector (see **Capacity**).

During the operation of a sector, where controller workload is predicted to be exceeded, then flow control measures are applied [27]. This involves reducing traffic demand on the sector by routing flights away from the sector, either horizontally (re-routing) or vertically (capping), or delaying inbound traffic at departure airports.

Therefore the geometrical shape of the sector, the routes placed within it and the traffic levels on those routes are all configured so that the sector can be safely and efficiently overseen by sector controllers. It can already be seen at this stage however, that the need to maintain high levels of safety within the system imposes constraints on where aircraft can fly and in what densities.

Capacity

Sector capacity is generally defined as the number of aircraft that can safely pass through a given piece of airspace in a given time period [29]. Sector capacity is usually declared as the maximum number of aircraft that can be controlled by a sector per hour and is a value largely determined by controller workload [29]. The capacity of a sector is established through a process of fast time simulation, real time simulation and operational review.

When the airspace capacity is predicted to be exceeded, then controlling authorities, and in particular Air Traffic Flow Management (ATFM), act to adjust traffic flows in a manner that alleviates the predicted overload while minimising total network delay. For individual aircraft however this can result in temporal and spatial constraints in the form of ground delays at departure airports, rerouting of aircraft away from their initial

flight plan or the allocation of inefficient flight levels (flight level capping) [27, 30]. The classic way of adding more capacity is to add more sectors, thereby adding more air traffic controllers to share the burden of managing the traffic demand [19]. However, this strategy has a point of diminishing returns, where further sub-division of airspace no longer provides capacity gains [19].

The capacity of a sector itself is defined by the physical characteristics of the airspace and typical levels of controller workload required to operate the sector. Therefore sectors are designed such that their shape and volume are capable of handling the projected volumes of traffic, but not so large and busy that they cannot be efficiently managed by sector controllers [19].

It has been determined that there is an inverse relationship between the number of conflict points within a sector and sector capacity, so there is an incentive to keep routes separated especially with regards to mixed climbing and descending traffic that requires high levels of controller supervision [19] (See Figure 2.1). To achieve high levels of capacity, the operation of a sector is designed to be as systemised as possible, with well designed standing agreements, which define standardised flight levels, speeds and headings for aircraft presented from one sector to the next, acting to streamline the management of the airspace and the Radio Transmission (RT) necessary between controllers and also between controllers and the flight crew.

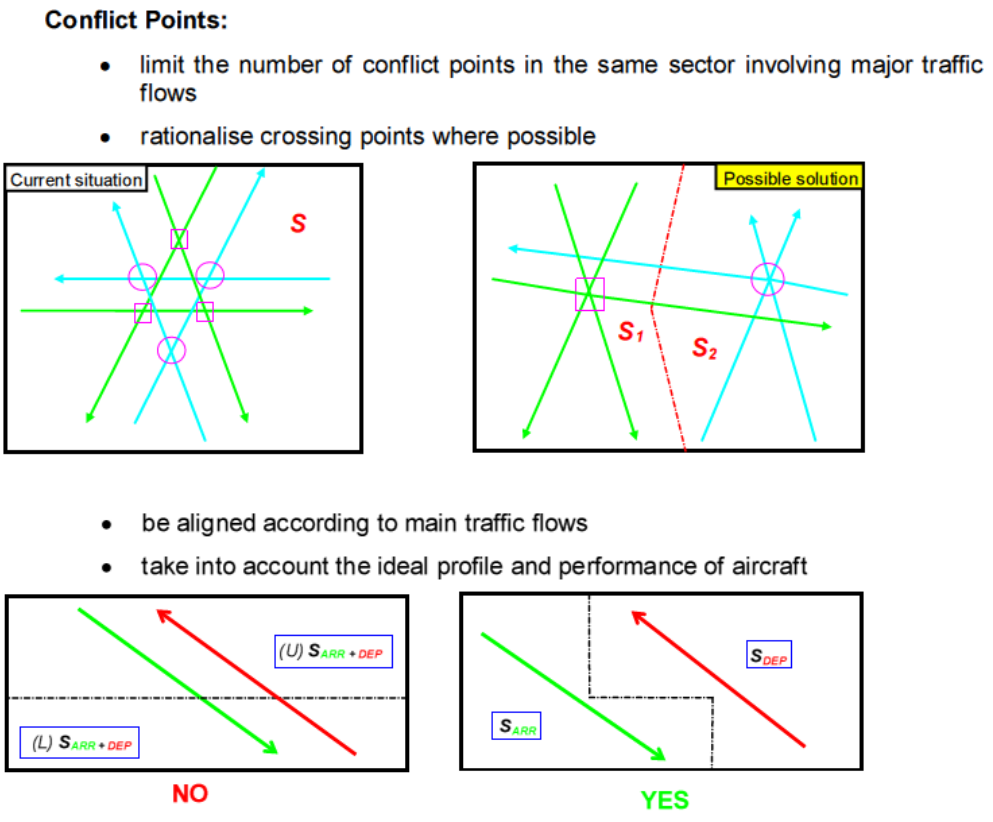

FIGURE 2.1: Airspace design guidelines. Source, Eurocontrol [19]

Further, to maintain sector throughput, controllers, especially under high levels of traffic demand, usually focus on safely clearing a flight through a sector as quickly as possible to free up airspace capacity for the next inbound flight. Therefore in satisfying traffic demand there is a certain acceptance of flight inefficiency due to an emphasis on delivering time efficient flights rather than specifically fuel or environmentally efficient flights. This necessary focus on expediency to provide capacity can then limit the flight or environmental efficiency of trajectories.

Environmental Impact

Whereas optimising the airspace design and operation for safety and capacity can involve the strategic and tactical use of spatial and temporal constraints, optimising for reduced fuel and CO2 operations frequently involves the removal of constraints from the system. Very generally, to reduce CO2, the aim is to allow the aircraft to fly the most direct path possible at the most fuel efficient cruising levels [18, 31, 32].

The closest realisation of this principle in the current system is free routing airspace [18, 31, 32]. Within free routing airspace, aircraft are free to choose direct routings between defined entry and exit points without the need to follow intermediate waypoints [18, 31, 32]. Separation assurance is still provided by air traffic control. However, currently, due to issues with controller workload, free routing airspace is generally only available at higher flight levels and when air traffic demands is low, such as night time and weekends [18, 31, 32].

Environmental impact within the Terminal Manoeuvring Area (TMA) however is more complicated than direct routing. Within the TMA it is recognised that there are significant trade-offs between aviation related environmental impacts [17, 33]. When approaching and departing the terminal area, aircraft fly along Standard Instrument Departure routes (SIDs), Standards Arrival Routes (STARs) and Instrument Approach Procedures (IAPs). SIDs, STARS and IAPS are pre-planned IFR arrival and departure procedures used to simplify clearance delivery, expedite traffic flow, and reduce pilot/controller workload [27]. Of increasing importance is the environmental design of the routes and how the horizontal placement of the route and the aircraft vertical profile along the route can be optimised for environmental impact.

A significant enabler of improvements in arrival and departure route design is Performance Based Navigation (PBN) [34, 35, 36]. PBN allows routing and vertical profile constraints to be more accurately followed and therefore can allow for the planned environmental benefits of the procedures to be more fully realised [34, 35, 36].

The approach and departure routes can be designed to enable aircraft noise abatement operating procedures such as NADP 1 and 2[i] for departures and Continuous Descent

[i]NADP1 and NADP2 are noise abatement departure aircraft operating procedures.

Approaches (CDA)[ii] on arrivals, that have been shown to mitigate noise and/or emissions impact [37, 23, 38]. There also has been increased focus on facilitating fuel efficient climbs through the removal of airspace restrictions that may impede optimal climb trajectories [39, 40]. However the removal of restrictions for environmental reasons may have airspace safety and capacity implications, or even environmental implications for other inbound and outbound traffic flows.

Therefore when considering the planning of environmentally efficient arrival and departure operations, the most noise optimal route may not be the most fuel or operator optimal route and there may also be inter-noise trade-offs between mitigating noise at specific sites and mitigating overall noise impact. Therefore there is a need to manage all these trade-offs in an informed manner [33, 23, 38].

The sequencing and merging of arrival traffic approaching an airport in high traffic conditions is a problem that has a relationship to arrival procedure design but is distinctive enough to warrant separate discussion. Sequencing and merging is a problem where careful airspace design and operation can be used to mitigate environmental impact. At airports such as Heathrow, where runway capacity is limited, airborne holding is used to create reservoirs of traffic used to maintain the runway rate. In these cases, it is considered good practice to design holding stacks so that aircraft can begin their descent to the runway from as high as possible [41]. This facilitates CDA delivery from as high as possible, minimising noise impact and fuel burn. Below the stacks, when merging traffic onto final approach, air traffic controllers issue open loop instructions to aircraft. Although the use of tactical vectors gives ATC a great degree of flexibility in merging the traffic flows, the flight crew cannot accurately predict distance to touchdown and therefore cannot determine optimum fuel efficient descent rates [27, 41]. The point merge concept has been proposed to merge the traffic in a more environmentally efficient manner.

For point merge, a family of approach procedures of differing lengths are developed, where the aircraft are allocated routes depending on the sequence desired and the delay that each flight needs to absorb. As all the route lengths are known to the flight crew, the distance to touchdown can be more accurately predicted by the crew potentially resulting in more environmentally efficient descent planning [43, 44]. The point merge concept is currently in trial at a number of airports and has shown promising benefits [45]. Whether these benefits continue to be realised with larger traffic densities and for all types of airports remains an area of research [45]. Trajectory optimisation work in this thesis did not examine the point merge concept explicitly, but the methods could be applied, investigating the concept as an extension of the standard arrival procedure development problem.

[ii] CDAs are continuous descents approach aircraft operating procedures with minimum thrust usage.

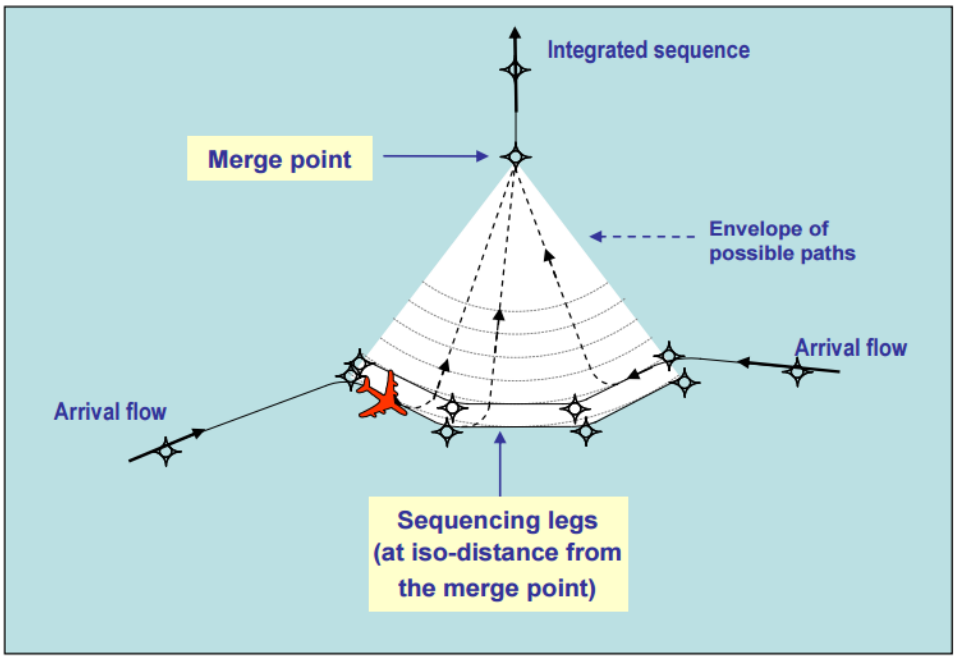

FIGURE 2.2: Point merge concept. Source, Eurocontrol [42]

2.2.1 Summary: Conflicting Demands

In summary then, when traffic demand on the ATM system is low, air traffic control will typically clear aircraft to fly direct routings at their optimal cruising levels. With increased traffic however there is an increased demand on limited resources such as airports and airspace that must be managed in a way that safely and efficiently satisfies that demand. There is also a related and accumulative increase in environmental impact in the form of noise and aircraft emissions. In response to this, there is an increased reliance on systemisation, where airspace is designed and operated in a more constrained manner to balance the goals of safety, capacity and environmental impact. Therefore, the systemisation must be carefully planned and needs to be supported by methods and tools that support decision making.

It is also discussed in this section how the pull of satisfying these sometimes conflicting goals can introduce operational constraints to trajectories that limit the environmental efficiency of any given flight. How the system constraints typically affect a flown flight is discussed in full detail in Section 2.3. In this current section it is further highlighted, that within the TMA, the environmental efficiency of a flight means different things to different people, and that neither the most community noise optimal operation of a flight nor most fuel optimal operation may result in a flight trajectory that best balances overall environmental impact.

After exploring how system design and constraints typically impact the efficiency of flown flights in Section 2.3, Section 2.4 takes a look at what concepts, tools and practices are being applied in the near term towards improving flight efficiency and supporting

continued air traffic growth. Section 2.5 then considers the more long term SES trajectory based concept of operation. Section 2.6 then, in the context of the thesis goals, will consider how trajectory optimisation methods can be used in support of both near and long term ATM system improvement plans.

2.3 Flight Efficiency - The Operator Perspective

This section will consider how airspace users are typically affected by the constraints that arise from the design and operation of the airspace around the goals of safety, capacity and environmental impact. This section also considers how aircraft operators factor ATM constraints into their flight planning and how those constraints typically impact the operation of aircraft within European airspace on a busy day.

Pre-Flight

For each commercial IFR flight within controlled airspace, a flight plan must be prepared by the aircraft operator for submission to air traffic control. The flight plan defines the intention of flight to air traffic control, with information regarding time and date of operation, intended route, requested flight levels and cruising speeds [27, 46]. The routes that aircraft are required to fly, the flight level and speed restrictions along the routes are a result of the airspace design and management goals discussed in Section 2.2. Details of each country's route network constraints are contained in the Aeronautical Information Publication (AIP) of each nation. The RAD is intended as a master flight planning document that consolidates en-route AIP information with traffic flow restrictions and ATFM routing requirements designed to make the most effective use of Air Traffic Management (ATM) capacity [47, 27, 46]. The RAD is a result of strategic balancing of demand and capacity. Operationally significant updates to both AIPs and the RAD occur every 28 days in what is known as the AIRAC (Aeronautical Information Regulation And Control) cycle. The intended route submitted to air traffic control as part of the flight plan is then listed as a sequence of arrival and departure routes, waypoints and airways [48, 27]. The requested cruising levels and speeds along the airways are all planned to maximise airline performance while also satisfying restrictions published in the AIP/RAD. Operators then submit flight plans to ATM units and European flow control through the Flight Plan Processing System (IFPS).

On the day of the flight, a number of hours before the flight's estimated departure time, the impact the flight will have on downstream airport and airspace capacity is assessed by the Network Manager Unit [49]. The Network Manager[iii] provides European air traffic flow control, and is concerned with balancing the traffic demand with the capacity of the

[iii]The Network Manager provides pan european ATFM and has a mission to strategically, operationally and technically support european air navigation service provision.

system at a pan-European level. Network Manager software will use the flight plan and aircraft performance data to create a 4D trajectory projection of the flight to check that the flight does not pass through airspace or arrive at an airport with a regulation applied to it [49, 48, 27]. If airspace or airport capacity is predicted to be exceeded then it will be subject to a regulation that limits the number of aircraft that can use that airspace or airport for a period of time. If an aircraft is predicted to fly through airspace or to an airport with a regulation applied then a constraint is applied to the departing aircraft designed to prevent capacity from being exceeded [50, 30].

The constraint can take the form of a hold at the departure airport, a flight level restriction (flight level capping) or the Network Manager may propose an alternative routing for the aircraft. The Network Manager then issues a Calculated Time of Takeoff (CTOT) slot to the aircraft operator and ground control at the departure airport, who collaborate to deliver the slot. The CTOT is a 15 minute window that airport ground control aim to have the aircraft depart within. The aircraft's flight plan is also sent to ATM units along the flight route to advise them of the aircraft's intent and impending presence in their airspace [49].

In-Flight

ATC Clearances are permissions to proceed under specified conditions and/or to a specified point, waypoint fix, or airspace boundary [27, 25]. Once the aircraft is cleared for take-off, each step of aircraft's intended flight as expressed in the flight plan is approved through tactical ATC clearances.

For a typical flight, the aircraft will be cleared to fly along a SID, which as discussed in Section 2.2, is a path with procedural constraints designed with safety, environmental and traffic separation considerations. Where the constraints permit, operators can choose to fly noise abatement operating procedures along the SID. From the end of the SID, the aircraft will be cleared in steps to its requested flight level. Aircraft may not be cleared to a higher flight level if that flight level is already occupied or if that clearance would result in a loss of separation between aircraft. It may also be constrained to lower flight levels due to procedural constraints such as standing agreements which may not be published in the AIPs.

Once in cruise, the flight will continue along the designated airways at the RFLs unless otherwise instructed by ATC. ATC may tactically alter the flights trajectory to ensure safe separation, or, for instance, use a direct routing to expedite the flight through the airspace. Descent from higher to lower flight levels and onto STARs frequently involves tactical judgement by air traffic control to best descend aircraft through airspace shared with other climbing, descending and cruising traffic flows.

Although there is an increasing awareness of the impact of tactical instructions on the ideal flight profile, the flexibility of tactical instructions allows ATC to minimise flight

time and maximise capacity to satisfy the throughput demands of the system. Descent onto the STAR and sometimes along the STAR will be approved through ATC clearances. At the end of the STAR, it may be possible for the aircraft to transition directly to an approach procedure. Where there are high levels of runway utilisation, aircraft may be subject to airborne holding. Descent from the hold or from the end of the STAR is frequently tactically controlled with radar vectors to allow the sequencing and merging of the aircraft with other inbound traffic. Where there is a policy of delivering CDAs the aircraft will be issued an estimated distance to go as they begin their descent to the airport to help the flight crew support their descent flight planning [41].

2.3.1 Summary: Highly Constrained

Therefore from an operator's perspective, air traffic control imposes a wide number of constraints that limit the optimal routing, height, speed, aerodynamic and control settings. Aircraft usually still have some scope within the ATM constraints to mitigate their emissions impact through the choice of climb and descent rates and cruising speeds. The aircraft speed schedule and therefore the climb and descent rates are largely determined by the cost index, which defines a priority ratio between the fuel burn cost and the operating time cost of the aircraft [51, 52]. In general faster flight times reduce operating time costs but increase fuel consumption costs. The FMS then adjusts the aircraft speed schedule according to the cost index entered by the aircraft operator. Operators are incentivised then by the cost of fuel and the cost of CO2 permits to reduce CO2 emissions over others forms of environmental impact, but also need to strike a balance between mitigating CO2 and profitably operating the aircraft.

For mitigating noise impact at low flight levels, aircraft routings are again highly constrained by the noise preferential routing of the SIDs, STARs and IAPs, making it critical that these routing constraints are environmentally optimised. Although there may be height, speed and climb rate constraints along the routings, they may not have been environmentally optimised, and are often intended more for safety and traffic separation. Aircraft may choose to fly noise abatement aircraft operating procedures, such as NADP1, NADP2 and CDAs, along the routings. However, because the aircraft operating steps have not been harmonised with the routings at the design stage, even with all actors acting in the same direction, the full environmental benefits of the combined ATM and aircraft operating procedures may not be realised.

Sections 2.4 and 2.5 consider what impact near and long term evolutions in the air traffic management system will have on environmental flight efficiency. Section 2.6 then, in the context of the thesis goals, considers how trajectory optimisation methods can be used in support of those system improvement plans.

2.4 Near Term Solutions

The European air traffic management system has already begun a transition to the Single European Sky (SES) and Trajectory Based Operations (TBO) [2]. However, before considering the impact of TBO on the environmental efficiency of a flight, a number of more immediate technologies and initiatives are considered. How these more immediate changes will remove, reduce the impact of, or alter the constraints that limit environmental flight efficiency will now be discussed.

Functional Airspace Blocks (FAB)

Airspace in Europe is highly fragmented with air traffic services provided by a range of states and agencies [53]. Functional Airspace Blocks aim to harmonise the air traffic service provided to aircraft operators. The intention of a FAB is that ANSPs cooperate to plan airspace not with regard to national boundaries but with regard to delivering safe and efficient traffic flows through multiple European regions.

It is proposed that by standardising procedures, training and systems within a FAB a more seamless service can be offered to the airspace user [54]. It is anticipated that by designing airspace around major traffic flows that large improvements in flight efficiency may be possible due to the greater availability of direct routing at efficient cruise levels [54]. FABs however have had difficultly in demonstrating concrete operational gains [55]. There have also been difficulties in determining how roles and responsibilities are shared among the ANSPs, nations and the Network Manager charged with operating the FAB [55].

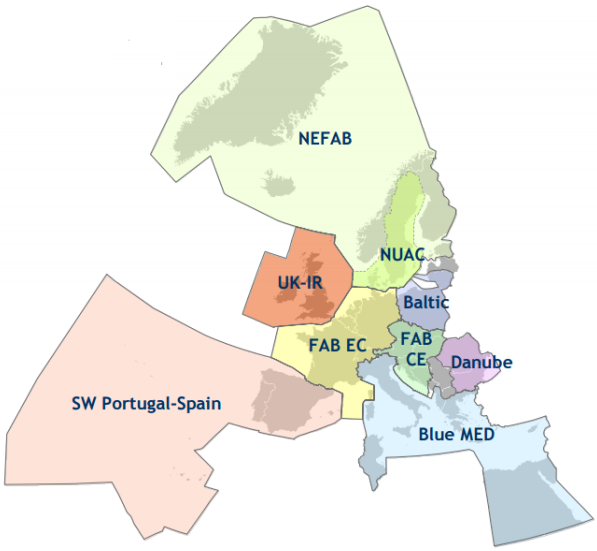

FIGURE 2.3: Planned Functional Airspace Blocks (FABs). Source, Eurocontrol [19]

As of 2014, there are two implemented FABs, namely the UK-Ireland and Denmark-Sweden FABs, with plans for nine in total covering European airspace [56]. Therefore, it

is proposed that FABs be a key feature of European ATM moving forward, and be a mechanism that drives greater airspace integration and harmonisation. However, in practice, it may be possible that greater harmonisation be achievable through other means.

Flexible Use of Airspace (FUA)

The concept of Flexible Use of Airspace (FUA) is that airspace is neither designated as either civil or military but can be flexibly scheduled by either user based around strategic and operational needs [19, 57]. Therefore, when not being used for military operations, airspace can be freed to be used by civil flights, enabling an increase in airspace capacity, routing and flight levels options.

Free Routing

Free routing is a concept where, in airspace with low to medium levels of traffic, aircraft are not constrained to the ATM route network, and are therefore free to route directly between specified airspace entry and exit points [18, 31, 32]. In free routing airspace, separation continues to be provided by air traffic control according to rules on minimum separation. Free routing airspace currently exists over much of central Europe above flight levels FL330 and is utilised mostly at night when controller workload is low [18, 31, 32].

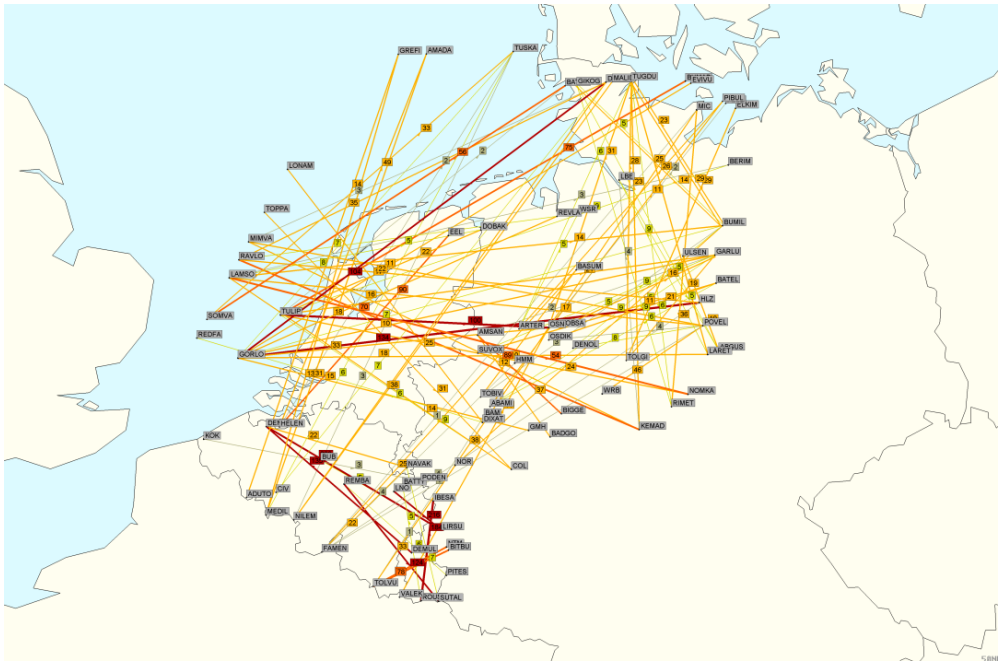

FIGURE 2.4: Free Route Airspace Maastricht (FRAM). Source, Eurocontrol [32]

Mid Term Conflict Detection (MTCD)

In free route airspace, aircraft are not required to operate along traffic corridors and potential conflict points between traffic are more difficult to predict [58, 59]. Therefore several agencies have developed Mid Term Conflict Detection tools to support air traffic controllers in identifying and resolving potential losses of separation between aircraft [58]. It is hoped that continued developments in conflict detection and resolution tools can be used to improve traffic predictability in a way that enables free routing in airspace with higher levels of traffic demand [58, 59].

Arrival Management (AMAN)

It is anticipated that Arrival Management (AMAN) technologies will offer greater supervision of inbound traffic allowing for the calculation of optimum aircraft arrival sequences and times that will subsequently allow for optimised use of runway capacity, more efficient and reduced levels of arrival stacking, smoothed traffic flow into stacks, and the absorption of arrival delay prior to stacks [60, 2, 61]. Once aircraft are within the AMAN planning horizon, the tool can calculate the ideal sequence and associated waypoint crossing times. Near term, air traffic controllers cooperating across sectors can issue advisories to speed up or slow down the aircraft in order to deliver the aircraft to the designated metering waypoints at those times[2, 61]. For more long term developments, the tools may issue constraints directly to the aircraft. AMAN tools in particular are being trialled where the AMAN issues CTA (Controlled Time of Arrival) constraints directly to the aircraft FMS, with the flight crew then adjusting the aircraft's speed schedule to meet the constraint. It is proposed then that the use of AMAN crossing times to schedule aircraft will reduce the need for vectoring on final approach facilitating better delivery of CDAs [2, 61, 54].

Airport Collaborative Decision Making (A-CDM)

The concept of integrating the best sources of information as well as sharing the objectives of each actor in the system is expected to enhance efficiency and safety by allowing the most informed choices of action by all parties. This is referred to as Collaborative Decision Making (CDM) [2, 59, 54]. For Airport CDM information is shared between airport operators, aircraft operators, ground handlers, air traffic control and the Network Manager. This information is used predict and schedule demand on airport resources and also used for accurate estimation of turnaround, pushback, and runway slot demand times. CDM is considered to be an necessary enabler for departure management (DMAN) tools [2, 60].

Departure Management (DMAN)

DMAN tools will provide decision support for airport ground control [2, 60]. Its goal will be to provide optimum runway allocation, sequencing, and ground routing advice, which will allow for the lessening of inefficient engine use related to queuing and ground holding. The use of Departure Management (DMAN) interacts closely with Airport-CDM. DMAN requires accurate and timely information to work effectively to predict runway demand, pushback and taxi times. The scheduling of airspace resources further downstream of the departure airport are usually dependant on flights meeting predicted take-off times. Where slots are not met then flight efficiency may be negatively affected. It is anticipated that ground controllers will increasingly rely on DMAN tools to support the delivery of aircraft to runways to satisfy departure slot times [2, 60, 54].

Communication, Navigation and Surveillance (CNS)

Developments in navigation technologies and standards will be dominated by the transition from a predominantly ground-based to a satellite-based infrastructure [2, 59, 60]. Improvements in navigation technologies are intended to permit aircraft operation on any desired flight path using ground or satellite based navigational aids, or a combination of both. This will support 3D and 4D track keeping and the realisation of free routing [2, 59, 54].

In busy airspace where route constraints will be required, then the increased precision of the navigational technologies should allow aircraft to more accurately maintain their routings and therefore allow for controllers to issue fewer ATC vectors to maintain aircraft on given routes [62, 60, 61].

Air-ground and ground-ground communication technologies are important enablers for air traffic management and airline operational control. They enable voice and data communications between controllers and the flight deck, between computers and between multiple airports.

Moving forward, cooperative surveillance is expected to become the system standard [2, 59, 54]. ADS-B (Automatic dependent surveillance-broadcast) and multilateration are expected to augment radar coverage and also to provide a replacement for radar coverage where it is poor or non-existent, filling in gaps particularly at low levels [2, 59, 54]. Extra information down-linked from aircraft can be utilised in controller support tools to provide applications, such as improved trajectory prediction information and precision approach monitoring.

2.4.1 Summary: Solution Evolution

Highlighted here are a number of near term concepts and technologies that are likely to be enablers of more environmentally efficient flights. There are two key themes for improving aircraft flight efficiency emerging from this overview. These are, the desire to enable more direct flight routings at fuel efficient flight levels, and the improved scheduling of airport and airspace resources in order better realise those efficient flights at the times they are desired. These two themes have also been been major drivers in the development of the SES target concept. A full discussion of how the technologies and concepts summarised here will continue to evolve as part of the SES target concept can be found in [2, 59, 54].

2.5 Trajectory Based Operations (TBO)

The most defining element of the SES target concept is the transition to trajectory based operations supported by System Wide Information Management (SWIM).

A central tenet of TBO is that large flight efficiency gains can be made by replacing currently used flight plans with highly accurate flight intent information in the form of user generated 4D flight trajectories. The principle behind this is that by giving ANSPs highly accurate information on where the aircraft is and where it wants to be, they can plan for the flight over longer time horizons, only applying constraints where necessary for the safe and efficient operation of the flight. It is intended that a common 4D reference trajectory be maintained between ANSPs and airspace users at all times, and that cooperative revisions to the trajectory be achievable in real time. Therefore, where constraints must be applied, they can be precisely applied, and done with the cooperation of the airspace user so as to minimise disruption to the intended flight.

The majority of the following description of TBO and the SES operating concept is taken, with some authorial interpretation, from the SES target concept document [59], with supporting information from the SES master plan [2], the SES concept of operations [63] and SES TBO analysis provided in [64].

Within SES, the business trajectory is the term used to describe the highly accurate 4D reference trajectory shared between ANSPs and the airspace users. A key goal of SESAR trajectory orientated operations is that airspace users business needs are accommodated to the greatest extent that is possible. Therefore, the business trajectory is a 4D flight trajectory that is considered to express the desired business outcome of the airspace user, and is also considered to be the target trajectory for ANSPs to deliver.

The business trajectory will have three distinct life cycle phases. The phases are;

- Business Development Trajectory (BDT),
- Shared Business Trajectory (SBT) ,
- Reference Business Trajectory (RBT).

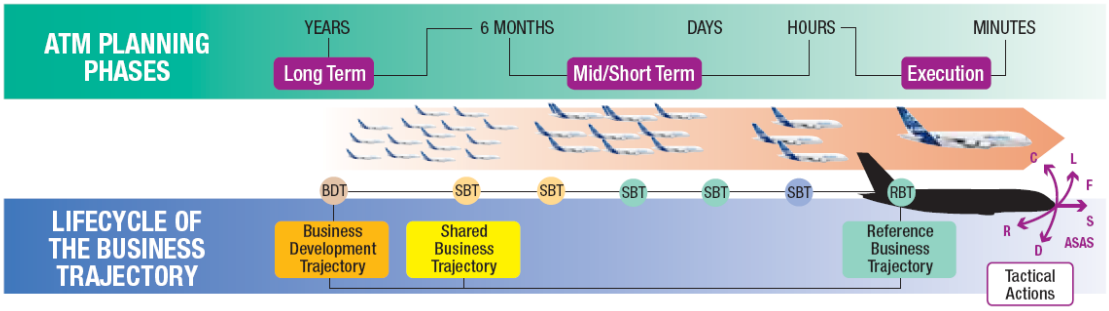

FIGURE 2.5: SES Business trajectory life cycle. Source SESAR [42]

The Business Development Trajectory is a planned trajectory that is not shared outside the airspace users organisation. The BDT results from the airspace users own internal business planning, resource management and schedule development. The development of the BDT can be viewed as corresponding to the preparation of the flight plan in current day operations. As with flight plan preparation (see Section 2.3), the RBT is planned to be commensurate with the constraints found in the various AIPs. For SESAR trajectory based operations, it is anticipated that controllers will be equipped with conflict detection and resolution tools (such as evolutions of the MTCD, see Section 2.4), which, supported by Performance Based Navigation, will reduce controller reliance on route structures to provide traffic predictability and procedural separation of traffic. Therefore, it is anticipated that, in airspace with low to medium traffic densities, aircraft will be less constrained to fly along an ATM route network. This should reduce the number of AIP ATM constraints that must be factored into the planning of the flight.

However, in high density airspace such as terminal areas, there is likely to be a continued need to plan trajectories according to route constraints (SID, STAR, IAP charts) and with consideration of environmental abatement procedures (noise and emissions) in order to support the safety, capacity and environmental management of the airspace [59, 62].

The Shared Business Trajectory is a planned trajectory that is sufficiently developed to be shared with other airspace users. In a process that is analogous to submitting a flight plan to the Network Manager and ATM units, the Shared Business Trajectory will need to be shared with ATFM, ANSPs and airports. Once shared, ATFM and ANSPs will plan their traffic flows in order to provide the necessary capacity and services to support the flight, potentially reorganising airspace to do so. However, where resource demand exceeds capacity, there may be still be positional and/or temporal ATM constraints applied to the planned trajectory.

The airspace user has responsibility for revising the SBT in order to best meet the ATM constraints. Through an iterative process between ANSPs and users, a SBT is reached that provides the optimum balance between user needs and ATM flow management. The final Shared Business Trajectory then becomes the Reference Business Trajectory. The RBT is defined as the trajectory the airspace user agrees to fly and the ANSP agrees to facilitate. The RBT is now a highly accurate 4D description of the aircraft's intended trajectory, inclusive of ATM constraints. The RBT is also the trajectory that the aircraft's flight management system will track to.

The RBT is not a de-conflicted trajectory. En-route separation assurance is still provided by human air traffic controllers. Therefore, the RBT may not be the actual trajectory flown but is a common reference trajectory that is shared by ANSPs and the operator, which is cleared in stages by ATC. This is similar to how a flight plan is cleared in stages in current day operations (see Section 2.3). The difference is, that by having very accurately defined aircraft intents complemented by MTCD controller support tools (see Section 2.4), that the trajectory can be cleared to be conflict free for much longer durations of flight than is currently possible. This, it is proposed, will considerably reduce controller workload, increasing airspace capacity and reducing the number of capacity related constraints imposed on the trajectory.

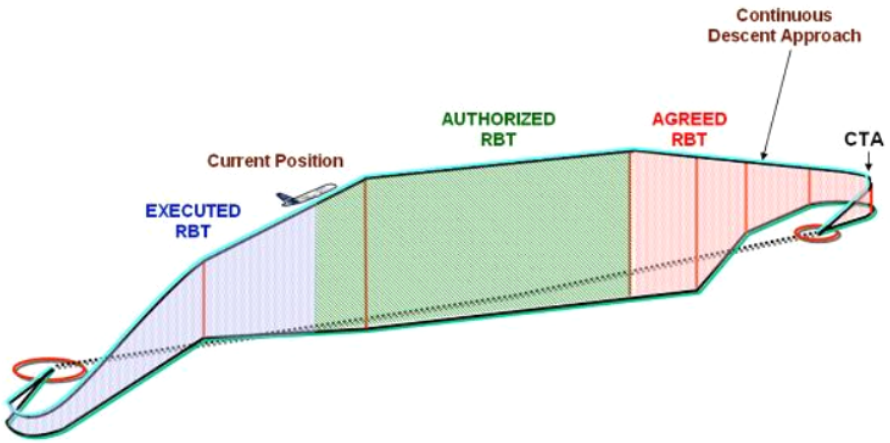

FIGURE 2.6: SES RBT, cleared in stages. Source SESAR [59]

The RBT may be subject to revision while the aircraft is in flight to ensure safe separation, or to reflect updates in 4D metering constraints. Where controllers need to take action to ensure safe separation, ATC can issue a 2D, 3D or 4D routing constraint to the aircraft, which is then used by the on-board flight management system to advise the flight crew in revising the common RBT and to support them in flying the aircraft to meet the constraint. In these cases the flight crew would control the free dimensions to avoid any conflict. The function within the FMS to support self separation is referred to as Airborne Separation Assurance System (ASAS). In the SES concept of operation it is also proposed that, in certain circumstances, ATC be able to delegate the role of separation provision to the aircraft, where ASAS tools would be used as the sole source of separation provision.

In addition to constraints for separation provision, 4D metering constraints may be updated due to arrival and departure queue management tools (AMAN and DMAN, see Section 2.4) that can use the most up to date traffic information to issue constraints aimed at improving the traffic situation at and around airports.

As the RBT is commonly shared with ANSPs and airports, changes to the RBT are automatically visible to all actors in the system. This allows the ANSPs to judge the impact of the change in the RBT in terms of conflicts, delay and controller workload and to adjust the traffic flows in order to minimise the negative impact of any RBT changes. It also allows airports to update estimated landing and in-blocks time and to allocate resources accordingly.

System Wide Information Management (SWIM)

SWIM is a concept where ANSPs, aircraft operators, flow management centres, airports and other stakeholders are connected within a single information sharing infrastructure.

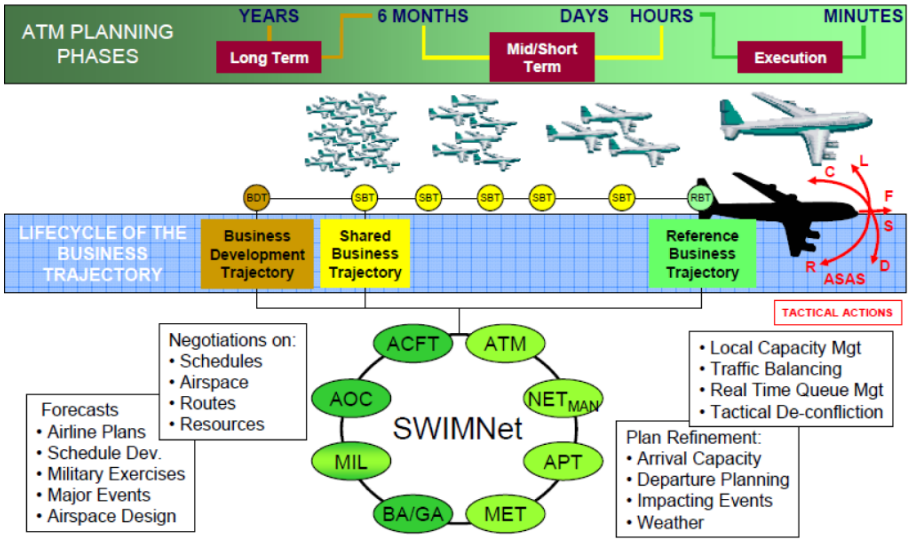

FIGURE 2.7: SWIM supporting trajectory based operations. Source, SESAR [65]

SWIM is intended as an enabler of trajectory based operations by supporting the sharing and revision of highly detailed information on intended or active business trajectories.

In the current system stakeholders involved in the planning, management, and execution of flights do not share information among each other. This can lead to decision making by individual stakeholders based on best guesses of the intents of other stakeholders. This in turn can lead to inefficient aircraft trajectory operations. For the SWIM concept, stakeholders such as ANSPs, operators, flow management, airports and the military share the most up to date information on intent and status in real time to allow the decisions to be made on the best information available.

It is intended that SWIM will enable Collaborative Decision Making (CDM). CDM takes the Airport-Collaborative Decision Making concept and applies it to the entire ATM stakeholder system. For CDM, stakeholders work from a harmonised data-set of information, collaborating in real time to balance the demand and capacity of the system. It is expected that this approach will lead to improvements in both ATM system and flight efficiency.

2.5.1 Single European Sky: Target Concept

Summarising, currently European ANSPs and the Network Manager use flight plans updated with radar data to predict the 4D trajectory of a flight. This is a highly inaccurate form of prediction. Therefore, predicting the demand on resources such as airspace, flight levels and runways can only be achieved with a large margin of error. This necessitates that the system is be managed in a tactical manner in a first come first served approach.

By having highly accurate 4D trajectory information shared between ANSPs and aircraft, the situational awareness of both air traffic control and the flight crews are increased. Therefore demand on airspace resources can be accurately predicted over longer time horizons. This enables the system to be managed in a more strategic manner where, in combination with enhanced conflict detection and resolution, flights can be cleared on direct routings for long portions of time. It is anticipated that this will reduce controller workload, which will lead to increased system capacity and safety. This will in turn lead to more availability of direct routings, reducing fuel costs and CO2 emissions. It is also anticipated that the increased predictability of 4D trajectories will enable runway resources to be better scheduled so that delay costs in the form of time and fuel burn from inefficient flight operations like airborne holding can be minimised.

2.6 Trajectory Optimisation

In the Introduction, in Section 1.2, it is highlighted that the work in this thesis is orientated around two specific goals. Recapping from Chapter 1, the goals are to support

- airspace designers in environmentally optimising the ATM system constraints,
- airline operators in planning environmentally efficient flight trajectories within the ATM system constraints.

Therefore, in the following sections, it will be considered how control based trajectory optimisation methods can be utilised in the context of these goals. Sections 2.6.1 and 2.6.2 review the current state of the art with regard to trajectory methods and their application to procedure optimisation and trajectory based flight planning. How work in this thesis contributes to each of the goals is set out in Section 2.6.3 on Climb and Descent Procedure Optimisation and in Section 2.6.4 on Flight Trajectory Planning. The conclusions of this thesis in section 10.1 and 10.2 are then similarly considered relative to both of these goals.

2.6.1 Procedure Optimisation

While approaching and departing the terminal area, aircraft are constrained to fly along procedural routes developed by Air Navigation Service Providers (ANSPs). The routes termed Standard Instrument Departures (SIDs), Standard Arrival Routes (STARs) and Instrument Approach Procedure (IAPS) serve a number roles. They provide obstacle avoidance and navigational support to aircraft departing and arriving to airports. They also support the safe management of traffic by creating corridors of traffic that can be easily managed by air traffic control. Of increasing importance is the environmental design of the routes and how the horizontal placement of the route and the aircraft vertical profile along the route can be optimised for environmental impact. As identified in [59] and [62], arrival and departure procedures will continue to be a feature of the SES operations, therefore the environmentally efficient planning of these procedures will continue to be a concern for the forseeable future.

The environmental optimisation of preferential routings and aircraft operating steps along those routings fall under the broad category of environmental procedure optimisation. If we consider how preferential routings and recommended flight operating steps constrain the horizontal paths and vertical profiles of trajectories. Then we can see that the goal of environmental procedure optimisation is to optimise the system constraints so that they best support the delivery of the most environmentally efficient climb and descent trajectory operations.

The Continuous Descent Approach (CDA) is an operating procedure for aircraft descents. The CDA ideally involves the aircraft descending from cruise to the runway while eliminating or minimising level approach segments, minimising thrust settings and maintaining a low drag configuration for as much of the descent as possible.

In recent years the CDA trajectory has been studied in detail [23]. CDAs can be implemented with current day technology levels and hence require relatively low investment costs. CDA procedures have been implemented at a number of major airports as they offer reductions in community noise impact as well as reductions in fuel consumption leading to lower emissions and lower airline operating costs [66].

The ICAO Procedures for Air Navigation Services - Aircraft Operations (PAN-OPS) document [37] contains recommended departure flight procedures that are complementary to CDAs and aim to reduce aircraft noise impact on departure. The PAN-OPS procedures, termed Noise Abatement DeParture climbs (NADP) 1 and 2 are aircraft operating procedures that can be applied to any departure routing. NADP 1 is designed to minimise noise near to the airport and NDAP 2 is designed to minimise noise further out from the airport.

Further to the PAN-OPS guidance, the Sourdine project [67] used expert analysis to develop noise optimised take-off procedures for a representative medium narrow-body and large wide-body aircraft. The PAN-OPS and Sourdine recommended flight procedures offer general guidance on minimising departure noise impact. However, neither set of recommended procedures are optimised to local conditions, and may require significant modification to fully realise potential noise benefits.

A number of ICAO CAEP reviews have looked at the factors that limit the realisation of environmentally efficient climb and descent procedures [23, 33]. Work in [33] showed that in addition to needing to adapt procedures to local conditions, there are multiple trade-offs to be made among various noise objectives and also between noise and emissions objectives. In the summary conclusions of the work, the authors observed,

"The results indicate that, of the procedures included in this study, no single departure procedure minimizes overall noise and emissions simultaneously."

In [23], further factors that limit the realisation of the best environmental climb and descent operations are reviewed. Highlighted in particular was a lack of harmonised planning between operators and ANSPs. Historically, ANSPs and operators have planned environmentally efficient climbs and descents separately, with ANSPs planning the horizontal routing and operators defining the aircraft operating steps along the routing [23]. This fragmented approach has prevented the full realisation of potential environmental benefits from the resulting combined operating procedure [23]. Specifically the authors in [23] make the following observations

"Lack of harmonising guidance - As noted above, PANS-OPS establishes minimum altitudes for aircraft configuration change and thrust reduction within an NADP, but leaves

development of specific aircraft profiles to the operator. Appendix B illustrates the diversity of aircraft procedures. These variations make the quantification of noise and emissions benefits very difficult and drive the requirement for very sophisticated modelling to determine the effects of the different profiles. With respect to CDA, there is no single definition of the procedure, nor is there a commonly agreed methodology or toolset for assessing the benefits."

Therefore, summarising from the ICAO review papers, there is a need for a sophisticated simulation methodology that can provide a harmonised approach to procedure planning, simultaneously optimising both routing and aircraft operating steps in determining the most environmentally efficient climb or descent operating procedures. The methodology will need to support detailed trade-off analysis between several environmental objectives, including multiple noise and emissions objectives. For noise objectives, in order to tailor procedures to local conditions, the methodology must be capable of supporting sophisticated analysis of noise impacts on community populations local to airports. The methodology must be able to do all this and also for a range of different aircraft types.

Although there is no methodology observed in the literature that meets all the requirements set out in the ICAO reviews, some steps towards these goals have been taken. These steps are now discussed.

Looking to mitigate aircraft noise impact specific to local populations, Visser [68, 69], treating the problem as an optimal control problem, used a direct collocation method [70], noise mapping software and a Geographic Information System (GIS) database to define 3D noise abatement procedures that minimise sleep disturbance related to aircraft noise at communities surrounding Amsterdam Airport Schiphol. Hebly [71] extended this approach by including an aircraft emissions model and comparing the results of optimising an RNAV (Area Navigation) SID for emissions relative to optimising for awakenings from aircraft noise. The NOISHHH tool used in this work however was a single objective optimisation tool, and while an analyst may continually run scenarios, looking at different objectives or different objective weightings, this approach is time consuming, requires specialist knowledge of weighting and is not suitable for providing a comprehensive trade-off analysis. Also, due to computational issues with the multi-phase formulation of the problem, researchers needed to specify parts of the horizontal path, vertical trajectory or both prior to the optimisation. The need to pre-specify part of the answer in this manner, limiting the trajectory space prior to optimization, reduces the tools effectiveness in supporting a harmonised approach to procedure planning.

Similar to the NOISHHH work, Houacine [72] uses single objective optimisation to look at reducing noise and fuel consumption from aircraft operations. In this work a Gauss pseudospectral collocation method was used to compare the differences between trajectories optimised for noise levels under the aircraft centreline and for fuel consumption. The use of global polynomials and the SNOPT solver only required specification of the

boundary values and constraints prior to optimisation. However, the scenarios considered were idealised 2-D scenarios offering no consideration of realistic community noise impact. The suitability of utilising pseudospectral methods for multi-objective procedure design optimisation for realistic local conditions remains unexplored.

Prats [73] recognising that environmental impacts are multi-objective in nature, treats the procedure planning problem as an multi-objective optimal control problem. The work aims to provide a harmonised approach for departure only procedure planning, principally considering trade-offs between noise objectives. The problem was then solved with a direct collocation method and a lexicographic based optimisation technique. The lexicographic method required the definition of a hierarchy of objectives. The approach then uses a series of single objective optimisations, first finding the minimum of the objective highest in the hierarchy, before moving on to minimise objectives lower in the hierarchy, adding constraints at each optimisation step to preserve the higher ranked objectives. The lexicographic method produces a single Pareto optimal point and requires the definition of the importance of each objective prior to the optimisation.

In many cases, the decision maker finds it desirable to be informed about the full range of trade-off options in order to choose the solutions that provide the best balance between the objectives. In this case, the Pareto optimal set is desirable. The solution vector $\mathbf{s}^*$ of a multi-objective optimisation problem is Pareto optimal, if there is no other solution vector $\mathbf{s} \in S$ where $s_i \leq s_i^*$ for all objectives $i = \{1, \ldots, k\}$ and where $s_j^* < s_j$ for at least one index j, $j \in \{1, \ldots, k\}$. A Pareto front is a set of Pareto optimal solutions. The Pareto optimal set identifies the extrema of the objectives and allows trade-offs between the objectives to be quantified and visualised. This is especially useful to decision makers trying to balance conflicting objectives. The set represents the best available trade-offs between different objectives.

In the work on the lexicographic method in [73] and [74], Prats expresses a desire to use global optimisation methods to solve the multi-objective optimal control problem, enabling procedure trade-off analysis using global Pareto fronts. However, it is concluded that this is unachievable using current day technology due to computational burden [73]. The lexiographic method is then offered as a compromise approach that offers a good solution in an acceptable computational time. In the conclusions of the work in [73], the author summarises,

"Global optimisation packages were tested unsuccessfully due to the high computational burden required. At this point, we conclude that with this technique the optimisation of this kind of trajectories, in a fully automated way, is not possible with nowadays technology."

So, although the work takes a step closer to the requirements defined by the ICAO review papers, it also recognises that solving the multi-objective control problem for Pareto fronts with global methods would be a more desirable approach if attainable.

Pervier et al [75] and Torres et al [76] treated the problem of calculating Pareto optimal fronts between environmental objectives as a discrete parameter optimisation problem with bounded variables to maintain feasibility. Although not explicitly formulated as optimal control problems, the solution approach in both studies most closely resembles the direct shooting method, with explicit integration of the differential equations [77]. The NLP problem was then solved in [76] using a derivative free solver, and in [75] using a genetic algorithm. The limitations of direct shooting are well understood and the approach is most successfully applied to simple control problems having a small number of optimisation variables as there are significant trade-offs in trajectory accuracy and run-time performance [77].

Correspondingly, the environmental trade-off analysis involved idealised 2-D optimisation scenarios, a small numbers of optimisation variables, and simplified trajectories represented by small numbers of linear segments. Pervier et al [75] do not consider noise metrics, and while Torres et al [76] do, it is for a simplified noise measure with no consideration of local population distribution. Increasing trajectory accuracy for procedure optimization can potentially require thousands of optimisation variables and nonlinear constraints [73]. The trade-offs then in accuracy and computational performance would seem to limit the application of shooting methods for more complex, realistic procedure optimisation studies.

2.6.2 SES Trajectory Planning

In the current ATM system, operators have two principle goals when developing their flight plans. These are, to plan the fuel consumption of the flight and to prepare an expression of intent and required resources for submission to the ANSP [27, 52]. Typically operators plan the fuel usage for a flight according to a cost index. The cost index is usually a ratio of fuel costs to time operating costs for the flight [51]. Time costs may include staffing costs or maintenance cost related to hours of operation. The cost index is then used by the operator with AIP information to define the speed and altitude schedule for the flight. In general, aircraft operators prioritise faster cruise speeds to minimise time costs and lower cruise speeds to minimise fuel consumption costs.

For the Single European Sky (SES) concept of operation, there is a transition to trajectory based flight planning, where a 4D trajectory is used define the operators intended flight and operation within the ATM system. Implicit in the 4D trajectory is also the cost priorities of the operator. The FMS therefore no longer uses a cost index in-flight but instead tracks to the planned 4D trajectory.

Business Development Trajectory (BDT)

The business trajectory is the first step in developing the trajectory an operator will fly. The business trajectory is planned around the strategic business goals of the aircraft operator. At this stage business decisions are made regarding operating schedules, destinations, operating routes, aircraft fleet types, load factors crew costs and ANSP airspace charges.

Trajectory optimisation methods can be used at this point in a decision support capacity to evaluate the trade-offs associated with the different ways of operating the aircraft [59, 2, 63, 64]. This could include evaluating different routing options or different ways of operating the aircraft on the same route. Through the use of trajectory optimisation methods, flight planners should then be able to develop trajectories that reflect the operating priorities of the airline.

While SES trajectory based operations discussed in Section 2.5 are intended to reduce the number and severity of ATM constraints, 4D trajectories, at this point, will still need to be planned according ATM route and flow control constraints that are known to the operator [59, 2, 63, 64].

Particularly in the TMA, due to high densities of traffic and for environmental abatement considerations, it is expected that flights approaching and departing airports will continue to be constrained by arrival and departure operating procedures [62].

As discussed in Section 2.5, it is hoped that the adoption of trajectory optimisation methods will harmonise the development of these procedures, defining environmentally efficient procedure height and speed constraints that, coupled with advancements in PBN, could largely determine the aircraft climb and descent operating steps. Improvements in this form of constraint definition could also open up the possibility of aircraft type specific or category specific procedure definition.

In circumstances where operators continue to have considerable freedom to plan climb and descent aircraft operating steps along preferential routings, then the operators can avail of trajectory optimisation methods to plan climb and descent operations that provide a balance between operating costs and mitigating the different forms of environmental impact.

Within the ATM constraints, the environmental goal most likely to be prioritised at this stage is CO2 [52]. CO2 is directly related to fuel consumption and reducing fuel related costs has always been a priority for operators. The prioritisation of CO2 over other forms of environmental impact is further reinforced by the CO2 aviation emissions trading schemes, both existing and proposed, that place caps on CO2 usage and incentivise aircraft operations that minimise CO2 usage [12].

By the end of this stage then, it is expected that the operator has availed of trajectory optimisation methods to develop a 4D trajectory that is defined by their strategic business goals and meets the constraints imposed by the ATM system.

Shared Business Trajectory (SBT)

When the business development trajectory is sufficiently mature to be shared with air traffic management, then the RBT becomes the SBT. For ATM, the SBT defines the users intended use of the system.

The submission of the SBT is then analogous to the submission of the flight plan in current ATM operations. Early in the RBT cycle, the SBT can be used to forecast the demand and capacity balance of the system and can be used by ATM and the military to plan staff schedules and the availability of shared airspace [59, 2, 63, 64].

On the day of operation, the RBT will likely be updated to accurately consider metrological conditions on the day. The RBT will also need to be revised on the day to reflect up to date ATM system constraints.

As with the current system, where predicted levels of traffic exceed the capacity of the system, RBTs will need to be revised to support the optimal network management. In this case, air traffic flow control will issue updated 4D flight and runway slot constraints to the operator. There is then an iterative cycle of ANSP issued constraints and operator generated trajectory revisions until a trajectory is defined that achieves the desired goals of the operator within the limits of the ATM network performance [59, 2, 63, 64].

Where RBTs need to be revised to reflect new and updated constraints there is a role for trajectory optimisation methods. The revision of the RBT at this stage is still part of the trajectory planning process, therefore the revision of the trajectory need not occur in real time and the calculation need not occur on-board the aircraft. However it is anticipated that both the operator and the ANSP will want to lock the trajectory in as quickly as possible and that this will require the calculation to be performed relatively quickly.

Therefore trade-offs in expensive cost functions such as overall noise impact would probably be computationally cost prohibitive at this stage. It is more likely that quickly evaluated cost functions, such as deviation from the initially planned trajectory could be used here, revising the trajectory while minimising disruption to trajectory intent.

Reference Business Trajectory (RBT)

The SBT is finalised before push-back from the terminal stand, at which point it becomes the Reference Business Trajectory. The RBT then exists as a common reference between air traffic control and the aircraft. The RBT is planned at this point to satisfy ATM

constraints including route structure, flow control, and CTAs issued by airport support tools such as AMAN and DMAN. However, at this point, the RBT is not a conflict free trajectory [59, 2, 63, 64]. Separation assurance is still provided by human air traffic controllers, where the flight is cleared along its trajectory using precision clearances.

Should a conflict be predicted or the aircraft be required to be re-routed due to weather conditions then ATC using controller support tools, such as MTCD, can issue a 2D/3D/4D constrained route updates to the flight [59, 2, 63, 64]. In order to discourage open loop ATC instructions, it is intended that the 2D/3D/4D constrained route updates be issued so that they have start and end at points on the initial planned trajectory. An RBT update may also be issued to the aircraft in the form of a CTA. In both cases the flight crew is then required to recalculate the RBT on-board in light of the new constraints and share the revised RBT with ATM [59, 2, 63, 64].

In circumstances where the RBT needs to be revised to include 2D/3D/4D revisions, then the problem can be reduced to a relatively straight forward trajectory optimisation problem where the trajectory must meet the new constraints while minimising a cost function. As revisions will need to occur in real time, relatively simple cost functions such as fuel, time or deviation from the planned trajectory are likely to be considered.

Along with speed, robustness will be a major concern, as methods here will need to have exceptional convergence performance and/or have steps in place to guarantee safe separation in the case of non convergence.

However, by having human in the loop air traffic control oversee the traffic management problem, then the on-board systems need only solve designated, human reviewed, short and mid-term constraint satisfaction problems. This control of the problem complexity, could enable the safe use of on-line direct trajectory optimisation methods for generating flight crew trajectory advisories.

The Airborne Separation Assurance System (ASAS) is a flight crew support tool concept that issues trajectory advisories to the flight crew [59, 2, 63, 64]. ASAS aims to provide a very advanced see and avoid system that enables flight crews to operate commercial aircraft in a similar manner to how pilots of small aircraft manage safe separation under Visual Flight Rules, by sensing and avoiding other aircraft. While ASAS concept may not be the most efficient way to operate the flight, it does enable operators to maintain safe separation in uncontrolled airspace and opens up the possibility of increased autonomy for flight crews in controlled airspace, where controllers delegate separation assurance responsibilities to pairs or small groups of aircraft [59, 2, 63, 64].

Real time trajectory optimisation methods can and have been applied to the problem of self separation, and would seem like good candidates for underlying ASAS systems [78, 79, 80, 81]. However, there are still many questions surrounding ASAS in controlled airspace. In particular, in high density traffic situation, where the complexity of the optimisation problem increases exponentially, and where the resolution of primary conflicts can cause

secondary conflicts, can ASAS systems guarantee provision of separation? Also, in such high density situations, if they can, are they really more efficient (fuel, time, etc) than a human centric organised air traffic management system?

2.6.3 Thesis Goals - Climb and Descent Procedure Optimisation

In summarising the literature, currently there is no methodology that meets all the requirements set out in the ICAO reviews, providing a harmonised approach and allowing for the identification of trade-offs in procedures optimised for several conflicting environmental objectives. Recapping the ICAO high level environmental goals as,

- to limit or reduce the number of people affected by significant aircraft noise,
- to limit or reduce the impact of aircraft engine emissions on local air quality,
- to limit or reduce the impact of aviation greenhouse gas emissions on the global climate,

then there is a need for an approach that can be used alongside the goals to manage the trade-offs in noise and emissions in an informed manner. This thesis proposes such a methodology. The method aims to be a computationally efficient approach that uses current best practice environmental cost models and treats the problem as a multi-objective optimal control problem. The problem is then solved for a set of Pareto optimal solutions by discretising the optimal control problem with a direct numerical method and solving the resulting Non Linear Programming (NLP) problem with the stochastic Differential Evolution solver.

Exploring the efficacy of the approach, the method is applied to complex, real world environmental procedure optimisation case studies. This is done to determine if the proposed method can identify procedures that provided better trade-offs between the environmental impacts than those proposed in the initial studies.

Also of note, the method will make use of the noise Annoyance Score metric. In the literature, noise abatement procedure optimisation has focused on awakenings from night time flights or maximum noise levels for sensitive sites [68, 69, 70, 71, 73]. There is a dearth of research on operational procedures optimised for reductions in overall levels of community noise impact. The Annoyance Score metric is a metric that was developed as part of the MIME (Market-based Impact Mitigation for the Environment) project to measure relative changes in community noise due to changes in flight operating procedures [24]. The metric was developed from the LEQ metric and therefore has a direct relationship with overall community noise impact (see Chapter 4). The metric is a single value, single event noise impact measure, making it easily integratable to noise optimisations studies (see Section 4). However, prior to work in this thesis, the Annoyance Score metric has not been investigated for this purpose.

None of the trajectory optimisation based research reviewed in the literature considered airspace structure constraints. Although a relatively trivial inclusion, a multi-objective case study in this thesis will include airspace structural constraints. It will be shown how airspace structural constraints can be represented as simple geometrical shapes for easy inclusion in optimisation studies.

A trend also observed in the literature is the use of commercial or proprietary software to solve the trajectory optimisation problem. All research conducted through the Delft University of Technology used the commercial optimiser EZopt [68, 69, 70, 71]. Houacine used the commercial optimiser SNOPT [72]. Prats uses the commercial optimiser GAMS [73]. GATAC, used in [75], is proprietary to the Cleansky project. The Mesh Adaptive Direct Search (MADS) optimiser used in [76] is open source, however the flight dynamics and noise calculation methods used in the work are Airbus proprietary.

In some cases, as with the SNOPT solver, alternative general NLP solvers can be easily swapped in for problem solution. However, in other cases, such as the NOISHHH and the lexiographic work, the definition of the NLP problem is tied closely to the solver used. Methods for calculating the environmental impact of procedure designs are generally open source or open method and utilise freely available data. Therefore, in keeping with this approach, no closed source or proprietary methods have been used for the work in this thesis.

2.6.4 Thesis Goals - Flight Trajectory Planning

The most obvious extension of the approach defined for the optimisation of climb and descent procedures discussed in Section 2.6.3, is the application of the IDVD-DE method to the planning of user-preferred flight trajectories in the Business Development Trajectory (BDT) phase of the reference trajectory development. Therefore work in this thesis will also look briefly at the application of trajectory optimisation methods to trajectory based flight planning that will form the basis of SES trajectory based operations.

Work in this thesis has always looked to adopt real world studies and reference solutions for comparison to IDVD-DE optimisation results. However, as SES trajectory based operations do not exist yet, there was a need to find an alternative reference for comparison.

The NATS 3Di metric is a flight efficiency metric based around the SES user preferred trajectory concept [39, 82]. Aligning with the SES CO2 flight efficiency target, the metric aims to measure fuel/CO2 inefficiencies in flight trajectories and is intended to drive improvements in flight operations. The metric works by comparing flown flight trajectories to NATS defined CO2 optimal flight trajectories to determine fuel/CO2 inefficiencies. The 3Di measure then, and it's associated benchmarking trajectories, provide ideal points of comparison for CO2 efficient trajectories generated by the IDVD-DE method. The added benefit of adopting the 3Di method as reference is that any observations made through

comparisons with the IDVD-DE method are equally relevant to both current and SES based trajectory planning.

Chapter 3

Aircraft Emissions

3.1 Introduction

Aircraft gaseous emissions resulting from the combustion of fossil fuel have significant effects on the concentrations of climate changing greenhouse gases in the Earth-atmosphere system. Aircraft gaseous and particulate emissions also have potential human health concerns for communities local to airports. Emissions are therefore discussed here relative to their impact on climate change and local air quality. Some emissions, such as NOx, acting through different mechanisms, have both climate change and local air quality impacts, and are therefore discussed in both sections.

3.2 Climate Change

Radiative Forcing

Radiative forcing is a measure of the influence a climate change mechanism has in altering the energy balance of incoming and outgoing energy in the Earth-atmosphere system. The key climate changing mechanisms for aviation are the emissions carbon dioxide (CO2), oxides of nitrogen (NOx) and water vapor (H2O) [3].

CO2

CO2 and water vapor are the main products of aircraft engine combustion. As CO2 has a large, well understood impact on net radiative forcing, it is considered to be aviation's most important climate change contribution [3]. Therefore CO2 reduction or mitigation has become a key target for improvements in airframe/engine design and for improvements in aircraft operation.

For improvements in operations, ICAO has proposed a global market based cap and trade system for CO2 emissions to control greenhouse gas emissions produced by aircraft [11, 12]. In Europe a CO2 emissions trading scheme already exists for aircraft operating within the EU [8, 9]. There are also European targets set by ACARE (from a 2000 baseline) and the Single European Sky (from a 2005 baseline) for a 10% improvement in the operational CO2 efficiency of flights, to be achieved through combined improvements in air traffic management and airline operating procedures [16, 13].

H20

Water vapor is the other main product that accompanies CO2 when fuel is burned completely the engine. When the hot water vapor is released into particularly cold air within a certain humidity range, the water vapor can freeze forming contrails [83].

It has been proposed that contrails have significant radiative forcing effects both in them selves and through their contribution to cirrus cloud formation [3]. This is thought to be due to contrail and cloud cover blocking more outgoing radiation energy than incoming radiation energy, leading to a climate warming effect. However, there is a lack of scientific understanding of the mechanisms by which contrails affect radiative forcing and therefore there is substantial uncertainty in the radiative forcing values associated with contrails. Consequently, there are are no current aviation related policies regarding the mitigation of contrail formation [84].

Aviation soot and sulphur are thought to have negligible direct climate changing effects. However, both can interact with water vapor to increase the likelihood of contrail formation [85].

NOx

Nitrogen gas (N2) is the largest constituent of the Earth's atmosphere. Oxides of nitrogen (NO and NO2, collectively referred to as NOx) are formed by the oxidation of atmospheric nitrogen in the high temperature of the combustion chambers [86]. NOx does not have a direct radiative forcing impact, but it's emissions at different altitudes can result in positive radiative forcing (warming) through ozone production and negative radiative forcing (cooling) via methane reduction [3, 87]. Aviation emissions of NOx are considered to have a greater radiative forcing effect than NOx emissions from other forms of transport as NOx emissions at cruise altitudes have a longer atmospheric lifetime [3].

NOx has been proposed as an emission suitable for building an emissions trading scheme around, and has been proposed for inclusion in existing and planned emissions trading schemes [88, 89, 87]. However, uncertainties in determining the radiative forcing of NOx

has generally led to CO2 emissions being the preferred climate change measure for anthropogenic global warming. CO2 is also then they key target for reductions in anthropogenic climate change.

While NOx reduction remains a target for engine manufactures, largely due to the human health impacts of NOx emissions at lower altitudes [3], there are currently no policies regarding aircraft operations and the climate changing effects of NOx.

3.3 Local Air Quality

Local Air Quality (LAQ) can be described generally as the condition of the ambient air to which humans and nature are typically exposed [3]. Poor air quality has an adverse impact on human health, particularly with regard to the development respiratory illnesses. While in theory, the complete combustion of fuel in the engine results in only water and CO2 emissions, in reality, the burning of fuel in the engine provides a range of other gaseous and particle emissions. These are regarded as pollutants [90]. Emissions that occur from incomplete burning of fuel include carbon monoxide (CO), hydrocarbons (HC), black carbon (soot) and volatile organics. Other emissions occuring from impurities in the fuel are oxides of sulphur and sulphates.

Since the 1970s, ICAO, through it's Committee on Aviation Environmental Protection (CAEP) has aimed to control potential adverse effects of air pollutants on LAQ, primarily pertaining to human health and welfare [90]. To this end, as part of the engine certification process, engine pollutant rates are determined at a series of thrust settings. The thrust settings correspond to those typically used for take-off, climbing, landing and taxiing to and from the stand. Combined, the rates are used to determine the Landing and Take-Off (LTO) emissions performance for engine/airframe combinations. Around the LTO cycle, CAEP defined a series of limits to classify aircraft local emissions performance [3]. The LTO performance and associated limits have then been used by ICAO states to define policy and by airport operators to incentivise the use of aircraft with the best LTO performance. In general, the standards are intended to encourage the adoption of the new and efficient engines into the operating fleet. The CAEP standards have been credited for driving improvements in the emissions profiles of engines and for acting to mitigate aviations impact on local air quality [91]. However, the LTO cycle and associated limits, are historically intended to drive technology improvements rather than operational improvements. There is, for instance, no LAQ equivalent of the recommended noise abatement operating procedures.

The European Union strategy on ambient air quality is set out in Directive 96/62/EC (EC, 1996 and 2003) [22]. Related to the management of LAQ by states and airports Local air quality thresholds relevant to airports are then set out the daughter directive 99/30/EC, which take the form of limit values, target values, alert thresholds and margins of tolerance

for the emissions of sulphur dioxide SO2, oxides of nitrogen NOx and particulate matter PM10 [92, 93].

Gaseous Emissions

Gaseous emissions having an impact on air quality are carbon monoxide, hydrocarbons, oxides of nitrogen and oxides of sulphur (SOx).

Carbon monoxide CO is a toxic gas to human beings. Carbon monoxide occurs where there is not enough air for the complete combustion of hydrocarbon fuel and oxygen to CO2, with the incomplete combustion resulting in CO. Hydrocarbons are also unburned fuel or fractions of fuel that come from incomplete combustion.

NOx is the combined term for oxides of nitrogen, which are nitric oxide (NO) and nitrogen dioxide (NO2). NOx is formed in the combustion process where the high combustion temperatures cause oxidation of atmospheric nitrogen. NOx can have a direct human health effect, but also has an indirect impact due its role in the formation of ozone and acid rain.

There are a number of NOx and HC charging schemes utilised by airports within Europe. Where applied, airports charge a fee for each kilogram of NOx or HC emitted by an aircraft during its LTO cycle [94, 95]. In general the charges pertain to NOx emissions, with hydrocarbon (HC) charges only utilised at airports that cater for older aircraft that tend to have low NOx but relatively high HC emissions [96]. The estimate of the emissions is done with the "times in mode" calculation, which produces a NOx estimate based on how long the aircraft spent in each part of the LTO cycle [96]. The times in mode calculation is relatively insensitive to the efficiency of aircraft climb and descent operation. Therefore there is currently no incentive to minimise NOx emissions through aircraft operation. However, when aircraft NOx was calculated using a more accurate, "advanced" emissions calculation method equivalent to the Boeing Fuel Flow Methodology (BFFM) [97], it was found that aircraft climb out and descent operations could play a significant role in minimising NOx emitted [98, 99]. Reducing NOx from operations though will likely require trade-offs with other emissions and noise.

SOx is the combined term for oxides of sulphur, which are sulphur dioxide (SO2) and sulphur trioxide (SO3). Gaseous SOx is produced when sulphur impurities in the fuel are oxidised to either SO2 or SO3. SOx emissions are a function of the quantity of sulphur in the fuel and are proportional to the fuel burned. As with CO2 and water vapor, the direct relationship between fuel and SOx means that fuel burned is often used as a surrogate metric to show how operational changes will increase or decrease SOx emissions [94, 92].

Particulates

Particulate matter (PM) is a pollution composed of very small particles. PM(10), for example, is composed of particles that are 10 micrometer or smaller in diameter [85].

Soot/black carbon consists of ash and unburned solid carbon from the fuel. Volatile organics are other solids and liquids resulting in the incomplete combustion of fuel. Sulphates are solid particles that result from the oxidisation of sulphur impurities in the fuel during the combustion process.

Particulate emissions from aircraft flight operations are considered secondary to other sources of particulate emissions at airports, which include auxiliary power units, air-side ground support vehicles, heating plants and land-side road traffic [100]. Therefore, flight trajectory particulate emissions are currently not seen as a driver for changes in flight operations [22]. However, particulate emissions from aircraft climbs and descents to and from 3000 ft are typically included in measuring the airports particulate emissions footprint [100].

3.4 Gaseous Emissions Modelling

For the calculation of aircraft emissions, all emissions are calculated as a function of fuel burn. Specifically, for turbofan aircraft, the rate of fuel burn $\dot{W}_f(t)$ (kg/s) is calculated as

$$\dot{W}_f(t) = \begin{cases} \dot{W}_{f_{nom}}(t) = \eta(t)T(t) & \text{:if } \dot{W}_{f_{nom}}(t) > \dot{W}_{f_{min}}(t) \\ \dot{W}_{f_{min}}(t) = \frac{C_{f3}}{60}\left(1 - \frac{\mathsf{H_p}(t)}{C_{f4}}\right) & \text{:if } \dot{W}_{f_{nom}}(t) \leq \dot{W}_{f_{min}}(t) \end{cases} \tag{3.1}$$

where the nominal fuel flow $\dot{W}_{f_{nom}}(t)$ is calculated by multiplying the aircraft thrust T (kN) by η, the BADA thrust specific fuel flow coefficient (kg/(s-kN)), which is itself a function of flight level. The fuel flow for the aircraft is set as the nominal fuel flow unless the nominal fuel flow is less than the minimum fuel flow $\dot{W}_{f_{min}}(t)$, which is calculated from the geopotential height $\mathsf{H_p}$ (ft) and the BADA fuel coefficients C_{f3} and C_{f4}. For further information on the fuel model, see Appendix B.

For the calculation of aircraft emissions, the emissions CO2, water vapor (H2O) and sulphur oxides (SOx) are calculated using direct multipliers on fuel burn, such that their rate of emissions in (g/s) is

$$\dot{e}(t) = f_e(\dot{W}_f(t), \boldsymbol{\alpha}) \tag{3.2}$$

where $\boldsymbol{\alpha}$ is a set of fuel burn multipliers [92, 101] such that

$$\begin{aligned} \alpha_{CO2} &= 3155.0 \text{ g/kg} \\ \alpha_{SOx} &= 1 \text{ g/kg} \\ \alpha_{H2O} &= 1237.0 \text{ g/kg} \end{aligned} \tag{3.3}$$

The emissions for hydrocarbons (HC), carbon (CO), and oxides of nitrogen (NOx) are calculated using the Boeing Fuel Flow Methodology (BFFM) [97]. The BFFM is a method recommended by ICAO for the calculation of aircraft emissions from fuel burn information [92].

The BFFM uses open source information from the ICAO engine emissions databank (EDB). The ICAO emissions databank contains empirical information obtained as part of the engine certification process that relates fuel burn to emissions indices at 4 different engine thrust settings. Figure 3.1 shows the EDB datasheet for the IAE V2530 engine. The BFFM then offers a procedure for interpolating this data to define a set of in flight emissions indices **EI** that allow for the calculation of in-flight rates of NOx, HC, and CO emissions as

ICAO ENGINE EXHAUST EMISSIONS DATA BANK

SUBSONIC ENGINES

ENGINE IDENTIFICATION: V2530-A5
UNIQUE ID NUMBER: 1IA005
ENGINE TYPE: MTF
BYPASS RATIO: 4.54
PRESSURE RATIO (π_{oo}): 32.1
RATED OUTPUT (F_{oo}) (kN): 133.4

REGULATORY DATA

CHARACTERISTIC VALUE:	HC	CO	NOx	SMOKE NUMBER
D_p/F_{oo} (g/kN) or SN	0.4	24.2	67.2	11.6
AS % OF ORIGINAL LIMIT	2.0 %	20.5 %	64.5 %	53.0 %
AS % OF CAEP/2 LIMIT (NOx)			80.7 %	
AS % OF CAEP/4 LIMIT (NOx)			94.4 %	
AS % OF CAEP/6 LIMIT (NOx)			106.4 %	
AS % OF CAEP/8 LIMIT (NOx)			123.7 %	

DATA STATUS

- \- PRE-REGULATION
- x CERTIFICATION
- \- REVISED (SEE REMARKS)

TEST ENGINE STATUS

- x NEWLY MANUFACTURED ENGINES
- \- DEDICATED ENGINES TO PRODUCTION STANDARD
- \- OTHER (SEE REMARKS)

EMISSIONS STATUS

- x DATA CORRECTED TO REFERENCE (ANNEX 16 VOLUME II)

CURRENT ENGINE STATUS

(IN PRODUCTION, IN SERVICE UNLESS OTHERWISE NOTED)

- \- OUT OF PRODUCTION
- \- OUT OF SERVICE

MEASURED DATA

MODE	POWER SETTING (%F_{oo})	TIME minutes	FUEL FLOW kg/s	EMISSIONS INDICES (g/kg) HC	CO	NOx	SMOKE NUMBER
TAKE-OFF	100	0.7	1.331	0.045	0.45	33.8	-
CLIMB OUT	85	2.2	1.077	0.041	0.52	27.1	-
APPROACH	30	4.0	0.377	0.056	1.81	10.1	-
IDLE	7	26.0	0.138	0.1	10.95	5	-
LTO TOTAL FUEL (kg) or EMISSIONS (g)			504	35	2620	7732	-
NUMBER OF ENGINES				1	1	1	1
NUMBER OF TESTS				3	3	3	3
AVERAGE D_p/F_{oo} (g/kN) or AVERAGE SN (MAX)				0.26	19.7	58	9
SIGMA (D_p/F_{oo} in g/kN, or SN)				-	-	-	-
RANGE (D_p/F_{oo} in g/kN, or SN)				-	-	-	-

ACCESSORY LOADS

POWER EXTRACTION 0 (kW) AT - POWER SETTINGS
STAGE BLEED 4.5 % CORE FLOW AT 0.07 POWER SETTINGS

ATMOSPHERIC CONDITIONS

BAROMETER (kPa)	101.0-101.8
TEMPERATURE (K)	288 - 296
ABS HUMIDITY (kg/kg)	.0086-.0114

FUEL

SPEC	Jet A
H/C	-
AROM (%)	-

MANUFACTURER: International Aero Engines
TEST ORGANIZATION: Pratt & Whitney
TEST LOCATION: East Hartford, CT, USA
TEST DATES: FROM 12 Aug 92 TO 13 Aug 92

REMARKS

1. Engine Type description originally wrongly given as TF; amended in Issue 15 to MTF

FIGURE 3.1: ICAO Engine emissions datasheet for V2530 engine

$$\dot{e}(t)_{BFFM} = f_{BFFM}(\dot{W}_f(t), \mathbf{EI}) \tag{3.4}$$

The first step in the BFFM is to adjust the ICAO engine database fuel flows $\dot{\mathbf{W}}_{f_{EDB}} = [\dot{W}_{f_{TO}}, \dot{W}_{f_{CL}}, \dot{W}_{f_{APP}}, \dot{W}_{f_{IDLE}}]$ by a set of correction factors to adjust the engine fuel flows for the effects of installing the engine on an aircraft. $\dot{W}_{f_{TO}}, \dot{W}_{f_{CL}}, \dot{W}_{f_{APP}}, \dot{W}_{f_{IDLE}}$ are the engine database fuel flows for take-off, climb, approach and idle. The set of terms for the installation adjustment, common for all aircraft types, is $\mathbf{i}_a = [1.010, 1.013, 1.020, 1.100]$ such that the adjusted engine database fuel flows $\dot{\mathbf{W}}_{f_{AEDB}}$ are

$$\dot{\mathbf{W}}_{f_{AEDB}} = \mathbf{i}_a^T \circ \dot{\mathbf{W}}_{f_{EDB}}^T \tag{3.5}$$

where $\circ$ designates the element wise multiplication of the vectors.

The engine database emissions indices of HC, CO and NOx (g/kg) are not adjusted.

The next step is to create plots of the adjusted engine database fuel flows versus the engine database emissions indices of HC, CO and NOx. Figure 3.2 shows the two standard Log-Log interpolation plots for fuel flow ($\dot{\mathbf{W}}_{f_{AEDB}}$) against the emissions indices of NOx, HC, and CO at sea level static conditions.

For NOx emissions, a least squares linear line is fitted to the 4 emissions database data points, reflecting the linear relationship between fuel burn and NOx. For the emissions HC and CO, a two segment plot is created consisting of a diagonal line connected to a flat horizontal line. The diagonal line reflects how emissions of HC and CO decrease with increasing thrust and fuel flows. The inverse relationship between the emissions and fuel flow occurs because of the more efficient burning of fuel that occurs at high combustion temperatures related to higher aircraft thrust settings. The connected horizontal lines reflect that there is a point where the reduction in emissions plateaus and that beyond this point there are no further reductions in the emissions for higher levels of fuel flow.

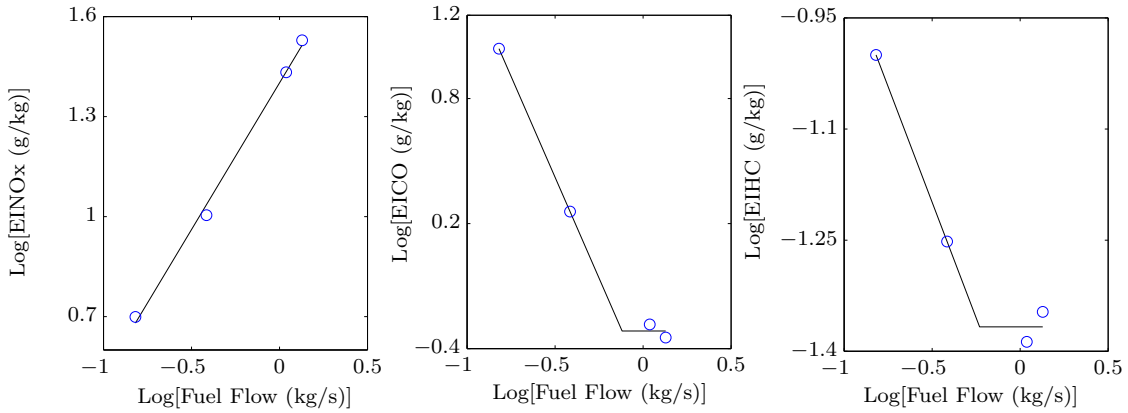

FIGURE 3.2: Log-Log plots of fuel burn against emissions indices for a A320/V2530 aircraft

For the next step the, in flight, operational fuel flow ($\dot{W}_f$) is corrected to sea level conditions, and emissions indices for the reference conditions are interpolated off the Log-Log plots.

The in flight fuel flow is corrected for the number of engines N_E and to the database reference conditions by

$$\dot{W}_{ff} = \frac{\dot{W}_f}{N_E\, \delta_{amb}} \theta_{amb}^{3.8} \mathrm{e}^{0.2 {M_a}^2} \tag{3.6}$$

where $\dot{W}_{ff}$ is the corrected fuel flow factor, M_a is Mach number and

$$\begin{aligned}\theta_{amb} =& \frac{\mathsf{T}_{amb}}{288.15}\\ \delta_{amb} =& \frac{\mathsf{P}_{amb}}{101.325}\end{aligned}$$

where P_{amb} and T_{amb} are the ambient pressure (kPa) and Temperature (Kelvin) at the geopotential pressure altitude $\mathsf{H_p}$ in feet, such that θ_{amb} is the ratio of inlet temperature to sea level temperature and δ_{amb} ratio of inlet pressure to sea level pressure. The ambient conditions are calculated as

$$\begin{aligned}\mathsf{P}_{amb} &= \mathsf{P}_{amb_o} - 3.6197 \cdot 10^{-3} \cdot \mathsf{H_p} + 4.9061 \cdot 10^{-8} \cdot {\mathsf{H_p}}^2 - 2.5191 \cdot 10^{-13} \cdot {\mathsf{H_p}}^3\\ \mathsf{T}_{amb} &= \mathsf{T}_{amb_o} - 1.9813 \cdot 10^{-3} \cdot \mathsf{H_p}\end{aligned} \tag{3.7}$$

The corrected fuel flow factor $\dot{W}_{ff}$ can then used with the log-log plots to interpolate the reference emissions indices (g/kg) *REIHC*, *REICO*, $REINO_x$. The reference emissions indices are the emissions indices for the data bank reference conditions. The reference indices then need to be corrected back to operational conditions, such that the in flight emissions indices *EIHC*, *EICO*, $EINO_x$ are given as

$$\begin{aligned}EIHC =& REIHC \frac{\theta_{amb}^{3.3}}{\delta_{amb}^{1.02}}\\ EICO =& REICO \frac{\theta_{amb}^{3.3}}{\delta_{amb}^{1.02}}\\ EINO_x =& REINO_x \mathrm{e}^{H_f} \left(\frac{\delta_{amb}^{1.02}}{\theta_{amb}^{3.3}}\right)^{1/2}\end{aligned} \tag{3.8}$$

where, for the calculation of $EINO_x$, the humidity correction factor H_f is calculated as

$$H_f = -19\,(\omega - 0.0063) \tag{3.9}$$

where ω is the specific humidity

$$\omega = \frac{0.62198\,\varrho\,P_v}{\frac{\mathsf{P}_{amb}}{6.895} - \varrho\,P_v} \tag{3.10}$$

and where the saturation vapour pressure is $P_v = 0.014504 \cdot 10^{\beta}$ (psia), and where the relative humidity ϱ (%) is assumed to be 0.6. β is then calculated as

$$\beta = 7.90298\left(1 - \frac{373.16}{\mathsf{T}_{amb}}\right) + 3.00571 + (5.02808)\log\left(\frac{373.16}{\mathsf{T}_{amb}}\right) + (1.3816 \cdot 10^{-7})\left[1 - 10^{11.344\left(1-\frac{\mathsf{T}_{amb}+273.16}{373.16}\right)}\right] + (8.1328 \cdot 10^{-3})\left[1 - 10^{3.49149\left(1-\frac{373.16}{\mathsf{T}_{amb}+273.16}\right)} - 1\right]$$

Quantities of carbon monoxide S_{CO}, hydrocarbons S_{HC} and oxides of nitrogen S_{NO_x} for the entire trajectory in kilograms are then given as

$$S_{HC} = N_E \cdot \int_{t_f}^{t_o} EIHC(t)\,\dot{W}_f(t) \cdot 10^{-3}\,dt$$
$$S_{CO} = N_E, \cdot \int_{t_f}^{t_o} EICO(t)\,\dot{W}_f(t) \cdot 10^{-3}\,dt$$
$$S_{NO_x} = N_E, \cdot \int_{t_f}^{t_o} EINO_x(t)\,\dot{W}_f(t) \cdot 10^{-3}\,dt$$

3.5 Particulate Emissions Modelling

The First Order Approximate (FOA) method is a recently developed approach proposed by ICAO to calculate the emissions of particulate matter, namely black carbon, sulphates and volatile organics [92]. As the engine data in the emissions databank does not contain reference emissions indices for black carbon, sulphates and volatile organics, the approximate method defines a series of work arounds, supplementing data from the databank with external information in order to estimate the required emissions indices. ICAOs own summary of the method [92] includes the caveat

“ the user should be aware that not all physical concepts are well understood and data for many of the parameters are sparse. This leads to uncertainties in the estimation methodology including a lack of data in the ICAO EDB.”

Further, the first order approximate method is only valid for sea level conditions. Due to missing input parameters and uncertainty in the veracity of the method, there is no policy or recommendations related to the impact of individual trajectory operations on particulate matter. For this reason, particulate matter emissions were never set an an objective in any of the real world case studies taken as baselines for the optimisation studies in this work. However, the FOA method is included here for completeness and for any future research wishing to include particulate emissions modelling.

Black Carbon/Soot

To estimate the smoke created in the engine, the engine operating air:fuel ratio and smoke number are used. The smoke number is a dimensionless term that quantifies smoke emissions. As the engine operating air:fuel ratio is not contained in the ICAO EDB and as the smoke number information in the EDB is highly fragmented [92], the parameters need to be estimated. The relationship between the corrected fuel flow factor $\dot{W}_{ff}$ and the air:fuel ratio is taken as an air fleet average, and is shown in Table 3.1.

TABLE 3.1: Relationship between ICAO engine database fuel flows and representative AFRs

$\dot{\mathbf{W}}_{f_{AEDB}}$	AFR
$\dot{W}_{f_{IDLE}}$ (idle)	106
$\dot{W}_{f_{APP}}$ (approach)	83
$\dot{W}_{f_{CL}}$ (climb-out)	51
$\dot{W}_{f_{TO}}$ (take-off)	45

A similar table is required to define the relationship between the fuel flows in $\dot{\mathbf{W}}_{f_{AEDB}}$ and the smoke numbers at the different power settings. This information may be available in the engine datasheets but is typically incomplete or not available at all. Therefore there are are a number of methods suggested for estimating the missing engine data base data, covered in detail in [93, 92]. Once the relationships between fuel flow and air:fuel ratio (AFR), and between fuel flow and smoke number (SN), have been defined, the smoke number and the air:fuel ratio for the corrected fuel flow factor $\dot{W}_{ff}$ can be determined through linear interpolation. The emissions index of black carbon (g/kg) is then calculated as

$$EIBC = (0.776\,AFR + 0.877)\frac{CL}{1000} \tag{3.11}$$

where the carbon index CL (mg/m^3 produced by burning 1 kg of fuel) can be determined from

$$ln(CL) = 1.23357\,\ln(SN) - 2.66997 \tag{3.12}$$

The emissions rate (g/s) of black carbon for the instantaneous operating fuel flow can then be determined as

$$\dot{e}_{BC} = N_E \cdot \dot{W}_f \cdot EIBC \tag{3.13}$$

Sulphates

The EI for sulphate particulate matter $EI_{vol-FSC}$ (g/kg) is given as

$$EI_{vol-FSC} = 3 \cdot 10^3 \times FSC \cdot \varepsilon \tag{3.14}$$

where Fuel Sulphur Content $FSC = 0.00068$, $\varepsilon = 0.033$ defines the fractional conversion of SO2 (S^{IV}) to SO3 (S^{VI}) via oxidation.

The emissions rate of sulphates (g/s) for the instantaneous operating fuel flow can then be determined as

$$\dot{e}_{vol-FSC} = N_E \cdot \dot{W}_f \cdot EI_{vol-FSC} \tag{3.15}$$

Volatile Organics

Measurements of condensable organics in the engine exhaust are very limited [92]. Based on the assumption that condensable organics are directly related to unburned hydrocarbons, an estimate is made by scaling the engines hydrocarbon (HC) EI by a factor δ_f [92]. The fuel flow to δ_f relationship is shown in Table 3.2 [92].

$\dot{\mathbf{W}}_{f_{AEDB}}$	δ_f (g/kg)
$\dot{W}_{f_{IDLE}}$ (idle)	0.115
$\dot{W}_{f_{APP}}$ (approach)	0.076
$\dot{W}_{f_{CL}}$ (climb-out)	0.05625
$\dot{W}_{f_{TO}}$ (take-off)	0.00617

TABLE 3.2: Relationship between ICAO engine database fuel flows and δ_f

Using Table 3.2, the emissions index of volatile organics (g/kg) is

$$EI_{volFuelOrganics} = \delta_f \cdot EIHC \tag{3.16}$$

where δ_f is determined by linear interpolation of the fuel flow δ_f relationship in Table 3.2 using the corrected fuel flow factor $\dot{W}_{ff}$. $EIHC$ is the emissions index (g/kg) of hydrocarbons at $\dot{W}_{ff}$.

The emissions rate (g/s) of volatile organics for the instantaneous operating fuel flow can then be determined as

$$\dot{e}_{volFuelOrganics} = N_E \cdot \dot{W}_f \cdot EI_{vol-FSC} \tag{3.17}$$

3.6 Summary: Aircraft Emissions

In summary, the aircraft emission type whose effects are best understood and that is currently subject to the greatest regulatory attention is carbon dioxide (CO2). Therefore in all of the case studies in this thesis where aircraft emissions are considered, CO2 is taken as the primary emission type of interest. In the multi-objective case studies, CO2 is also the primary antagonist to noise impact objectives (see Sections 7.3,8.2,8.3). As, for work done in this thesis, CO2 is calculated as approximately three times fuel burn, for a number of case studies herein, fuel consumed is used directly as a surrogate for CO2.

The primacy of CO2 as the emission type of focus is subject to constant review by the scientific community. As the understanding of in particular the climate impacting effects of different aircraft emissions types improves, policy focus may shift from CO2 to other forms of aircraft emissions such as NOx. Also, although not traditionally a driver for flight operational changes, aircraft/airport LAQ emissions are monitored as a matter of European law. So, although CO2 is the primary emission of interest in this work, it would be remiss not to include other emission types in the inventory of emissions calculation methods to be used along the IDVD-DE trajectory method.

It was found in the literature that there was a relative sparsity of procedure design studies that looked at the trade-offs between noise and non-CO2 emissions. This may be due to the difficulty already inherent in managing the trade-offs in CO2 and noise. In response to this, for the Luton departure case study explored in Section 8.3, the IDVD-DE method is applied to not just consider trade-offs in CO2 and noise, but also to examine how air quality emissions (NOx, HC and CO) change relative to fuel/CO2 values. By supporting this sort of analysis, it can be seen that the IDVD-DE approach has the potential to support a more varied trade-off analysis than is typically considered.

Chapter 4

Aircraft Noise Impact

4.1 Introduction

Noise is defined as unwanted sound perceived as disturbing or even painful [102]. While there are no strategic ATM related targets for noise mitigation, noise is covered EU legislation (EU) No 598/2014, which is in line with the rules and procedures for reducing aviation noise laid out in the ICAO Balanced Approach to Aircraft Noise Management Doc 9829 AN/451 [21]. The Balanced Approach recognises that reducing the impact of aviation noise is not achieved through any one approach to noise mitigation, but by adopting a combination of approaches. These include the reduction of noise at source, land-use planning and management, noise abatement operational procedures and operating restrictions. EC Directive 2002/49 also requires airport operators to regularly develop action plans, where appropriate using noise mapping, for the mitigation of noise around airports.

4.2 Noise Measures

For assessing noise impact at a trajectory level, single event noise metrics are used. A single event is a single aircraft operation, such as an departure, arrival or over-flight. The extent of the unacceptability of the sound depends on the physical characteristics of the sound, notably its intensity, duration and frequency. A number of metrics have evolved to capture one or more of these physical characteristics. Single event metrics commonly used to examine aircraft noise are L_{Amax}, Sound Exposure LeveL (SEL), and the Effective Perceived Noise Level (EPNL) [102]. L_{Amax} and SEL are derived from the Sound Pressure Level (SPL) metric and EPNL is derived from the Perceived Noise Level (PNL) and the SPL.

The Sound Pressure Level

Sound is caused by changes in air pressure. The SPL (SPL or L_p) is the difference in sound pressure from a reference level caused by a sound wave. SPL is the base metric used for constructing single event noise metrics. The SPL is a function of the logarithmic ratio between the measured Root Mean Squared (RMS) sound pressure p_{RMS} and a reference sound pressure p_{ref}, both measured in Pascals (Pa). The units of the SPL are decibels, and the logarithmic nature of the metric allows the metric to capture the sensitivity of the human ear to a wide range of sound intensities. The SPL is calculated as:

$$L_p = 10 \log_{10} \left(\frac{p_{RMS}^2}{p_{ref}^2} \right) \tag{4.1}$$

with the root mean square pressure p_{RMS}^2 calculated from

$$p_{RMS}^2 = \frac{1}{t_2 - t_1} \int_{t_1}^{t_2} p(t)^2 dt \tag{4.2}$$

where t_1 and t_2 are the start and end times of the sound event and p is pressure in Pascals. The reference pressure p_{ref} is then

$$p_{ref} = 20 \cdot 10^{-6} \tag{4.3}$$

$\mathbf{L_{Amax}}$

To better correlate sound levels to perceived sound, the SPL (L_p) is weighted with respect to loudness to produce the A-weighted SPL (L_A) measured in A-weighted decibels (dBA). The maximum A-weighted sound pressure level for a noise event during a specific time interval is then the L_{Amax} (dBA).

SEL

The Sound Exposure Level (SEL), measured in dBA, considers both the loudness and the duration of the noise event above a low cut off point of $L_{Amax} - 10$dB . The total noise energy for the event is then normalised to reference time t_{ref} of 1 second. The SEL is calculated as

$$SEL = 10 \log_{10} \left(\frac{1}{t_{ref}} \int_{t_1}^{t_2} 10^{\frac{L_A(t)}{10}} dt \right) \tag{4.4}$$

where t_1 to t_2 is the duration of the noise event above the cut-off limit.

PNL and PNLT

The Perceived Noise Level (PNL) was developed to be a measure of the noisiness of aircraft [103]. It accounts for how the human ear perceives both the loudness and frequencies of broadband sound. The PNL, measured in PNdB, is calculated as:

$$PNL = 40 + \frac{10}{\log_{10} 2} \log N \tag{4.5}$$

where N is the total perceived noisiness, in Noy, as defined in [104]. The PNL is further augmented by a tone correction C (dB) [104] to produce the Tone-corrected Perceived Noise Level ($PNLT$) metric (PNdB). The tone correction is added to PNLs to reflect how certain discrete tones are perceived to be more annoying than indicated purely by PNL values.

$$PNLT = PNL + C \tag{4.6}$$

EPNL

The Effective Perceived Noise Level (EPNL) is an extension of PNLT in that it accounts for the duration of a noise event through the integration of the PNLT noise energy normalized to a 10 second time interval t_{ref}.

$$EPNL = 10 \log_{10} \left(\frac{1}{t_{ref}} \int_{t_1}^{t_2} 10^{\frac{PNLT(t)}{10}} dt \right) \tag{4.7}$$

The integration interval t_1 to t_2 is defined by the time that the PNL is greater than $PNL_{max} - 10$dB. The unit of the EPNL is denoted as EPNdB. EPNL is used by the FAA for certification of large transport and turbojet aircraft, helicopters and heavy propeller driven aircraft.

Noise Annoyance Score

The noise Annoyance Score, developed as part of the MIME study [24], is a single event noise metric that attempts to include consideration of the annoyance experienced by populations overflown. The Annoyance Score was developed from the Equivalent Noise Level (LEQ) metric by considering the relationship between the SEL measure and LEQ. LEQ (4.8) is a long term noise metric that is used with dose response functions to describe the relationship between noise levels and community annoyance [102]. The discrete form of LEQ involves the sum of SELs from multiple single events averaged over a specific time period such as 1 day. LEQ is calculated as

$$L_{AeqT} = 10 \log_{10} \left(\frac{1}{t_{ref}} \sum_{i} 10^{\frac{SEL_i}{10}} \right) \tag{4.8}$$

Granoien et al [105], noting the relationship between LEQ and annoyance that exists for multiple events, proposed that each event makes a contribution to the total annoyance. Therefore, in [105] it was proposed that the annoyance experienced at a single point from a single event is proportional to the anti-log of the SEL (SE) for that event,

$$\text{Annoyance} \propto 10^{\frac{SEL}{10}} = SE \tag{4.9}$$

The proposed Annoyance Score is then the linear sum of the SE for each grid cell weighted by the population of that cell (4.10). Where the cell population is Pop_c and c is the grid cell index,

$$\text{Annoyance Score} = \sum_{c} Pop_c SE_c \tag{4.10}$$

While it is common to use single event SEL footprints to asses noise impacts from operational changes, there may be trade-offs between the different SEL noise contour levels [22, 69]. For instance, a reduction in the population within one SEL contour level, may only be realisable through population increases within other SEL contour levels. Therefore, there is a degree of subjective judgement as to the relative importance of different SEL noise contour levels when analysing changes in SEL footprints due to procedure changes.

There is however a well defined relationship between very high SEL levels $\geq$ 90 dB(A) and sleep disturbance [22, 69]. This relationship has allowed high SEL values to be consolidated into a single value, Awakenings metric, which provides a measure of the number of people woken at night due to a flight event. As single value metrics lend themselves well to inclusion in optimisation studies, prior noise abatement trajectory optimisation research has had a considerable focus on mitigating Awakenings [69, 106, 107, 108, 109].

Awakenings however provides a poor surrogate for overall community noise impact, as the metric was not designed to account for how different SEL noise levels contribute to the overall community noise impact [24]. The Annoyance Score on the other hand provides a single value, single event measure that was specifically designed to measure relative changes in community noise impact due to flight operational changes [109]. The Annoyance Score also overcomes the limitations of analysing SEL footprints by effectively using population numbers to provide an implicit, self adaptable weighting to the SEL noise values.

The noise Annoyance Score was therefore used in trajectory optimisation studies in this thesis to assess overall community noise impact. The application of the metric for this purpose has not been explored prior to work in this thesis.

4.3 Noise Modelling

For the calculation of civil aircraft noise impact local to airports, the dominant method used in the field is the Noise Power Distance (NPD) method [110, 111]. The NPD method utilises, for a number of noise metrics, tables of empirical data that relate the noise level calculated on the ground to the power utilised by the aircraft and the distance from the aircraft to the noise assessment point.

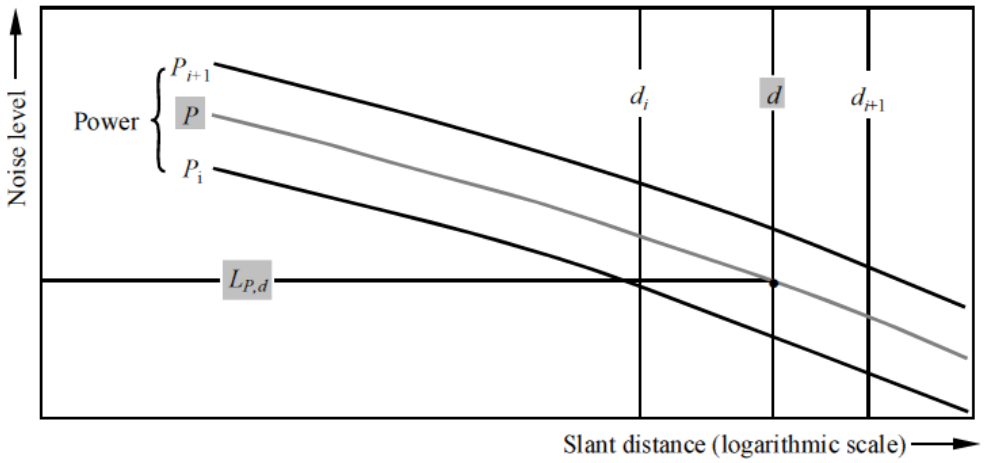

FIGURE 4.1: Noise Power Distance curves for aircraft noise calculation. Source, ECAC Doc 29 [110]

For the NPD approach, the flight trajectory is represented as a series of linear flight segments. The noise contribution of each segment is then considered relative to observer points. Figure 4.2 shows the segmented flight path and the geometrical relationships between the segments (S_1, S_2, S_3, S_4) and a single noise observation point.

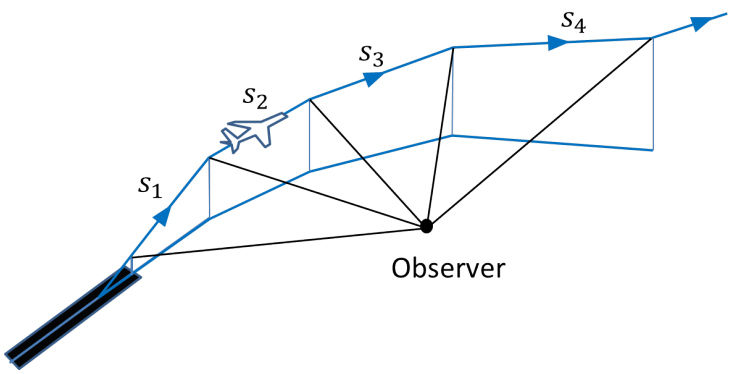

FIGURE 4.2: Segmented flight path for NPD noise calculation

Noise significant flight segments are determined relative to the observation point using a minimum noise level cut-off to exclude segments that will not make an appreciable change to the event noise level.

Where the metric is a maximum noise metric (L_{Amax}), noise values for all noise significant segments are assessed relative to the observation point, with the largest evaluated noise level set as the point noise value.

For noise exposure metrics, where the noise impact is duration dependant (SEL and EPNL), then the noise exposure value at the observation point is the decibel sum of the noise value contributions from each of the noise significant segments.

To evaluate the noise from a flight segment to any point, it is assumed that the aircraft is flying along a flight segment of infinite length. Figure 4.3 shows the flight segment S_2 treated as an infinite flight path segment where the observer point is alongside the segment and d_s and d_e are distances to the start and end of the flight path segment that define the segment geometry used in the evaluation of lateral attenuation.

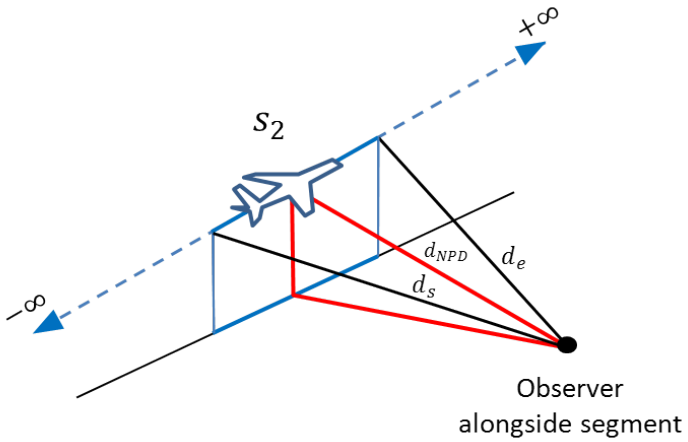

FIGURE 4.3: Noise calculation along an infinite segment

For maximum level noise, as shown in Figure 4.3, d_{NPD} is the slant distance between the observation point and the closest point on the flight path segment, where the aircraft power setting P_{NPD} is also evaluated. The distance d_{NPD} and the power P_{NPD} are then used to used to determine the baseline (infinite) noise level from the NPD curve. The actual (finite path) segment noise level value is determined by adjusting the baseline noise value for the effects of lateral directivity and attenuation. Lateral directivity defines the lateral radiation pattern of the sound, while lateral attenuation reduces the baseline noise value to reflect the reduction in noise experienced at the observer point due to ground absorption of the propagating noise.

For noise exposure, the slant distance d_{NPD} is the perpendicular distance between the observation point and the infinite flight path. Where the observer is alongside the segment, as in Figure 4.3, the slant distances are coincident for both maximum level and exposure type noise measures. Aircraft power P_{NPD} is evaluated at the point on the flight segment

closest to the observer point, where as before, P_{NPD} and d_{NPD} are used to determine the baseline noise value from the appropriate NPD curve. The actual (finite path) segment noise level values are then determined by applying adjustments to the baseline values.

For exposure metrics, in addition to lateral attenuation and directivity adjustments, the baseline noise is adjusted for finite segment length. Because the flight segments are of finite length, the sound energy radiated to the observer point from the segment is only a fraction of that radiated from the hypothetical infinite segment [112]. Lastly, the segment noise value is then further adjusted for differences between the flown flight speeds and the test condition flight speeds used for the derivation of the NPD data.

In summary then, for the NPD method, the noise level at a point for any noise metric is given as

$$\text{noise level} = f(P_{NPD}(\mathbf{x}, \mathbf{u}, t), d_{NPD}(\mathbf{x}, \mathbf{u}, t), \boldsymbol{\beta}), \ldots \forall t \in [t_o, t_f] \tag{4.11}$$

where P_{NPD} is power d_{NPD} is slant distance and $\boldsymbol{\beta}$ is a set of segment level correction terms. t_o and t_f are respectively the start and end times of the trajectory. Power, as used in the NPD method, is the corrected net thrust.

The NPD method noise model chosen for use in this work is the Integrated Noise Model (INM). INM is a model developed by the Federal Aviation Administration (FAA) to assess the impact of civil aircraft noise on communities local to airports. INM version 7 [111] is fully compatible with ECAC Doc 29 guidance [110], that provides a standardized methodology for the computation of noise impact around civil airports.

4.4 Summary: Aircraft Noise

In summary, there are a wide number of noise metrics available to the analyst when assessing changes to operational procedures. However, there may be trade-offs in how useful a metric is to an experienced analyst and how easily it is understood by the general public.

Therefore, it is thought here useful to consider the noise metrics introduced relative to their use in the case studies explored later in this work. In the Sourdine case study in Section 7.2, EPNL under the centreline was used as the metric of choice in the initial (baseline) study. Among other goals, the Sourdine study aimed to develop a series of recommended aircraft departure operating procedures that minimised noise impact on communities either near or more distant to the airport. As the procedures were designed as general guidance, they were not developed to a specific population. Without specific reference to a population, the EPNL metric, which includes an accounting of sound level, exposure time and is weighted for the human perception of noise, was taken as most

representative of actual noise impact. Therefore, it was inferred, that where reductions in EPNL values were achieved, that this would communicate real world noise benefits to communities surrounding airports. However, EPNL and the effects of changes in EPNL, are relatively complex to communicate to the public. Therefore, other single event metrics are often preferred when communicating changes in noise to those who will ultimately experience it.

In the Luton departure procedure optimisation case study in Section 8.3, two noise metrics were used, these were the peak noise L_{Amax} and the noise Annoyance Score metrics. In the initial study that formed the baseline for the scenario, no noise metrics were used. Instead, it was reasonably assumed that routing aircraft away from the largest population centres would result in reductions in noise for the most amount of people. In the IDVD-DE optimisation study, the peak noise measure, L_{Amax} metric, was measured at the population centres Milton Keynes and Leighton Buzzard. Seeking reductions in peak noise at those centres was considered to best represent the intent of the original work. However, for the optimisation study, the extra noise Annoyance Score metric was considered necessary. The Annoyance Score was chosen as the single event metric most representative of overall community noise impact. By including the Annoyance Score, it was intended that any trade-offs between noise at major populations and overall community noise impact would be represented on the Pareto front.

In the Luton case study in Section 8.3, when selecting the best balance Pareto solutions from the front, there was a general preference for solutions that reduced Annoyance Score, even at the expense of reductions in noise at specific sites. However, any decisions made were informed directly and transparantly from the Pareto based data in the study. This was considered to be a key advantage of such a Pareto based approach.

The Nottingham and East Midland (NEMA) case study in Section 8.2 looks at approach procedure optimisation. In the baseline study that formed the basis for the optimisation study, noise was assessed under the centreline of the flight. In this case, peak noise (L_{Amax}) under the centreline rather then EPNL was used. This is likely due to L_{Amax} being a more intuitive metric to understand.

One of the key findings of this case study was that, for the particular design of experiment detailed in Section 8.2.1, noise under the centreline was not a useful performance measure for guiding decision making. This was evidenced both in the distribution of the Pareto front solutions and in the difficulty that arose in relating data values to meaningful changes in community noise impact. Therefore, the IDVD-DE approach was shown not only to inform decision making on procedure optimisation, but also to inform decision making on metric choice and overall design of experiment.

Only single event noise metrics were used in this work. In Europe noise impact is governed through the use of the LEQ metric, which is applied to assess an airports overall noise impact. LEQ is discussed in this chapter, but is not used in this work due to it being a multi-event metric. However the relationship between the multi-event LEQ and single

event Annoyance Score metric is highlighted here. It is due to this relationship, and that the Annoyance Score allows single event noise impact to be defined as a single value, that the Annoyance Score metric is used extensively in the optimisation studies in this thesis.

Typically, arrival and departure procedures are developed individually, with design changes often initiated by non-environmental drivers such as capacity and safety. The trajectory and Pareto based approaches developed here are intended to support decision making typically required of these studies. However, by treating each procedure individually, noise inefficiencies due to interactions between procedures are not optimised for. Therefore, there is still a need for very careful analysis to ensure that what is achieved through optimising individual procedures is in fact an accumulative noise reduction when the noise impact of all procedures are considered.

Chapter 5

Trajectory Optimisation Methods

This chapter provides an overview of trajectory generation techniques and how they are used to convert the optimal control problem into a non-linear programming problem. It has been established in Sections 3.4 and 4.3 that the environmental impacts of flights can be determined as functions of the states and controls of the aircraft. Trajectory optimisation methods can then be applied to find the states and controls of flight trajectories that minimise environmental performance measures.

5.1 The Optimal Control Problem

The time history of the input controls to a dynamical system and the corresponding state histories of the system over the interval $[t_o, t_f]$ are denoted as $\mathbf{u}$ and $\mathbf{x}$ respectively.

The goal then of trajectory optimisation is to find the states $\mathbf{x}^*(t) \in \mathbb{R}^n$ and controls $\mathbf{u}^*(t) \in \mathbb{R}^n$ that minimise or maximise a measure of performance

$$J = \int_{t_0}^{t_f} L(\mathbf{x}(t), \mathbf{u}(t), t)dt + \varphi(\mathbf{x}_f, t_f) \tag{5.1}$$

where L is the running cost and φ is the terminal cost. The system dynamics are defined by a set of ordinary differential equations

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}, t) \tag{5.2}$$

with the prescribed initial conditions

$$\mathbf{x}(0) = \mathbf{x}_0 \tag{5.3}$$

and the terminal conditions

$$\boldsymbol{\psi}(\mathbf{u}_f, \mathbf{x}_f) = 0 \tag{5.4}$$

The multi-objective trajectory optimisation problem can be stated as the problem of minimising an array of scalar objective functionals

$$\min_{\mathbf{x},\mathbf{u}}[J_1(\mathbf{x},\mathbf{u}), J_2(\mathbf{x},\mathbf{u}), \ldots, J_k(\mathbf{x},\mathbf{u})]^T \tag{5.5}$$

A functional is a function that takes as its input another function and produces as an output a single real number. The individual objectives of the array are then minimised subject to the inequality constraints to be satisfied

$$\mathbf{c}(\mathbf{x},\mathbf{u}) \leq 0 \tag{5.6}$$

Because the optimisation free variables are continuous functions, optimal control problems have infinite dimensions. Generally to solve an optimal control problem, the infinite dimensional problem must be converted to a finite dimensional problem that can be solved with standard numerical methods. Betts classified two approaches for transforming the optimal control problem into a Nonlinear Programming (NLP) problem as direct and indirect methods. Indirect methods involve forming the Hamiltonian of the system, estimating the costate variables, and finding the root of the two point boundary value problem. Direct methods involve discretising the optimal control problem and solving for the states and controls at a series of dividing nodes.

5.2 Indirect Methods

Calculus of variations may be applied to optimisation problems to determine a function that minimises a given functional [113, 77, 114, 115]. Applying calculus of variations to the objective functional to find the first order optimality conditions, the augmented performance index becomes

$$\hat{J} = \varphi(\mathbf{x}_f, t_f) + \boldsymbol{\upsilon}^T\boldsymbol{\psi}(\mathbf{x}_f, t_f) + \int_{t_0}^{t_f}[L(\mathbf{x}(t),\mathbf{u}(t),t) + \boldsymbol{\lambda}^T(t)(\mathbf{f}(\mathbf{x}(t),\mathbf{u}(t)) - \dot{\mathbf{x}})]dt \tag{5.7}$$

where the differential constraints are adjoined to L by the adjoint variable vector $\boldsymbol{\lambda}^T$. Further Lagrange multipliers $\boldsymbol{\upsilon}^T$ are introduced for the boundary condition constraints $\boldsymbol{\psi}$. In Hamiltonian form the performance index becomes

$$H = L(\mathbf{x}(t),\mathbf{u}(t),t) + \boldsymbol{\lambda}^T(t)\mathbf{f}[\mathbf{x}(t),\mathbf{u}(t)] \tag{5.8}$$

where path constraints may be included by using additional Lagrange multipliers and adjoint equations. The first-order optimality conditions are the adjoint equations

$$\frac{\partial H}{\partial \mathbf{x}} = -\dot{\boldsymbol{\lambda}} \tag{5.9}$$

the transversality conditions

$$\boldsymbol{\lambda}(t_0) = 0 \tag{5.10}$$

$$\mathbf{x}(t_0) = \mathbf{x}_0 \tag{5.11}$$

$$\boldsymbol{\lambda}(t_f) = \left(\frac{\partial \varphi}{\partial \mathbf{x}} + \boldsymbol{\upsilon}^T \frac{\partial \boldsymbol{\psi}}{\partial \mathbf{x}}\right)^T_{t=t_f} \tag{5.12}$$

$$\left[\frac{\partial \varphi}{\partial t} + \boldsymbol{\upsilon}^T \frac{\partial \boldsymbol{\psi}}{\partial t} + H\right]_{t=t_f} = 0 \tag{5.13}$$

$$\boldsymbol{\psi}(t_f, \mathbf{x}_f) = 0 \tag{5.14}$$

and the stable condition

$$\frac{\partial H}{\partial \mathbf{u}} = 0 \tag{5.15}$$

For the global optimum, Pontryagins Maximum Principle is applied, which requires the Hamiltonian to be a minimum with respect to the control at every point of the trajectory. It is stated as

$$H(\mathbf{x}^\star, \mathbf{u}^\star, \boldsymbol{\lambda}^\star) \leq H(\mathbf{x}^\star, \mathbf{u}, \boldsymbol{\lambda}^\star), \quad t \in [t_0, t_f] \tag{5.16}$$

The optimal control problem can now be solved as a Boundary Value Problem (BVP) with numerical methods commonly applied to its solution [79]. The convergence of the solution is highly sensitive to the choice of the initial costate values, which are generally non-intuitive [79]. The problem is compounded when considering multiobjective optimisation. Therefore direct methods were considered more appropriate for this work due to their good convergence properties and ease of application to the multi-objective problem. Therefore indirect methods were not further pursued.

5.3 Direct Methods

For direct methods, the states and/or controls in the optimal control problem are parameterised using finite dimensional approximations to the continuous functions. Often polynomials interpolants are chosen to approximate the functions. The optimal control problem is then transformed to a NLP problem by discretising the approximations and solving for the states and controls at each of the dividing nodes. Typical methods of transcription include shooting, collocation, pseudospectral and inverse dynamics.

Time marching (Shooting, Multiple Shooting) typically involve parameterisation of the controls and the explicit integration of the state differential equations. The methods requires computationally expensive numerical integration and can be numerically unstable [116] . Their use is usually limited to problems with only a few control changes [117].

For collocation methods, time is discretised

$$t_0 < t_1 < \ldots < t_M = t_f \tag{5.17}$$

and both the states and controls are parameterised by polynomials such that the optimisation variables become the state and control values at the discretisation nodes

$$\mathbf{nlp} = [\mathbf{x}_0, \mathbf{u}_0, \mathbf{x}_1, \mathbf{u}_1, \ldots, \mathbf{x}_M, \mathbf{u}_M] \tag{5.18}$$

Local collocation typically uses piecewise polynomials derived from numerical integration schemes. The widely adopted K-stage Runge-Kutta method for discretising the differential equations is stated generally as

$$\mathbf{x}_{i+1} = \mathbf{x}_i + \mathrm{h}_i \sum_{j=1}^{K} \mathrm{b}_j \mathbf{f}_{ij} \tag{5.19}$$

where

$$\mathbf{f}_{ij} = \mathbf{f}\left[\mathbf{x}_{ij}, \mathbf{u}_{ij}, \boldsymbol{\rho}, t_{ij}\right] \tag{5.20}$$

and

$$\begin{aligned} \mathbf{x}_{ij} &= \mathbf{x}_i + \mathrm{h}_i \sum_{l=1}^{K} \mathrm{a}_{jl} \mathbf{f}_{il} \\ \mathbf{u}_{ij} &= \mathbf{u}(t_{ij}) \\ t_{ij} &= t_i + \mathrm{h}_i \rho_j \\ \rho_j &= (0, 1] \end{aligned} \tag{5.21}$$

where K defines the number of stages, determining the number of function evaluations, and h is the step size. The coefficients of the specific Runge Kutta method are then given by the weights $(\mathrm{b}_j)_{j=2,\ldots,K}$, the nodes $(\rho_j)_{j=1,\ldots,K}$ and the Runge Kutta matrix $(\mathrm{a}_{jl})_{1 \leq l < j \leq K}$

[118, 77]. The state and control trajectories are related through the systems dynamics, which are then enforced at each step through the residual constraint

$$\boldsymbol{\varsigma}_i = \mathbf{x}_{i+1} - \mathbf{x}_i - \mathrm{h}\sum_{j=1}^{K} \mathrm{b}_j \mathbf{f}_{ij} \tag{5.22}$$

Collocation schemes exist in two forms, local and global [117]. Local collocation schemes vary the number of piecewise polynomials to better approximate the function, Global collocation, such as pseudospectral methods, use global polynomials, where the degree of the polynomial is varied. For pseudospectral methods the states and controls are generally parameterised by Lagrange polynomials over the normalised time interval $\boldsymbol{\tau} \in [-1, 1]$

$$\begin{aligned} \mathbf{x}(\boldsymbol{\tau}) \approx \mathbf{X}(\boldsymbol{\tau}) &= \sum_{k=1}^{N_p} \mathbf{x}(\boldsymbol{\tau}_k)\mathcal{L}_k(\boldsymbol{\tau}) \\ \mathbf{u}(\boldsymbol{\tau}) \approx \mathbf{U}(\boldsymbol{\tau}) &= \sum_{k=1}^{N_p} \mathbf{u}(\boldsymbol{\tau}_k)\mathcal{L}_k(\boldsymbol{\tau}) \end{aligned} \tag{5.23}$$

where N_p is the order of the polynomial and $\mathcal{L}_k(\boldsymbol{\tau})_{k=1,\ldots,N_p}$ are the Lagrange basis functions [119].

The orthogonal polynomials may be discretised using nodes obtained from a Guassian quadrature. The collocation points are then the roots of the orthogonal polynomial [79]. Choosing the nodes in this manner maximises the accuracy of the quadrature interpolation. The 3 types of collocation points commonly applied to the global Legrande or Chebyschev polynomials over the interval $[-1, 1]$ are the Gauss, Gauss-Radau and Gauss-Lobatto points. As with the local collocation schemes, to transform the optimal control problem to an NLP problem, the differential algebraic equations are similarly collocated at the nodes using residual constraints

$$\boldsymbol{\zeta}_k = \sum_{i=0}^{N_p} \mathcal{D}_{ki}\mathbf{X}(\boldsymbol{\tau}_i) - \frac{t_f - t_0}{2}\mathbf{f}\left(\mathbf{X}(\boldsymbol{\tau}_k), \mathbf{U}(\boldsymbol{\tau}_k), \boldsymbol{\tau}_k\right) = 0 \tag{5.24}$$

where $\mathcal{D}_{ki} = \dot{\mathcal{L}}_i(\boldsymbol{\tau}_k)$ is the derivative of the ith Lagrange polynomial with respect to the normalized time at $\boldsymbol{\tau}_k$ [119, 120].

Yakimenko [121] proposed an inverse method where the position states of the aircraft and their derivatives are parameterised using 7th degree polynomials. Controls are then determined by inverting the state equations. The method significantly reduces the number of optimisation variables required by analytically determining the polynomial coefficients from the prescribed states and controls at the boundaries ($t = 0$ and $t = t_f$). Instead of parameterising by time, Yakimenko adopted Taranenkos method [122] of parameterising the polynomials by τ, creating a virtual arc $\tau \in [\tau_0, \tau_f]$. The relationship between time

t and τ is defined as $\lambda = d\tau/dt$. The use of the relationship parameter τ allows the definition of aircraft velocity using a separate reference function, enabling the creation of a virtual speed profile along the trajectory path of the aircraft. The method, termed the Inverse Dynamics in the Virtual Domain (IDVD) method, is a fast trajectory optimisation method and has been considered for real time implementation [78].

The IDVD method was favoured for this work due its compatibility with the multiobjective Differential Evolution (DE) NLP solver chosen for this work. DE uses a computationally intensive trial and error approach to finding the Pareto front. Collocation schemes, which have large numbers of optimisation parameters act to further increase the computational complexity of the problem. The inverse method on the other hand, when applied with a 3 Degree of Freedom (DOF) aircraft dynamics model, results in a system with 6 aircraft states, 3 controls and 9 optimisation variables. This allows the combination of the IDVD and DE methods to quickly progress solutions from infeasible to feasible to a point on the global Pareto front. The principle drawback of the approach is the accuracy offered by the global interpolating polynomials, which are only of degree 7. This is considerably less than what may be available from adopting collocation schemes. However, as commercial aircraft have relatively unagressive trajectories and as the environmental models adopted are sensitive to macro rather than micro changes in the trajectory, the low parameter space offered by the IDVD method was considered more beneficial than the higher fidelity potentially offered by collocation schemes.

5.4 Inverse Dynamics

Dynamical systems where all the states and inputs are expressed as functions of the output variables and their derivatives are referred to as differentially flat [123, 80, 124]. The key idea of the inverse dynamics method is to parameterise a set of output trajectories that, through the differentially-flat property of the dynamics system, fully define all states and inputs in terms of those outputs and their derivatives.

For the inverse dynamics method, the state and control vectors are expressed as functions of the output trajectory vector $\mathbf{r} = [r_1(t), r_2(t), r_3(t)]$ and its derivatives, where $r_1(t), r_2(t), r_3(t)$ are the Cartesian x, y, z coordinates $\forall t \in [t_0, t_f]$. Yakimenko recommended parameterising the flat earth Cartesian coordinates and their derivatives using 7th order polynomials, where the trajectory is generated in the output space and the input controls and the remaining states are then determined algebraically. For the inverse method, the Cartesian positional states r_j $(j = 1, 2, 3)$ of the aircraft and their derivatives are parameterised by the following reference function and its derivatives:

$$r_j(\tau) = \sum_{k=0}^{7} \frac{a_{jk}\tau^k}{\max(1, k(k-1))} \tag{5.25}$$

Determining the coefficients of the polynomials analytically requires that the degree of the polynomial be defined by the number of boundary conditions to be satisfied. This results in the number of equations equalling the number of unknowns such that an equation can be defined for each of the unknowns. The minimum degree of the polynomials is $n = d_0 + d_f + 1$, which is greater by one than the sum of the maximum orders of the time derivatives of the aircraft position at the boundaries. The flexibility of the functions can be increased in two ways, the degree of the polynomial can be increased by adding extra boundary conditions and some of the boundary conditions can be turned into optimisation variables.

The reference functions defined at the boundaries $[\tau_0, \tau_f]$ lead to the following system of linear equations [78],

$$\mathbf{b}_j = \mathbf{C}\mathbf{a}_j, \;\; j = 1, 2, 3 \tag{5.26}$$

where $\mathbf{C}$ is a m by n matrix, $\mathbf{a}$ is a vector of polynomial coefficients and $\mathbf{b}$ is a vector of initial and final boundary conditions.

$$\mathbf{b}_j = \begin{bmatrix} r_{j0} \\ r'_{j0} \\ r''_{j0} \\ r'''_{j0} \\ r_{jf} \\ r'_{jf} \\ r''_{jf} \\ r'''_{jf} \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & \tau_f & \frac{\tau_f^2}{2} & \frac{\tau_f^3}{6} & \frac{\tau_f^4}{12} & \frac{\tau_f^5}{20} & \frac{\tau_f^6}{30} & \frac{\tau_f^7}{42} \\ 0 & 1 & \tau_f & \frac{\tau_f^2}{2} & \frac{\tau_f^3}{3} & \frac{\tau_f^4}{4} & \frac{\tau_f^5}{5} & \frac{\tau_f^6}{6} \\ 0 & 0 & 1 & \tau_f & \tau_f^2 & \tau_f^3 & \tau_f^4 & \tau_f^5 \\ 0 & 0 & 0 & 1 & 2\tau & 3\tau^2 & 4\tau^3 & 5\tau^4 \end{bmatrix}, \mathbf{a}_j = \begin{bmatrix} a_{j0} \\ a_{j1} \\ a_{j2} \\ a_{j3} \\ a_{j4} \\ a_{j5} \\ a_{j6} \\ a_{j7} \end{bmatrix} \tag{5.27}$$

The coefficients of the reference polynomials are then determined analytically by inverting the matrix $\mathbf{C}$ and making $\mathbf{a}$ the subject of the linear equations $\mathbf{a}_j = \mathbf{C}^{-1}\mathbf{b}_j,\ j = 1, 2, 3$ such that

$$\begin{aligned} a_{j0} = r_{j0}, \quad a_{j1} = r'_{j0}, \quad a_{j2} = r''_{j0}, \quad a_{j3} = r'''_{j0} \\ a_{j4} = \frac{2x'''_{jf} + 8x'''_{i0}}{\tau_f} + \frac{30x''_{jf} - 60x''_{j0}}{\tau_f^2} - \frac{180x'_{jf} + 240x'_{j0}}{\tau_f^3} + 420\frac{x_{jf} - x_{j0}}{\tau_f^4} \\ a_{j5} = \frac{10x'''_{jf} + 20x'''_{j0}}{\tau_f^2} + \frac{140x''_{jf} - 200x''_{j0}}{\tau_f^3} - \frac{780x'_{jf} + 900x'_{j0}}{\tau_f^4} - 1680\frac{x_{jf} - x_{i0}}{\tau_f^5} \\ a_{j6} = \frac{15x'''_{jf} + 20x'''_{j0}}{\tau_f^3} + \frac{195x''_{jf} - 225x''_{j0}}{\tau_f^4} - \frac{1020x'_{jf} + 1080x'_{j0}}{\tau_f^5} + 2100\frac{x_{jf} - x_{i0}}{\tau_f^6} \\ a_{j7} = 7\frac{x'''_{jf} + x'''_{j0}}{\tau_f^4} - 84\frac{x''_{jf} - x''_{j0}}{\tau_f^5} + 420\frac{x'_{jf} + x'_{j0}}{\tau_f^6} - 840\frac{x_{jf} - x_{j0}}{\tau_f^7} \end{aligned} \tag{5.28}$$

As with all direct methods, discretisation allows the functions to be considered as a finite set of variables. For the inverse method, Yakimenko used a node distribution that was evenly spaced over τ. However, Drury recommended in [125] a Chebyshev-Gauss-Lobatto node distribution as it avoided the ill-conditioning that occurs with high-degree polynomial interpolation using uniformly-spaced nodes. CGL interpolation points exist over the interval $[-1, 1]$ such that $\boldsymbol{\tau} \in [\tau_0, \tau_f] = [-1, 1]$ The interpolation points are generated as

$$\boldsymbol{\tau} = \cos\left(\frac{\pi l}{N}\right), \quad l = 0, \ldots, N \tag{5.29}$$

and the following transformation is used to convert to τ

$$\tau = \frac{(\tau_f - \tau_0)\boldsymbol{\tau} + (\tau_f + \tau_0)}{2} \tag{5.30}$$

A defining feature of the inverse method is the parametisation of the polynomials over the virtual arc τ. Parameterising by time explicitly links path and speed along the path.

Parameterising by τ allows the path optimisation to be separated from the speed profile. The speed profile may be predetermined or alternatively be defined by the reference function (5.25) producing a second set of algebraic equations where $v_{\mathtt{t}}$ is the aircraft true airspeed

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & \tau_f & \frac{\tau_f^2}{2} & \frac{\tau_f^3}{6} & \frac{\tau_f^4}{12} & \frac{\tau_f^5}{20} \\ 0 & 1 & \tau_f & \frac{\tau_f^2}{2} & \frac{\tau_f^3}{3} & \frac{\tau_f^4}{4} \\ 0 & 0 & 1 & \tau_f & \tau_f^2 & \tau_f^3 \\ 0 & 0 & 0 & 1 & 2\tau & 3\tau^2 \end{bmatrix} \begin{bmatrix} b_0 \\ b_1 \\ b_2 \\ b_3 \\ b_4 \\ b_5 \end{bmatrix} = \begin{bmatrix} v_{\mathtt{t}0} \\ v'_{j0} \\ v''_0 \\ v_{\mathtt{t}f} \\ v'_f \\ v''_f \end{bmatrix} \tag{5.31}$$

Defining a separate speed profile in this manner allows some of the higher order velocity components at the boundaries to be turned into optimisation variables. This adds an extra degree of freedom to the optimization and allows the speed profiles to be optimised along the trajectory path of the aircraft. The relationship between the true airspeed $v_{\mathtt{t}}$ and the speed along the virtual arc $\sqrt{x'^2 + y'^2 + h'^2}$ is then defined as

$$v_{\mathtt{t}} = \lambda\sqrt{x'^2 + y'^2 + h'^2} \tag{5.32}$$

where λ is the scale or speed factor. It follows then that

$$\lambda = \frac{v_{\mathtt{t}}}{\sqrt{x'^2 + y'^2 + h'^2}} = \frac{\dot{s}_{\mathtt{a}}(\tau)}{s'(\tau)} \tag{5.33}$$

where $\dot{s}_{\mathtt{a}}(\tau)$ is the rate of change of distance in the airmass frame parameterised by τ and $s'(\tau)$ is the rate of change of virtual arc parameterised by τ. Where values are assessed at each node, the parameterised true airspeed $v_{\mathtt{t}}$ is greater than then virtual arc speed $s'(\tau)$ when $\lambda > 1$. The virtual arc speed is less than $v_{\mathtt{t}}$ when $\lambda < 1$, and both the speed profiles are identical where $\lambda = 1$.

Higher derivatives of λ are calculated using the product and quotient rules where

$$\begin{aligned} \lambda' &= \frac{\dot{s}'_{\mathtt{a}} s' - s''\dot{s}_{\mathtt{a}}}{s'^2} \\ \lambda'' &= \frac{v''}{s'} - \frac{2s''\dot{s}'_{\mathtt{a}}}{s'^2} + \dot{s}_{\mathtt{a}}\left(\frac{2s''^2}{s'^3} - \frac{s'''}{s'^2}\right) \end{aligned} \tag{5.34}$$

and where

$$\begin{aligned} s'' &= \frac{2x'x'' + 2y'y'' + 2h'h''}{2\sqrt{x'^2 + y'^2 + h'^2}} \\ s''' &= \frac{(2x''^2 + 2h'''h' + 2x''^2 + 2x'''x' + 2y''^2 + 2y'''y')}{(2\sqrt{h'^2 + x'^2 + x'^2})} - \frac{(2h'h'' + 2x'x'' + 2y'y'')^2}{(4(h'^2 + x'^2 + y'^2))^{3/2}} \end{aligned} \tag{5.35}$$

Conversions between the virtual and time domain are then achieved by

$$\begin{aligned} \dot{r} &= \lambda r' \\ \ddot{r} &= \lambda(r''\lambda + r'\lambda') \\ \dddot{r} &= \lambda^3 r''' + 3\lambda^2\lambda' r'' + (\lambda^2 + \lambda\lambda'^2)r' \end{aligned} \tag{5.36}$$

The vector of true airspeed components is then $\dot{\mathbf{r}}(t) = \mathbf{v}_{\mathtt{a}}(t) = [\dot{x}_{\mathtt{a}}(t), \dot{y}_{\mathtt{a}}(t), \dot{h}_{\mathtt{a}}(t)]$ such that the ground speed $v_{\mathtt{g}}$ is given as $v_{\mathtt{g}} = \sqrt{(\dot{x}_{\mathtt{a}} + w_x)^2 + (\dot{y}_{\mathtt{a}} + w_y)^2}$, and where the wind speeds in the 3 coordinate directions are w_x, w_y and w_h. The subscripts $_{\mathtt{a}}$ and $_{\mathtt{i}}$ designate the air mass and the inertial frames respectively. The air mass frame is a frame that is aligned with the inertial frame but moves at a constant velocity with respect to the inertial frame. The inclusion of the air mass frame allows aircraft motion to be considered relative to both the earth and the moving air mass. As in [126], the wind's impact on velocity and path are considered but not its impact on acceleration. The relationships between the speeds in the inertial and air mass frames are described by

$$\begin{aligned} \dot{x}_{\mathtt{a}} &= \dot{x}_{\mathtt{i}} - w_x \\ \dot{y}_{\mathtt{a}} &= \dot{y}_{\mathtt{i}} - w_y \\ \dot{h}_{\mathtt{a}} &= \dot{h}_{\mathtt{i}} - w_h \\ \mathbf{v}_{\mathtt{a}} &= \mathbf{v}_{\mathtt{i}} - \mathbf{v}_w \end{aligned} \tag{5.37}$$

The time t is calculated from the following relationship

$$t = \int \frac{1}{\lambda} d\tau \tag{5.38}$$

To modify the polynomials, the variables iterated by the solver are then the initial and final jerks and the final tau, $\Xi = [x'''_{0,f}, y'''_{0,f}, h'''_{0,f}, v''_{0,f}, \tau_f]$. To transform the polynomials to the system dynamics, a point mass model is used. Therefore the state and controls are

determined by inverting the following equations

$$\begin{aligned} \dot{x}_{\mathtt{i}} &= v_{\mathtt{t}} \cos\gamma_{\mathtt{a}} \cos\chi_{\mathtt{a}} + w_x, & \dot{v}_{\mathtt{t}} &= \frac{T-D}{m} - g\sin\gamma_{\mathtt{a}} \\ \dot{y}_{\mathtt{i}} &= v_{\mathtt{t}} \sin\chi_{\mathtt{a}} \cos\gamma_{\mathtt{a}} + w_y, & \dot{\chi}_{\mathtt{a}} &= \frac{gn\sin\phi}{v_{\mathtt{t}}\cos\gamma_{\mathtt{a}}} \\ \dot{h}_{\mathtt{i}} &= v_{\mathtt{t}} \sin\gamma_{\mathtt{a}} + w_h, & \dot{\gamma}_{\mathtt{a}} &= \frac{g}{v_{\mathtt{t}}}(n\cos\phi - \cos\gamma_{\mathtt{a}}) \end{aligned} \tag{5.39}$$

such that the remaining states are given by

$$\gamma_{\mathtt{a}} = \sin^{-1}\left(\frac{\dot{h}_{\mathtt{a}}}{v_{\mathtt{t}}}\right), \qquad \chi_{\mathtt{a}} = \operatorname{atan2}\left(\dot{y}_{\mathtt{a}}, \dot{x}_{\mathtt{a}}\right) \tag{5.40}$$

$$\gamma_{\mathtt{i}} = \sin^{-1}\left(\frac{\dot{h}_{\mathtt{i}}}{v_{\mathtt{t}}}\right), \qquad \chi_{\mathtt{i}} = \operatorname{atan2}\left(\dot{y}_{\mathtt{i}}, \dot{x}_{\mathtt{i}}\right) \tag{5.41}$$

which are visualised in Figure 5.1, where $\mathbf{v}_{\mathtt{a}/x_{\mathtt{i}}y_{\mathtt{i}}}$ is the projection of the true airspeed vector onto the local x, y plane [127].

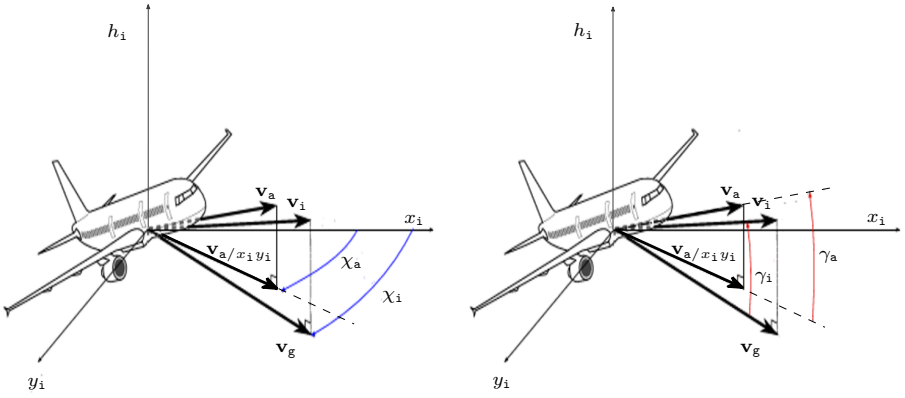

FIGURE 5.1: Dynamics model: χ and γ angles, with $x_{\mathtt{i}}$ pointing to the North

The controls are then given as

$$T = D + m\dot{v}_{\mathtt{t}} + mg\sin(\gamma_{\mathtt{a}}), \qquad n = \frac{\sqrt{(v_{\mathtt{t}}\dot{\gamma}_{\mathtt{a}} + g\cos(\gamma_{\mathtt{a}}))^2 + (v_{\mathtt{t}}\dot{\chi}_{\mathtt{a}}\cos(\gamma_{\mathtt{a}}))^2}}{g}$$
$$\phi = \tan^{-1}\left(\frac{\dot{\chi}_{\mathtt{a}} v_{\mathtt{t}}\cos(\gamma_{\mathtt{a}})}{g\cos(\gamma_{\mathtt{a}}) + v_{\mathtt{t}}\dot{\gamma}_{\mathtt{a}}}\right) \tag{5.42}$$

where

$$\begin{aligned}\dot{\chi}_{\mathtt{a}} &= \frac{(\dot{x}_{\mathtt{a}}\ddot{y} - \dot{y}_{\mathtt{a}}\ddot{x})}{(\dot{y}_{\mathtt{a}}^2 + \dot{x}_{\mathtt{a}}^2)} \\ \dot{\gamma}_{\mathtt{a}} &= \frac{\ddot{h}}{\sqrt{(\dot{x}_{\mathtt{a}}^2 + \dot{y}_{\mathtt{a}}^2)}} - \frac{\dot{h}_{\mathtt{a}}(\dot{x}_{\mathtt{a}}\ddot{x} + \dot{y}_{\mathtt{a}}\ddot{y} + \dot{h}_{\mathtt{a}}\ddot{h})}{\sqrt{(\dot{x}_{\mathtt{a}}^2 + \dot{y}_{\mathtt{a}}^2)}(\dot{x}_{\mathtt{a}}^2 + \dot{y}_{\mathtt{a}}^2 + \dot{h}_{\mathtt{a}}^2)}\end{aligned} \tag{5.43}$$

The flight path angle is $\gamma(t)$ and $\chi(t)$ is the heading angle. The subscripts $_{\mathtt{t,g,a,i}}$ designate true air speed, ground speed, air mass frame and inertial frame respectively. The aircraft controls are $\mathbf{u}(t) = [T(t), n(t), \phi(t)]^T$, where $T(t)$ is thrust, $n(t)$ is load factor and $\phi(t)$ is the bank angle. The drag D is modelled with the aid of the BADA drag polars [128], aircraft mass is m, g is gravitational acceleration.

Therefore, through inverse dynamics, the state and control vectors are expressed as functions of the output trajectory vector $\mathbf{r}$ and its derivatives such that

$$\mathbf{x} = \mathbf{f}_x(\mathbf{r}, \dot{\mathbf{r}}, \ddot{\mathbf{r}} \ldots), \qquad \mathbf{u} = \mathbf{f}_u(\mathbf{r}, \dot{\mathbf{r}}, \ddot{\mathbf{r}} \ldots) \tag{5.44}$$

Once the state and control histories are determined, constraints are applied to ensure that values lie between defined limits

$$\begin{aligned} r_1 \in [x_{min}; x_{max}], \quad & r_2 \in [y_{min}; y_{max}], \quad & r_3 \in [h_{min}; h_{max}] \\ T \in [T_{max}; T_{min}], \quad & n \in [n_{max}; n_{min}], \quad & |\phi| \leq |\phi_{max}| \\ v_{CAS} \in [v_{max_{CAS}}; v_{min_{CAS}}], \quad & v_{\mathtt{t}} \in [v_{\mathtt{t}_{max}}; v_{\mathtt{t}_{min}}], \quad & |a_l| \leq |a_{l_{max}}| \\ |a_n| \leq |a_{n_{max}}|, \quad & \gamma_{\mathtt{a}} \in [\gamma_{\mathtt{a}_{max}}; \gamma_{\mathtt{a}_{min}}], \quad & \gamma_{\mathtt{i}} \in [\gamma_{\mathtt{i}_{max}}; \gamma_{\mathtt{i}_{min}}] \\ \chi_{\mathtt{a}} \in [\chi_{\mathtt{a}_{max}}; \chi_{\mathtt{a}_{min}}], \quad & \chi_{\mathtt{i}} \in [\chi_{\mathtt{i}_{max}}; \chi_{\mathtt{i}_{min}}] \end{aligned} \tag{5.45}$$

where the positional and angle constraints on x, y, h, γ and χ are user defined and scenario specific. The constraints on roll angle $|\phi_{max}|$, longitudinal acceleration $|a_{l_{max}}|$ and normal acceleration $|a_{n_{max}}|$ are defined by the BADA dynamics model as

$$\begin{aligned} |\phi| &\leq 0.436 \quad \text{rad for takeoff and landing} \\ |\phi| &\leq 0.785 \quad \text{rad for all other phases} \\ |a_l| &\leq 0.6096 \quad \text{ms}^2 \text{ longitudinal acceleration} \\ |a_n| &\leq 1.524 \quad \text{ms}^2 \text{ normal acceleration} \end{aligned} \tag{5.46}$$

where

$$\begin{aligned} a_l &= \dot{v}_{\mathtt{t}} \\ a_n &= \dot{\gamma}_{\mathtt{a}} v_{\mathtt{t}} \end{aligned} \tag{5.47}$$

The minimum speed constraint, defined in Calibrated Air Speed (CAS) $v_{min_{CAS}}$ and the thrust constraint $T_{min;max}$ are determined from BADA functions as described in Appendix B.

For implementation, where the trajectory variables were subject to static higher and lower bounds at each node, the constraints were defined in the following linear algebraic form

$$\widehat{\mathbf{c}}_1 = \widehat{\mathbf{A}}\widehat{\mathbf{D}} + \widehat{\mathbf{b}} \tag{5.48}$$

where $\widehat{\mathbf{D}}$ was the matrix of trajectory variable values at each discretisation node

$$\widehat{\mathbf{D}} = [\mathbf{r}_1, \mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_2, \mathbf{r}_3, \mathbf{r}_3, \mathbf{v}_\mathtt{t}, \mathbf{v}_\mathtt{t}, \mathbf{a}_l, \mathbf{a}_l, \mathbf{a}_n, \mathbf{a}_n, \mathbf{n}, \mathbf{n}, \boldsymbol{\gamma}_\mathtt{a}, \boldsymbol{\gamma}_\mathtt{a}, \boldsymbol{\gamma}_\mathtt{i}, \boldsymbol{\gamma}_\mathtt{i}, \boldsymbol{\chi}_\mathtt{a}, \boldsymbol{\chi}_\mathtt{a}, \boldsymbol{\chi}_\mathtt{i}, \boldsymbol{\chi}_\mathtt{i}, \boldsymbol{\phi}, \boldsymbol{\phi}]^T \tag{5.49}$$

and where, prior to simulation initialisation, the appropriate constraint values from (5.45) were replicated $(1 \times N)$ times to create the rows of the constraint bounds matrix $\widehat{\mathbf{b}}$ such that

$$\begin{aligned}\widehat{\mathbf{b}} = [&\mathbf{x}_{min}, \mathbf{x}_{max}, \mathbf{y}_{min}, \mathbf{y}_{max}, \mathbf{h}_{min}, \mathbf{h}_{max}, \mathbf{v}_{\mathtt{t}_{min}}, \mathbf{v}_{\mathtt{t}_{max}}, \mathbf{a}_{l_{min}}, \mathbf{a}_{l_{max}}, \mathbf{a}_{n_{min}}, \mathbf{a}_{n_{max}}, \\ &\mathbf{n}_{min}, \mathbf{n}_{max}, \boldsymbol{\gamma}_{\mathtt{a}_{min}}, \boldsymbol{\gamma}_{\mathtt{a}_{max}}, \boldsymbol{\gamma}_{\mathtt{i}_{min}}, \boldsymbol{\gamma}_{\mathtt{i}_{max}}, \boldsymbol{\chi}_{\mathtt{a}_{min}}, \boldsymbol{\chi}_{\mathtt{a}_{max}}, \boldsymbol{\chi}_{\mathtt{i}_{min}}, \boldsymbol{\chi}_{\mathtt{i}_{max}}, \boldsymbol{\phi}_{min}, \boldsymbol{\phi}_{max}]^T\end{aligned} \tag{5.50}$$

where N is the number of nodes.

Because not all constraints were defined for every simulation run, the diagonal activation matrix $\widehat{\mathbf{A}}$ was defined prior to simulation initialisation and used during simulations to activate or cancel constraints

$$\widehat{\mathbf{A}} = \begin{bmatrix} \widehat{a}_1 & 0 & \cdots & 0 \\ 0 & \widehat{a}_2 & \cdots & 0 \\ 0 & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \widehat{a}_n \end{bmatrix} \tag{5.51}$$

The values on the diagonal for $i = 1 \ldots n$ where $n = \text{rows}(\widehat{\mathbf{b}})$ are

$$\widehat{a}_i = \begin{cases} 1 & : \text{if less than constraint} \\ -1 & : \text{if greater than constraint} \\ 0 & : \text{if no constraint is applied} \end{cases} \tag{5.52}$$

that is, 1 if a less than constraint was applied, -1 if a greater than constraint was applied or 0 if there was no constraint applied to the corresponding trajectory variable.

5.4.1 Path Constraints

For commercial aircraft trajectory optimisation, constraints imposed by the operating environment must be considered. Calculated trajectories must be able to adhere to Air Traffic Control (ATC) constraints imposed by airspace sectorisation, procedures and traffic flow corridors. Typically, operating environment and procedural restrictions manifest as constraints on the height, speed or path of the flight, or some combination of the three. Therefore a simple five dimensional constraints model was developed.

Waypoint fixes are defined by the user in the two horizontal dimensions as $\mathbf{W}_r = [W_x, W_y]$. The aircraft's trajectory path $\mathbf{r}_{ac_{2D}}(t) = [x(t), y(t)]$ can then be constrained to fly over the 2D fix. For each waypoint fix, the minimum distance between the trajectory path and the fix position is calculated as a 2D horizontal distance d_{min}, with a corresponding minimum time $t_{d_{min}}$,

$$d_{min} := \min_{t\in[t_0,t_f]} d(t) \;\text{ where }\; d(t) = \|\mathbf{r}_{ac_{2D}}(t) - \mathbf{W}_r\| \tag{5.53}$$

$$t_{d_{min}} := \min t \;\; s.t. \;\; d(t) = d_{min} \tag{5.54}$$

The aircraft is then constrained to fly within a distance radius of the centre point of the fix, where $\hat{d}$ is the upper constraint on d_{min},

$$\mathbf{c}_2(d_{min}) = d_{min} - \hat{d} \tag{5.55}$$

The aircraft can also be constrained to cross the fix at a specified height, speed and arrival time. The cross above constraints, $\underline{h}, \underline{v}, \underline{t}$, and cross below constraints, $\overline{h}, \overline{v}, \overline{t}$, constrain the minimum and maximum heights, speeds and time of the aircraft crossing the waypoint. The minimum and maximum can be constrained simultaneously to create height, speed and time windows at the waypoint.

$$\mathbf{c}_3(\mathbf{h}, \mathbf{v}_\mathbf{t}, \mathbf{t}, t_{d_{min}}) = \begin{bmatrix} h(t_{d_{min}}) - \overline{h}_{t_{min}} \\ \underline{h}_{t_{min}} - h(t_{d_{min}}) \\ v_\mathbf{t}(t_{d_{min}}) - \overline{v}_{t_{min}} \\ \underline{v}_{t_{min}} - v_\mathbf{t}(t_{d_{min}}) \\ t_{d_{min}} - \overline{t}_{min} \\ \underline{t}_{min} - t_{d_{min}} \end{bmatrix} \tag{5.56}$$

For implementation, a similar activation matrix approach as in (5.51) was used to to account for the fact that a waypoint can have multiple combinations of height, speed and time constraints active. As the constraints that are active may vary for each waypoint, the use of a pre-processed activation matrix allows the flexible inclusion or exclusion of constraints without needing to code multiple *if* statements.

Root Finding Method for Waypoint Constraints

The output trajectory vector of the inverse method provides the path coordinates at each discretisation node. Using the node constraint method described above, the node positions can easily be compared to the waypoint position to determine if the node and therefore a trajectory segment lies within the waypoint radius. However, this approach relies on a node occuring within a waypoint radius, which may not be the case, and it is possible that a trajectory path satisfies a waypoint crossing constraint without a node occuring in the waypoint radius. This can occur if the discretisation is not finite enough or if the local node distribution is sparse. An alternative approach that may be used to avoid the issue is to use root finding to determine if a trajectory crosses a waypoint fix.

The distance from any point on the trajectory path to the waypoint perimeter is defined as

$$d(t) = \sqrt{(x - W_x)^2 + (y - W_y)^2} - W_r \tag{5.57}$$

where W_r is the waypoint radius and x and y are the 7th degree position polynomials in the form

$$\begin{aligned} x &= a_7\frac{\tau^7}{42} + a_6\frac{\tau^6}{30} + a_5\frac{\tau^5}{20} + a_4\frac{\tau^4}{12} + a_3\frac{\tau^3}{6} + a_2\frac{\tau^2}{2} + a_1\tau + a_0 \\ y &= b_7\frac{\tau^7}{42} + b_6\frac{\tau^6}{30} + b_5\frac{\tau^5}{20} + b_4\frac{\tau^4}{12} + b_3\frac{\tau^3}{6} + b_2\frac{\tau^2}{2} + b_1\tau + b_0 \end{aligned} \tag{5.58}$$

When the trajectory path passes through the waypoint, the intersection of the polynomal and the waypoint perimeter occurs at $d(t) = 0$. Using squared distance to remove the square root in (5.57), points of intersection between the trajectory path and the waypoint occur where

$$d^2(t) = (x - W_x)^2 + (y - W_y)^2 - W_r^2 = 0 \tag{5.59}$$

Expression (5.59), may also be used as an equality constraint to require the aircraft path to cross the waypoint. To determine if the constraint is satisfied, it is only necessary to determine if (5.59), which factors out to be a 14th degree polynomial, has real roots. For waypoint following problems, this is sufficient to implement a Mixed Integer Linear Programming (MILP) approach, which allows the incorporation of logical statements as part of the trajectory optimisation formulation [129]. Here, the waypoint crossing constraint is defined such that it has has a binary solution, 1 if satisfied and is 0 if not. MILP based trajectory optimisation methods can be found in [129], [130].

Determining if a polynomial has real roots can be achieved analytically through Sturm's theorem [131]. However, if height, speed and time constraints are to be satisfied at the waypoint, or if the value of the constraint violation is desired, the intersection times between the trajectory and the waypoint must be determined. These can be determined, if they exist, from the real roots of equation (5.59). If the roots of the constraint are a complex number, then the real part of the complex number provides the time of the Closest Point of Approach (CPA) of the polynomial and the waypoint. From the time of

CPA, the CPA distance is easily determined and this can be used by the NLP method to reduce the constraint violation until the constraint is satisfied.

The choice of the constraints method for waypoint crossing constraints is largely scenario dependant. For the node constraint method, node numbers can be increased to ensure sufficient density. However increasing the number of nodes, increases the computational burden and time. Over larger flight distances with sparser node and waypoint distribution, it may be more efficient to use the root finding approach.

Corridor Constraint

In the current ATM system, aircraft are frequently required to fly along airways defined as airspace corridors. The corridor constraint from [132] was therefore adopted. Where for a straight line corridor, the start of the corridor is defined as $\mathbf{r}_{c_{start}}$ and the end of the corridor is defined by $\mathbf{r}_{c_{end}}$, the unit vector between the corridor start and end is $\tilde{\mathbf{p}}$. The perpendicular distance vector $\mathbf{d}_{\perp}(t)$ between the aircraft trajectory $\mathbf{r}_{ac}$ and the centreline of the corridor is defined as

$$\mathbf{d}_{\perp}(t) = (\mathbf{r}_{ac}(t) - \mathbf{r}_{c_{start}}) - ((\mathbf{r}_{ac}(t) - \mathbf{r}_{c_{start}}) \cdot \tilde{\mathbf{p}})\tilde{\mathbf{p}} \tag{5.60}$$

The trajectory is then constrained within the corridor of maximum distance width $\hat{d}$

$$\|\mathbf{d}_{\perp}(t)\| \leq \hat{d} \text{ while } t_{c_{start}} \leq t \leq t_{c_{end}} \tag{5.61}$$

Waypoint crossing and corridor constraints have been designed to restrict aircraft to permissable regions of airspace. However, it may also be useful to change the equalities on the flight path constraints from $\leq$ to $\geq$ to convert them to obstacle constraints to represent no-fly zones and or airways that must not be crossed.

5.5 Summary: Trajectory Planning Methods

In summary, the problem of planning environmentally efficient trajectories is treated as an optimal control problem. Two categories of numerical based approaches were considered for solving the problem. These were direct and indirect methods.

Indirect methods typically involve representing the problem as a boundary value problem, which is discretised and solved using numerical techniques. For indirect approaches, the necessary conditions for optimality must be explicitly satisfied. Direct methods are based on a discretised finite dimensional parameterisation of the infinite dimensional problem [133]. Direct methods were considered to be the most appropriate for the case studies

explored in the thesis due to their ease of application to the multi-objective procedure optimisation problem. Therefore, the use of indirect methods were not pursued further.

The inverse method utilising the virtual domain (IDVD), was the direct method carried forward and applied to the case studies in this thesis. This was due to its low parameter space and promising performance in prior studies when combined with the stochastic Differential Evolution (DE) solver.

The parameterisation then of the states and controls using the inverse dynamics method, the cost functions established in Chapters 3-4, and the constraint functions defined in Section 5.4, allow the infinite dimensional optimal control problem to be converted to an NLP problem, where the trajectory search is conducted in the output space and only the algebraic equations need to be solved.

The NLP problem can then be presented to the NLP solver in the general form defined in Chapter 6. The most effective NLP algorithm for solving the problem is also investigated in Chapter 6.

Chapter 6

Non Linear Programming

6.1 Introduction

The main idea behind direct methods is to discretise the states and controls of the original continuous time optimal control problem in order to obtain a finite dimensional nonlinear programming (NLP) problem [134]. Numerical methods for solving NLP problems then involve the iteration of a set of variables from a finite set of unknowns [117]. This Section provides an overview of different NLP methods, particularly with regard to their adaption to multi-objective optimisation. The qualities needed from a NLP method for solving the environmental trajectory optimisation problems posed in this thesis are considered. The decision to ultimately adopt the global Differential Evolution NLP method is also discussed. The method is then defined in detail, including the adaptions required to apply the algorithm to many objective optimisation problems.

6.2 Solving the NLP Problem

Once converted, the optimal control problem can be solved as an NLP problem. Generally a NLP problem can be stated as the problem of finding a vector of optimisation variables $\bar{\mathbf{z}} \in \mathbb{R}^m$ that minimises the cost function

$$f(\bar{\mathbf{z}}) \tag{6.1}$$

subject to the constraints

$$\begin{aligned} \mathbf{g}(\bar{\mathbf{z}}) &= \mathbf{0} \\ \mathbf{h}(\bar{\mathbf{z}}) &\leq \mathbf{0} \end{aligned} \tag{6.2}$$

where the objective $f : \mathbb{R}^m \to \mathbb{R}$ and the constraints $\mathbf{g} : \mathbb{R}^m \to \mathbb{R}^p$, $\mathbf{h} : \mathbb{R}^m \to \mathbb{R}^q$.

To solve the continuous optimal control problem, it is commonly assumed that the constraints and objectives are twice continuously differentiable. This allows the Lagrangian to be stated as

$$\mathcal{L}(\bar{\mathbf{z}}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = f(\bar{\mathbf{z}}) + \boldsymbol{\lambda}^T \mathbf{g}(\bar{\mathbf{z}}) + \boldsymbol{\mu}^T \mathbf{h}(\bar{\mathbf{z}}) \tag{6.3}$$

where $\mathcal{L}(\bar{\mathbf{z}}, \boldsymbol{\lambda}, \boldsymbol{\mu})$ is a scalar function and where $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$ are vectors of Lagrange multipliers. The first-order necessary conditions at a point $\bar{\mathbf{z}}^*$ for a local minimum are then the Karush-Kuhn-Tucker conditions [135]

$$\begin{gathered} \nabla \mathcal{L}(\bar{\mathbf{z}}^*, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \nabla f(\bar{\mathbf{z}}^*) + \boldsymbol{\lambda}^T \mathbf{g}(\bar{\mathbf{z}}^*) + \boldsymbol{\mu}^T \mathbf{h}(\bar{\mathbf{z}}^*) = \mathbf{0} \\ \mathbf{g}(\bar{\mathbf{z}}^*) = \mathbf{0} \\ \mathbf{h}(\bar{\mathbf{z}}^*) \leq \mathbf{0} \\ \boldsymbol{\mu} \geq 0 \\ \boldsymbol{\mu}^T \mathbf{h}(\bar{\mathbf{z}}^*) = 0 \end{gathered} \tag{6.4}$$

There remains a relationship between the optimal control and the NLP problems such that the KKT conditions and the Lagrange multipliers approximate the optimal control necessary conditions and the adjoint variables [77, 135]. Therefore there is an equivalence between a local minimum of the NLP problem and a local minimum of the optimal control problem [77, 136, 137, 138]. Where objectives are non-differentiable and multi-modal, then NLP methods that either approximate derivatives or do not require them may be used. However, if KKT conditions cannot be evaluated, then a measure of the optimality and convergence of the trajectory solution is lost. There are strengths and weaknesses to different NLP methods and these are discussed with reference to the environmental optimisation goals of the work in this thesis.

In Section 2.6.3 it is highlighted that a global NLP method is required that can support both the single and multi-objective optimisation of environmentally efficient trajectories.

Rao classified the numerical methods applied to the solution of trajectory optimisation problems as either gradient based or heuristic [117].

For constrained problems, gradient methods aim to find the root of the gradient of the Lagrangian [139]. To do this requires the analytic definition or estimation of the derivatives of the objective and the constraint functions. Sequential Quadratic Programming methods are often considered to be the most effective gradient based nonlinear programming approaches [140, 141, 142, 117]. SQP methods seek to approximate the optimisation problem as a series of quadratic sub-problems with linear constraints in the form

$$\begin{aligned} \min_{\mathbf{d}} \quad & \frac{1}{2} \mathbf{d}^T \nabla^2_{zz} \mathcal{L}(\bar{\mathbf{z}}_k, \boldsymbol{\lambda}_k, \boldsymbol{\mu}_k) \mathbf{d} + \nabla f^T(\bar{\mathbf{z}}_k) \mathbf{d} \\ & \nabla \mathbf{g}^T(\bar{\mathbf{z}}_k) \mathbf{d} + \mathbf{g}(\bar{\mathbf{z}}_k) = \mathbf{0} \\ & \nabla \mathbf{h}^T(\bar{\mathbf{z}}_k) \mathbf{d} + \mathbf{h}(\bar{\mathbf{z}}_k) \leq \mathbf{0} \end{aligned} \tag{6.5}$$

where k is the k^{th} iteration and $\nabla^2_{zz}\mathcal{L}(\bar{\mathbf{z}}_k, \boldsymbol{\lambda}_k, \boldsymbol{\mu}_k)$ is the Hessian of the Lagrangian. The quadratic problem in (6.5) is then solved for the step direction $\mathbf{d} = \bar{\mathbf{z}} - \bar{\mathbf{z}}_k$. Once the step direction $\mathbf{d}$ is determined a merit function ϕ is used to choose a step length α such that $\phi(\bar{\mathbf{z}}_k + \alpha\mathbf{d}) < \phi(\bar{\mathbf{z}}_k)$. The merit function combines both the objective and the constraints creating a scalar indicator for feasible reductions in the objective that provides a measure of progress towards convergence. The subproblems are then solved sequentially to converge on the solution $\bar{\mathbf{z}}^*$ [140]. For SQP methods where the Hessian of the Lagrangian is not calculated, it may be replaced by a quasi-Newton approximation derived from successive gradient vectors instead [79].

If the problem is unconstrained, then the SQP method reduces to Newton's method for finding the root of the gradient of the objective function [139]. SQP methods directly evaluate the KKT conditions to determine if the optimisation has converged [79].

Gradient solvers have advantages over other solvers in terms of speed and accuracy [79]. Gradient methods, under suitable conditions, are globally convergent, in that they assure convergence to a local solution from any starting point. SQP methods, using gradient and Hessian information display fast local convergence to a solution [140]. Gradient solvers however require the definition or estimation of the derivatives of the objective function and are susceptible to becoming trapped in a local minima if the objectives or constraints are nonconvex [135]. Gradient methods can be applied to the multi-objective problem by using a weighted means cost function, or used to generate a Pareto front by turning the multi-objective problem into a series of single objective optimisation problems, such as with the ϵ-constraint and lexiographic methods [143]. However, using any scalarisation approach requires the specification of objective weightings prior to the simulation, which assumes some advance knowledge of the solution trade-off surface. Common gradient solvers utilised in the literature include NPSOL, IPOPT and SNOPT [135].

The NLP methods most commonly utilised in the literature for environmental trajectory optimisation were single objective gradient based methods [144]. In [68, 69, 70, 71], derivative based gradient methods were used with weighted sum objective functions to look at noise and emissions impacts from procedure designs. To examine trade-offs in noise and emissions however, analysts had to run simulations multiple times, manually altering objective value weightings for each scenario. As the INM/Doc29 NPD noise calculation method used in the studies is grid based, highly nonlinear and multi-modal, derivative information is difficult to estimate. Therefore without precisely chosen starting points, solutions in such studies are likely to converge to local minima.

Derivative free optimisers such as Hooke Jeeves, Nelder-Mead, Simulated Annealing and Particle Swarm are heuristic methods that do not use gradient information to determine step size and compare only the objective function values when determining consecutive steps [145]. Derivative free algorithms do not evaluate the KKT conditions, which can lead to slower and less optimal convergence [125]. However, derivative free algorithms do

not require objectives and constraints to be twice differentiable and are therefore applicable to a wider range of problems [145]. Derivative free methods are also typically less sensitive to nonlinearities, discontinuities and numerical noise in the objectives or the constraints [145]. Similar to gradient methods, derivative free methods can be applied to the multiobjective problem by combining multiple objectives into a single performance measure such as a weighted mean cost function, or used to generate a Pareto front by turning the multiobjective problem into a series of single objective optimisation problems [143, 146]. However, as with the gradient methods, it is often difficult to estimate objective weightings or ϵ values prior to any simulation runs [143].

A form of derivative free heuristic NLP algorithms are evolutionary algorithms. Evolutionary algorithms are designed around the Darwinian principles of natural selection and survival of the fittest. In nature, random mutation in genes can lead to advantageous changes in an organisms pheneotype that allow it to better survive its environment. The better an organism is at surviving the more opportunity it has to reproduce and the more likely the genes and the mutations they contain will be passed on to future generations. As organisms not possessing advantageous gene mutations are more likely to die out relative to better adapted organisms, the surviving organisms are seen to be well adapted to the environment they find themselves in. Evolutionary algorithms adopt some of the principles of natural selection and involve random mutation of input variables, selection between candidate solutions and reproduction by the fittest individuals [117]. The aim of evolutionary algorithms is to evolve a solution that is well adapted to the problem they are applied to.

EA's as with other derivative free optimisers, are less sensitive to nonlinear, multi-modal objective functions and constraints than gradient methods [145, 147]. EA's, at each step of an optimisation, maintain a population of solutions, allowing the algorithms to simultaneously explore different parts of the solution space at once. This reduces the likelihood that the optimisation gets prematurely trapped in one specific part of the solution space (i.e. in a local minima). As EA's seek to converge to a global rather than a local minimum, they are referred to as a global optimisation technique [143]. However as they do not utilise any conditions for optimality other than the objective function value, there is no way of knowing if the algorithm has arrived at even a local solution [77]. Also the trial and error approach of evolutionary algorithms means that they are likely to require a large number of function evaluations before converging on a solution, which can be an issue if the objective functions are expensive in terms of time to evaluate [147, 143].

EA's are perhaps the most adaptable of all the solvers to the multiobjective optimisation problem. EAs, throughout even a single objective optimisation, natively maintain a set of solutions distributed in the solution space. For multiobjective problems, which are unlikely to have a single solution, the evolutionary algorithm may be adapted such that this solution set is required to converge to the Pareto optimal solution set[143]. This

potentially allows the entire trade-off surface between objectives to be defined after a single optimisation run [146].

It was decided to adopt a global stochastic NLP method for the environmental trajectory optimisation work in this thesis. As a number of the environmental objectives are non-linear, non-differentiable and multi-modal, stochastic methods offered a critical enhanced ability to escape local minima. The desire to have a Pareto trade-off front drove the specific adoption of an evolutionary algorithm. The NLP method ultimately chosen was the global, stochastic Differential Evolution method. The choice of the DE method, it's adaption and application are discussed further in the following sections.

6.3 Differential Evolution

Differential Evolution (DE) is a evolutionary algorithm that has already shown promise when used with the Inverse dynamics method. Differential Evolution is a simple and effective heuristic for global optimisation [148, 145]. As with many evolutionary algorithms, DE creates a randomly generated first parameter population. Elements of the initial population are then combined to form a trial population. The objective values from the initial population are tested against those of the trial population to determine the population members to be carried forward to a future generation. The process of mutation, combination and selection continues until no better objective value can be found.

One of the defining characteristic of DE is the differential mutation mechanism. Differential mutation is a self adaptive mechanism where a new mutant parameter vector $\bar{\mathbf{v}}$ is created by finding the difference between 2 parameter vectors $\bar{\mathbf{x}}_{r1}$, $\bar{\mathbf{x}}_{r2}$ and adding the scaled difference to a third $\bar{\mathbf{x}}_{r3}$. The mutation operation is $\bar{\mathbf{v}}_{i,G+1} = \bar{\mathbf{x}}_{r3,G} + C\left(\bar{\mathbf{x}}_{r1,G} - \bar{\mathbf{x}}_{r2,G}\right)$ for $i = 1, \ldots, NP$, where $r1, r2, r3 \in \{1, \ldots, NP\}$ are randomly chosen except that $r1 \neq r2 \neq r3 \neq i$. The scale factor $C \in (0, 1]$ and NP is the number of populations. Figure 6.1 graphically shows the differential mutation process in the solution space of a two variable objective function, where 3 target vectors $\bar{\mathbf{x}}_{r1}, \bar{\mathbf{x}}_{r2}, \bar{\mathbf{x}}_{r3}$ are used to create a mutant vector $\bar{\mathbf{v}}_1$ that moves the stochastic search closer to the function minimum.

Difference vectors are also employed by other NLP algorithms such as Nelder Mead. One of the advantages of using difference vectors is that they scale the step size to the objective function surface [145]. Large distances between the parameter vectors lead to larger step sizes and a wider sampling of the objective function. Smaller distances between the parameter vectors lead to smaller step sizes and a more localised sampling of the objective function. For multi-modal functions, the algorithm can use both small step sizes to locally explore basins of attraction and large step sizes to transport vectors between basins and to sample the objective space more broadly [145]. As the algorithm converges, the population contains more closely spaced parameter vectors, so differences

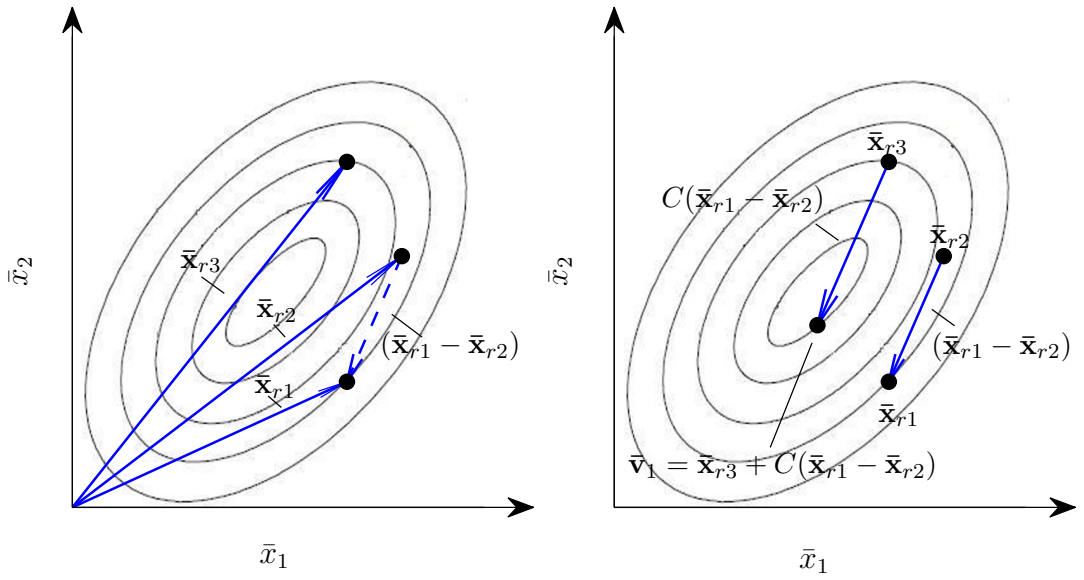

FIGURE 6.1: Differential mutation process

between vectors reduce and DE self adapts to search a smaller localised area of the objective space.

The other defining characteristic of DE is the Crossover Constant (CR). The crossover constant is a value between 0 and 1 that determines how much of the mutant vector $\bar{\mathbf{v}}$ is crossed with the target vector $\bar{\mathbf{x}}$ to form the trial vector $\bar{\mathbf{u}}$. The higher the CR value the more greedy the trial vector becomes for mutant parameters.

$$\bar{u}_{j,i,G+1} = \begin{cases} \bar{v}_{j,i,G+1} & \text{if } \mathrm{rnd}_j[0,1) \leq CR \vee j = k \\ \bar{x}_{j,i,G} & \text{otherwise} \end{cases} \qquad (6.6)$$
$$k \in \{1, \ldots, D\} \quad \text{randomly chosen index}$$

where D is the length of the parameter vector.

From a uniform distribution, a random number is generated for each individual of the parameter vector. If the number is higher than the CR value and the index j is not equal to a randomly chosen index k, then the target parameter is carried over to the trial vector. When the number is lower than the CR value, or when the index j is equal to a randomly chosen index k, the mutant parameter is carried over to the trial vector. Figure 6.2 shows the crossover of a target and mutant vector to form a trial vector, where 3 mutant vector parameters have a randomly generated index less that CR and are so carried through to the trial vector. The chosen value of CR between 0 and 1 controls the amount of new mutant and existing target parameters that are carried through to the trial population.

For single objective selection, DE has a greedy selection criteria. If the trial vector has a lower objective value than the target vector, is feasible when the target is not, or has a lower overall constraint violation than the target vector, then the trial replaces the target

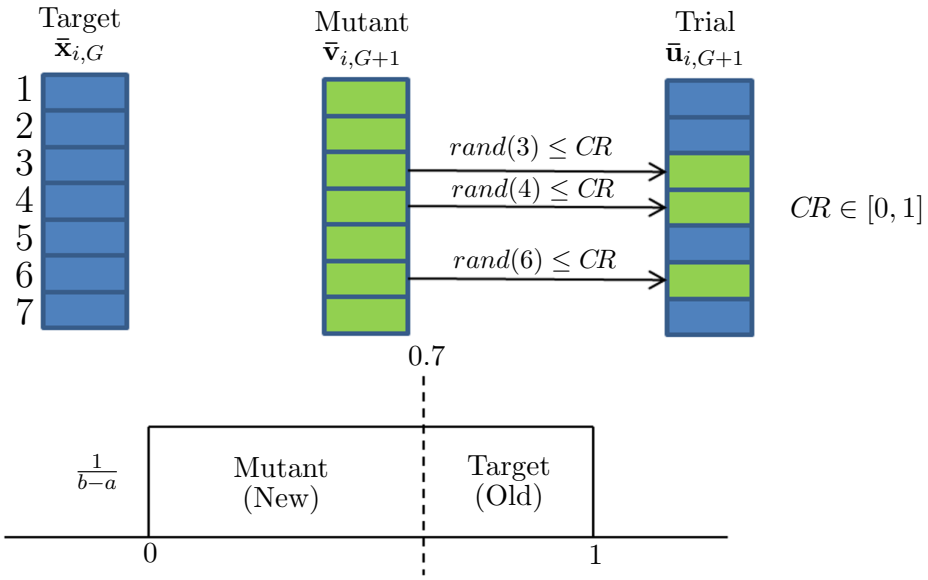

FIGURE 6.2: Differential Evolution crossover

in the population of the next generation. Where $\mathbf{g}$ is an array of inequality constraints, the overall constraint violation ζ for the target and trial vectors are calculated as follows

$$\begin{aligned} \zeta(\bar{\mathbf{x}}_{i,G}) &= \sum_{k}^{|\mathbf{g}|} \max[0, g_k(\bar{\mathbf{x}}_{i,G})] \\ \zeta(\bar{\mathbf{u}}_{i,G+1}) &= \sum_{k}^{|\mathbf{g}|} \max[0, g_k(\bar{\mathbf{u}}_{i,G+1})] \end{aligned} \tag{6.7}$$

with selection occuring as

$$\bar{\mathbf{x}}_{i,G+1} = \begin{cases} \bar{\mathbf{u}}_{i,G+1} & if \begin{cases} \begin{cases} \zeta(\bar{\mathbf{u}}_{i,G+1}) \leq 0 \wedge \zeta(\bar{\mathbf{x}}_{i,G}) \leq 0 \\ \wedge \\ f(\bar{\mathbf{u}}_{i,G+1}) < f(\bar{\mathbf{x}}_{i,G}) \end{cases} \\ \vee \\ \begin{cases} \zeta(\bar{\mathbf{u}}_{i,G+1}) \leq 0 \\ \wedge \\ \zeta(\bar{\mathbf{x}}_{i,G}) > 0 \end{cases} \\ \vee \\ \begin{cases} \zeta(\bar{\mathbf{x}}_{i,G}) > 0 \\ \wedge \\ \zeta(\bar{\mathbf{u}}_{i,G+1}) < \zeta(\bar{\mathbf{x}}_{i,G}) \end{cases} \end{cases} \\ \bar{\mathbf{x}}_{i,G} \end{cases} \tag{6.8}$$

DE requires only three configuration parameters for calibration and has been shown to be an effective solver when combined with the IDVD method for single objective trajectory optimisation problems. Drury [149] tested the performance of DE relative to the SNOPT gradient solver and the Nelder Mead and the Hook Jeeves derivative free solvers with the IDVD method for a minimum flight time problem with 2000 different boundary value sets.

DE outperformed all the other solvers in terms of robustness and relative optimality. Robustness was measured by the ratio of successful test cases to all test cases, where the defined convergence criteria was reached prior to the algorithm reaching an upper trajectory evaluation limit.

For the calculation of relative optimality, a reference optimal solution was defined for each test case. The reference optimal solution for each test case was arrived at by running the inverse method with each of the NLP algorithms. The shortest flight time found using any of the NLP methods was then set as the reference optimal flight time t_f^* for that test case. Relative optimality for each algorithm was then determined by measuring over all test cases how often the converged trajectory flight time t_f was within 2% of t_f^*.

Tested over a sample of 2000 pairs of boundary conditions, DE achieved a robustness of 99.8% and a relative optimality score of 94%. For the other algorithms investigated in the study, SNOPT had a robustness of 65-70% with Nelder Mead and Hook Jevees being approximately 95% robust. However, the relative optimality achieved by the Nelder Mead and Hook Jevees algorithms was 79% and 53% respectively, which was low compared to the 87% attained attained by SNOPT and the 94% acheived by DE.

DE however was an order of magnitude slower than the other NLP methods, requiring 8-12 times as many trajectory evaluations as the Hook Jeeves and SNOPT solvers, and requiring up to 4 times as many trajectory evaluations as the Nelder Mead solver [79].

Due to the offline planning nature of the problems considered in this thesis, the extra trajectory evaluations required by DE were not considered to be a significant issue in this work. As an evolutionary algorithm, DE was also considered to be very adaptable for application to the multiobjective trajectory optimisation problem.

Therefore, as DE showed considerable potential when combined with the IDVD method for UAV trajectory optimisation problems with short flight times, the combination of methods has again been considered in this work for solving commercial aircraft trajectory optimization problems over longer flight times.

The following section details the extension of the DE algorithm for it's application to the multiobjective optimisation problem.

6.4 Multiobjective DE

Extended from (6.1)-(6.2), the multiobjective NLP problem can be stated as the problem of minimising a vector of k scalar objectives

$$\min_{\bar{\mathbf{z}} \in \mathbb{R}^m} [f_1(\bar{\mathbf{z}}), f_2(\bar{\mathbf{z}}), \ldots, f_j(\bar{\mathbf{z}})] \tag{6.9}$$

subject to the constraints

$$\begin{aligned} \mathbf{g}(\bar{\mathbf{z}}) &= \mathbf{0} \\ \mathbf{h}(\bar{\mathbf{z}}) &\leq \mathbf{0} \end{aligned} \tag{6.10}$$

If the objectives are complementary, then there exists a vector of optimisation variables $\bar{\mathbf{z}}^*$ that provides the global minimum for all j objective functions [143]. In practice though a complementary global minimum for all objectives rarely exists [143]. Therefore the objective becomes finding the global Pareto front, which provides a solution set that identifies the best trade-offs between the objectives. DE, as an evolutionary algorithm can be adapted to converge on the Pareto optimal set as a whole [143].

Differential evolution, at each step of an optimisation, maintains a population of solutions, allowing the algorithm to simultaneously explore different parts of the solution space. This makes the algorithm well suited for adaption to multi-objective optimisation. Ranier and Storn proposed a multi-objective method with Pareto dominance selection in [145]. In general, for 2 feasible solutions where $\bar{\mathbf{p}}, \bar{\mathbf{q}} \in S$, $\bar{\mathbf{p}}$ dominates $\bar{\mathbf{q}}$ $(\bar{\mathbf{p}} \prec \bar{\mathbf{q}})$ if $\forall j : f_j(\bar{\mathbf{p}}) \leq f_j(\bar{\mathbf{q}}) \wedge \exists j : f_j(\bar{\mathbf{p}}) < f_j(\bar{\mathbf{q}})$. Therefore, applied to DE, the trial vector replaces the target vector in the population of the next generation if the trial is found to dominate the target vector. This method was found in tests to result in a very slow convergence on the Pareto front. The method also did not have any mechanism for preserving diversity along the Pareto front, resulting in solution sets that were a poor representation of the true front.

Fonseca and Fleming [150] with their MOGA algorithm and Srinivas and Deb [151] with their NSGA algorithm highlighted the importance of elitism and diversity in finding the Pareto front solution set. Therefore work in [146, 152] proposed supplementing DE's mutation mechanism with selection methods based on nondominated sorting and the crowding distance measure proposed by Deb for the Nondominated Sorting Genetic Algorithm II (NSGAII) [153]. It has been shown by Madavan et al [146] and Robic et al [154] that using nondominated sorting and the crowding distance metric with DE can be very effective at reducing the number of function evaluations required to reach the Pareto front, for improving optimisation convergence on the true Pareto front, and for maintaining solution diversity. The DEMO (Differential Evolution for Multiobjective Optimisation) [154] and GDE3 (Generalized Differential Evolution) [155] methods further supplement nondominated sorting and crowding distance with an additional greedy selection step.

Nondominated sorting uses domination to rank each member of a population and then sorts the population into fronts that are sets of solutions with equal dominance ranking. Population truncation then occurs where the most elite ranking solutions are retained. Where truncation splits a front, the solutions in the front are further sorted by their crowding distance such that the most diverse of the solutions are preserved.

As both methods utilise nondominated sorting and crowding distance, the mature PDE and DEMO algorithms have both been adopted and implemented here for use with the IDVD method. However, in [156] it is shown that the NSGAII crowding distance measure, used natively by both the PDE and DEMO methods, and proposed by Deb in [153], only provides good diversity in the case of 2 objectives, and does not perform well when applied to problems with 3 or more objectives.

As experiments in this thesis require the use of three or more objectives, a second crowding distance measure was implemented. For problems with many objectives, Kukkonen [157] proposes a k nearest neighbour crowding distance measure. With k being the number of objectives, the measure is based on finding the distances to the k nearest neighbors and multiplying the distances together to determine a crowding distance value cd_{kNN}

$$cd_{kNN} = \prod_{i=1}^{k} L_2^{NN_i} \tag{6.11}$$

where $L_2^{NN_i}$ is the distance to the ith nearest neighbor according to the L_2 distance metric. Figures 6.3 and 6.4 show Pareto fronts for the 3 objective ($k = 3$) DTLZ1 problem. Figures 6.3(a) and 6.4(a) show the DTLZ1 test problem solved by the PDE and DEMO solvers with the standard NSGAII crowding distance measure. Figures 6.3(b) and 6.4(b) show the DTLZ1 test problem solved by the PDE and DEMO solvers with the cd_{kNN} crowding distance measure. It can be visually seen that selection incorporating the kNN algorithm and the cd_{kNN} crowding measure provides a better distribution of solutions that gives a better approximation of the true Pareto front.

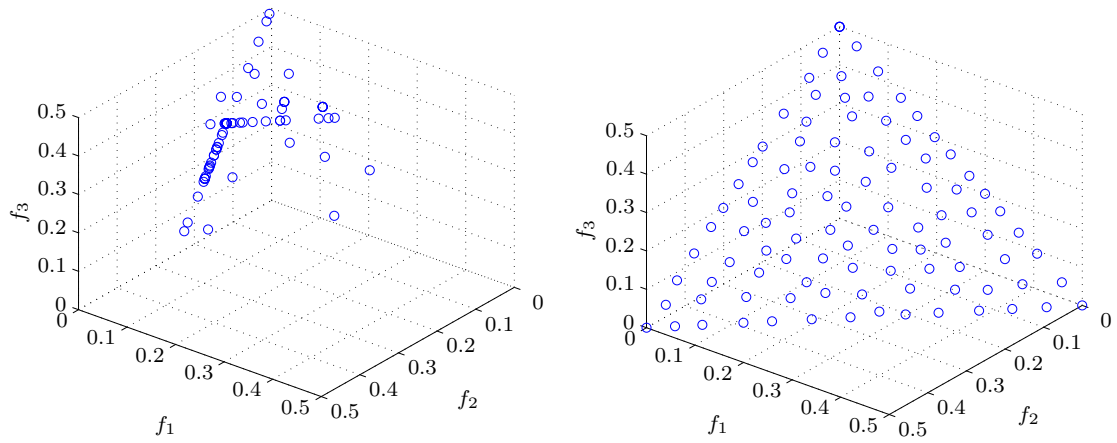

(a) PDE with NSGAII crowding

(b) PDE with cd_{kNN} crowding

FIGURE 6.3: DTLZ1 Problem solved with the PDE solver

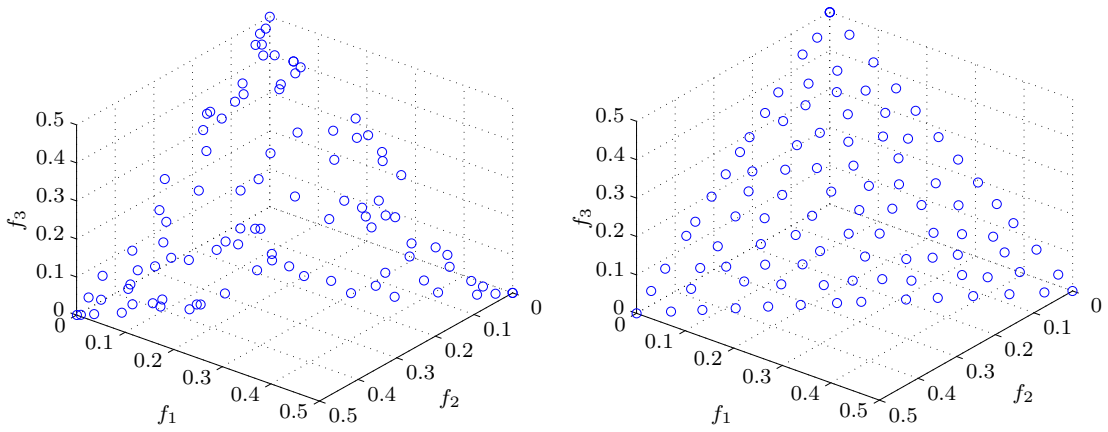

(a) DEMO with NSGAII crowding

(b) DEMO with cd_{kNN} crowding

FIGURE 6.4: DTLZ1 Problem solved with the DEMO solver

In addition to an improved crowding distance measure, the pruning algorithm proposed in [157] aims to avoid a brute force approach of calculating the distance from each solution to every other solution and uses pre-calculated projection values to provide a fast way to determine the distances between each solution and its k nearest neighbours.

Both the DEMO and PDE variants of multoobjective DE were adapted to use either the NSGAII crowding distance or the kNN crowding distance algorithms. The NSGAII crowding distance algorithm was retained as for 2 objective problems it was equally as effective and faster than the kNN algorithm. However, when more than two objectives were considered, the kNN algorithm was always used.

6.4.1 Nondominated Sorting

Nondominated sorting involves using domination to rank each solution into fronts that are sets of solutions with equal dominance ranking. Solutions that are not dominated by any other solutions are assigned to the first front $\mathcal{F}_1$, Solutions that are dominated by 1 other individual will appear on the next front $\mathcal{F}_2$, and so on until all solutions are assigned to the appropriate front. Figure 6.5 shows 3 individual fronts from the same population resulting from a 2 objective multiobjective optimisation. The solutions making up Front 1 are not dominated by any other solutions in the set. The solutions making up Front 2 are all dominated by a solution on Front 1 and all solutions on Front 3 are dominated by at least 2 solutions on the other fronts. The fast non dominated sorting algorithm implemented from [153] is shown in Algorithm 6.1

Algorithm 6.1 Fast-nondominated-sort(P)

For each $\bar{\mathbf{p}} \in P$
  $S_p = \emptyset$
  $n_p = 0$
  for each $\bar{\mathbf{q}} \in P$
    if$(\bar{\mathbf{p}} \prec \bar{\mathbf{q}})$ — If $\bar{\mathbf{p}}$ dominates $\bar{\mathbf{q}}$
      $S_p \frown \bar{\mathbf{q}}$ — Add $\bar{\mathbf{q}}$ to the set of solutions dominated by $\bar{\mathbf{p}}$
    elseif $(\bar{\mathbf{q}} \prec \bar{\mathbf{p}})$
      $n_p = n_p + 1$ — Increment the domination counter of $\bar{\mathbf{p}}$
    end
  if $n_p = 0$ — $\bar{\mathbf{p}}$ belongs to the first front
    $p_{rank} = 1$
    $\mathcal{F}_1 \frown \bar{\mathbf{q}}$
  end
$i = 1$ — Initialize the front counter
while $\mathcal{F}_i \neq \emptyset$
  $Q = \emptyset$ — Used to store the members of the next front
  for each $\bar{\mathbf{p}} \in \mathcal{F}_i$
    for each $\bar{\mathbf{q}} \in S_p$
      $n_q = n_q - 1$
      if $n_q = 0$ — $\bar{\mathbf{q}}$ belongs to the next front
       $q_{rank} = i + 1$
       $Q \frown \bar{\mathbf{q}}$
      end
    end
  end
  $i = i + 1$
  $\mathcal{F}_i = Q$
end

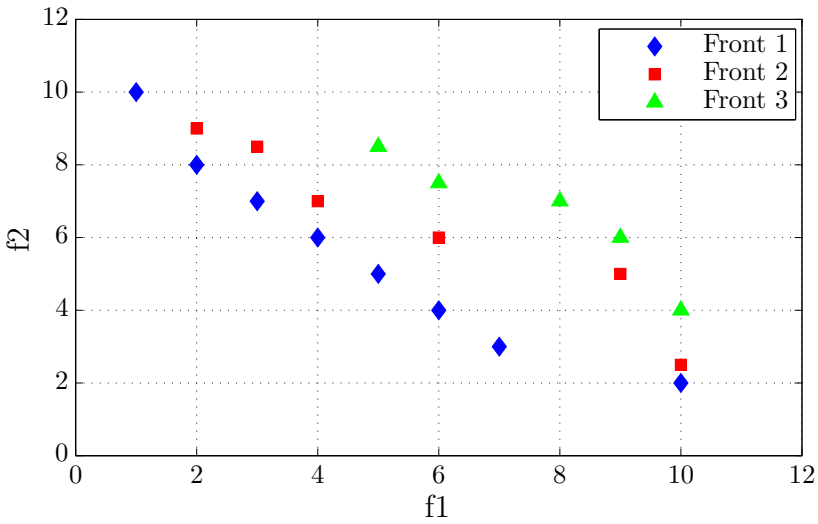

FIGURE 6.5: Nondominated sorting

6.4.2 Crowding Distance

Crowding distance measures how crowded a solution is by other nearby solutions. Selection using crowding distance encourages the even distribution of points along the Pareto front [158]. The NSGAII crowding distance metric is used to measure the distance along the same nondominated front from one solution to the 2 adjacent solutions, Figure 6.6 [153]. For each objective function, the greatest and smallest objective values are assigned an infinite crowding value, preserving the boundary value individuals in any crowding distance selection. For each intermediate individual, its proximity to other individuals is determined by taking the normalised difference between the solutions either side of that solution. When this measure is summed over all individuals objective functions, a measure of the closeness between solutions is reached.

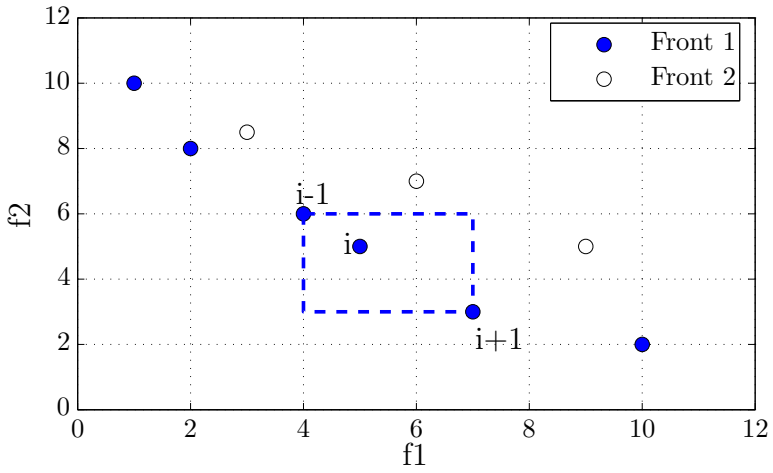

FIGURE 6.6: NSGAII crowding

Algorithm 6.2 Crowding-distance-assignment($\mathcal{F}$)

$l = |\mathcal{F}|$ Number of solutions in I
for each i, set $\mathcal{F}[i]_{cd_{NSGAII}} = 0$ **end** Initialise distance
for each objective m
  $\mathcal{F} = sort(\mathcal{F}, m)$ Sort using each objective value
  $\mathcal{F}[1]_{cd_{NSGAII}} = \mathcal{F}[l]_{cd_{NSGAII}} = \infty$ so that boundary points are always selected
  for $i = 2$ to $(l-1)$ For all other points

$$I[i]_{cd_{NSGAII}} = \mathcal{F}[i]_{cd_{NSGAII}} + \frac{\mathcal{F}[i+1].m - \mathcal{F}[i-1].m}{f_m^{max} - f_m^{min}}$$

  end
end

6.4.3 k Nearest Neighbour Pruning

In [157], Kukkonnen proposes a crowding distance measure based on multiplying together the distances to the k nearest neighbours. The solutions with the smallest products are the most crowded and those with with the largest products least crowded. In [157], k is set to be the number of objectives and the crowding distance is then calculated as $cd_{kNN} = \prod_{i=1}^{k} L_2^{NN_i}$ where $L_2^{NN_i}$ is the Euclidean distance to the ith nearest neighbor.

The method uses a projection vector and pre-calculated projection values to accelerate the nearest neighbour search reducing the number of distance comparisons that need to be made between solutions. The pruning algorithm proposed in [157] aims to avoid a brute force approach of calculating the distance from each solution to every other solution and proposes a fast way to determine the distances between each solution and its k nearest neighbours. A inequality constraint between the projection values and the Euclidean distance, derived in [159], places a limit on the number of distances that need to be calculated between nearby solutions as determined by their projection values. For two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ The inequality and the projection values are calculated as

$$(\mathrm{p}_x - \mathrm{p}_y)^2 \leq L_2(\boldsymbol{x}, \boldsymbol{y})^2, \quad \text{where} \quad \mathrm{p}_x = \frac{\tilde{\mathbf{p}} \cdot \boldsymbol{x}}{|\tilde{\mathbf{p}}|} \quad \text{and} \quad \mathrm{p}_y = \frac{\tilde{\mathbf{p}} \cdot \boldsymbol{y}}{|\tilde{\mathbf{p}}|} \tag{6.12}$$

The pruning algorithm (Algorithm 6.8) takes as its input a nondominated front $\mathcal{F}$ and a cutoff value that determines the number of solutions to be pruned from the front. The algorithm starts by normalising the objective values (Algorithm 6.3) to account for the possibility large value differences in objective values between different axes. To use the algorithm in its designed form [157], the normalised vales are then subtracted from 1 to determine each normalised objective value.

Algorithm 6.3 Normalise($\mathcal{F}$)

N = $|\mathcal{F}|$
for each objective m
$\mathcal{F} = \text{sort}(\mathcal{F}, m)$ — Sort front by objective values
$\mathcal{F}[1]_{cd_{kNN}} = \mathcal{F}[N]_{cd_{kNN}} = \infty$ — Set crowding distance of boundary points to ∞
$\mathcal{F}[1:N].m = 1 - \frac{(\mathcal{F}[1:N].m - f_m^{min})}{(f_m^{max} - f_m^{max})}$ — Normalise objective values
end

The next step in Algorithm 6.8 is to assigning each solution on the front a projection p value (Algorithm 6.4). Back and Sung [160] state that the projection vector should be chosen to represent the direction in which the objective vectors have the largest variance. Kukkonen [157] therefore recommends that for Pareto optimisation problems, where objective values are normalised over the range [0,1], that the projection axis be chosen to go through the points (0,0,...,0,1) and (1,1,...,1,0). Projection values are then calculated using (6.12)

Algorithm 6.4 Set-p-values($\mathcal{F}, k$)

$\mathbf{a} = \text{zeros}(1, k)$ — Generate a k length vector of zeros
$\mathbf{b} = \text{ones}(1, k)$ — Generate a k length vector of ones
$\mathbf{a}[k] = 1, \quad \mathbf{b}[k] = 0$
$\tilde{\mathbf{p}} = \mathbf{a} - \mathbf{b}$ — Create projection axis vector

for each $f \in \mathcal{F}$ — For each solution in the solution set
$f_{\text{p}_x} = \frac{\tilde{\mathbf{p}} \cdot f_{\mathbf{M}}}{|\tilde{\mathbf{p}}|}$
$f_{\text{p}_y} = f_{\text{p}_x}$ — Assign projection values to each solution
end

For Algorithm 6.4, the functions zeros$(1, k)$ and ones$(1, k)$ generate a $1 \times k$ sized vector of zeros and ones respectively.

Once each solution in the solution set has had it's objectives normalised and is assigned a projection $p_{x,y}$ value, an index list **I** of all solutions to be assigned a crowding distance measure is then created. In the first instance, every solution in the front will require a crowding value, so the initial list **I** is a index list of all the solutions on the front. As the pruning algorithm removes solutions from the front the index list **I** will be updated to contain only the indices of the solutions that need their crowding distance measure updated. Once **I** has been set or updated, the k nearest neighbour crowding distance cd_{kNN} can be assigned for each solution listed in **I** using Algorithm 6.5.

In Algorithm 6.5, for each solution, the k solutions upward and downward from the design solution are sampled and the squared distances between the normalised objective values are calculated. The maximum distance found is then set as d_{Emin}.

Further solutions are then sampled, where the distances between projection p values are compared. If the distance between the projection values satisfy the inequality $(p_x - p_y)^2 \leq d_{Emin}$, then actual distances d_E between solutions are compared. As further actual distances are calculated, if the kth largest d_E is less than d_{Emin} then the k^{th} largest d_E it is set to be d_{Emin}. Upward and downward steps continue to be taken as long as the inequality $(p_x - p_y)^2 \leq d_{Emin}$ holds. The use of the inequality removes the need to determine the distances between all the solutions and significantly reduces the number of comparisons needed to determine each individuals crowding distance.

Algorithm 6.5 kNN-search($\mathcal{F}, \mathbf{I}, k$)

$N =\mid \mathcal{F} \mid$
for each $i \in \mathbf{I}$ — For each solution of index i

  $m = i$, $n = m$
  $\mathcal{F}[i]_{\mathrm{d_{Emin}}} = 0$ — Initialise k^{th} squared euclidean distance
  $\mathcal{F}[i]_{\mathbf{d}_{k\mathbf{NN}}} = \emptyset$ — Initialise list of Nearest Neighbour (NN) distances
  $\mathcal{F}[i]_{\mathbf{idx}_{\mathbf{NN}}} = \emptyset$ — Initialise list of NN position indexes

  if $m = N$ **then** down = false **else** down = true **end**
  if $m = 1$ **then** up = false **else** up = true **end**

  while up **or** down — Search up and down front from current position
    if down
      $m = m + 1$
      $dm_{down} = \left\| \mathcal{F}[i]_{\mathrm{p_x}} - \mathcal{F}[m]_{\mathrm{p_y}} \right\|^2$ — Projection value distance bound in the down direction
      if $m \leq N$ and $m \leq i + k$ — Search k solutions down
        Set-kNN($\mathcal{F}, i, m$) — Set the k^{th} NN to the most distant
      else
        if $dm_{down} \geq \mathcal{F}[i]_{\mathrm{d_{Emin}}}$ — If k^{th} distance is greater than the distance bound
          down = false — Stop searching in the down direction
        else
          Update-kNN($\mathcal{F}, i, m, k$) — Update lists of NN distances and positions
        end
      end
      if $m = N$ **then** down = false **end**
    end
    if up
      $n = n - 1$
      $dm_{up} = \left\| \mathcal{F}[i]_{\mathrm{p_x}} - \mathcal{F}[n]_{\mathrm{p_y}} \right\|^2$ — Projection value distance bound in the up direction
      if $n \geq i - k$ and $n \geq 0$ — Search k solutions up
        Set-kNN($\mathcal{F}, i, n$) — Set the k^{th} NN to the most distant
      else
        if $dm_{up} \geq \mathcal{F}[i]_{\mathrm{d_{Emin}}}$ — If k^{th} distance is greater than the distance bound
          up = false — Stop searching in the up direction
        else
          Update-kNN($\mathcal{F}, i, n, k$) — Update lists of NN distances and positions
        end
      end
      $n = 1$ **then** up = false **end**
    end
  end
$[\mathcal{F}[i]_{\mathbf{idx}_{\mathbf{NN}}}, \mathcal{F}[i]_{\mathbf{d}_{k\mathbf{NN}}}] = \mathrm{sort}(\mathcal{F}[i]_{\mathbf{idx}_{\mathbf{NN}}}, \mathcal{F}[i]_{\mathbf{d}_{k\mathbf{NN}}})$
$\mathcal{F}[i]_{\mathbf{idx}_{\mathbf{NN}}} = \mathcal{F}[i]_{\mathbf{idx}_{\mathbf{NN}}}[1:k]$
$\mathcal{F}[i]_{cd_{kNN}} = \prod_{j=1}^{k} \mathcal{F}[i]_{\mathbf{d}_{k\mathbf{NN}}}[j]$ — Calculate crowding distance for current solution
end

Algorithm 6.6 Set-kNN($\mathcal{F}, i, l$)

$d_E = \|\mathcal{F}[i]_{\mathbf{M}} - \mathcal{F}[l]_{\mathbf{M}}\|^2$	Calculate L_2 distance between solution objectives
$\mathcal{F}[i]_{\mathbf{idx_{NN}}} \frown l$	Add position index to list of NN positions
$\mathcal{F}[i]_{\mathbf{d}_k\mathbf{NN}} \frown d_E$	Add distance to list of NN distances
if $d_E > \mathcal{F}[i]_{\mathrm{d_{Emin}}}$	If distance is greater than k^{th} distance
$\mathcal{F}[i]_{\mathrm{d_{Emin}}} = d_E$	Set k^{th} distance
end	

Algorithm 6.7 Update-kNN(F, i, l, k)

$d_E = \|\mathcal{F}[i]_{\mathbf{M}} - \mathcal{F}[l]_{\mathbf{M}}\|^2$	Calculate L_2 distance between solution objectives
if $d_E < \mathcal{F}[i]_{\mathrm{d_{Emin}}}$	
$\mathcal{F}[i]_{\mathbf{idx_{NN}}} \frown l$	Add position index to list of NN positions
$\mathcal{F}[i]_{\mathbf{d}_k\mathbf{NN}} \frown d_E$	Add distance to list of NN distances
$\mathcal{F}[i]_{\mathrm{d_{Emin}}} = \mathcal{F}[i]_{\mathbf{idx_{NN}}(k)}$	Update k^{th} distance,
end	where $_{(k)}$ is the k^{th} largest value in the list

Once crowding distance values are assigned, Algorithm 6.8 details the removal of the most crowded solution and how references to nearest neighbours are managed due to the solution removal. In [157], Kukkonen used linked lists and heaps to manage the references to nearest neighbours. The pruning algorithm (Algorithm 6.8) uses a simpler method for managing lists of nearest neighbour positions. The approach is slower than Kukkonens' but simplifies the implementation and still provides an order of magnitude improvement over a brute force approach.

Looking at Algorithm 6.8, the removal of the most crowded solution f_j, will cause the position index of solutions with higher position references to reduce their position index by one. Therefore, in preparation for the removal of a solution, all solutions in the front population are looped through and the position values in each solutions nearest neighbour index $f_{\mathbf{NNidx}}$ are reduced by 1 if they are greater than j.

This step can easily be vectorised, and Algorithm 6.8 describes this step with a pseudocoded version of Matlab's logical indexing. Therefore the expression $\mathbf{L} = [f_{\mathbf{NNidx}} > j]$ creates a vector $\mathbf{L}$ of equal size to $f_{\mathbf{NNidx}}$ whose elements are logical 1 where the elements $f_{\mathbf{NNidx}}$ are greater than j and whose elements are zero otherwise. The logical vector $\mathbf{L}$ then, when subtracted from a solutions $f_{\mathbf{NNidx}}$, simultaneously adjusts all the relevent pointers in $f_{\mathbf{NNidx}}$ in a single, simple step. As the length of each solutions nearest neighbour index is limited to the number of objectives $|f_{\mathbf{NNidx}}| = k$, the index update loop is generally very fast.

One remaining housekeeping task is needed before the solution is removed and that is to check if the most crowded solution is listed as a nearest neighbour in the current solutions nearest neightbour index $f_{\mathbf{NNidx}}$. If it is, then the index pointer n to the current solution is added to the list of solutions $\mathbf{I}$ whose crowding distance must be updated.

Once the nearest neighbour index $f_{\mathbf{NNidx}}$ and the crowding distance assignment index list $\mathbf{I}$ have been updated the most crowded solution can be removed from the front. The removal of the most crowded solution, updating of the crowding distance measures and the management of the indexes continues until the front is pruned to the desired size.

Algorithm 6.8 Pruning($\mathcal{F}$, cutoff)

```
N =| F |
set k to number of objectives
Normalise(F)                          Normalise objective values of nondominated set
SetPValues(F)                         Calculate projection values
sort(F, F[1 : N]_{p_y})               Sort F in descending order using projection values

I = (l)_{l=1}^{N}                     Indices list of solutions needing cd assignment
kNN-search(F, I, k)                   Assign crowding distance to all solutions indexed in I
Until N < cutoff                      Until front is pruned
    I = ∅                             Reset cd assignment list
    arg min_{j∈{1,...,N}} (F[j]_{cd_kNN})   Get index of most crowded solution
    n = 1                             Current solution index
    for each f ∈ F                    For each solution on the front
        L = [f_NNidx > j]             Logical index of NN solutions that must be updated
        f_NNidx = f_NNidx − L         All position indices greater than j must be reduced by 1
                                      to account for solution removal
        if j ∈ f_NNidx                Solutions referencing j as a NN must be updated
          if n > j                    Solutions with an index greater j must be removed
            I ⌢ n − 1                 Reduce position index by 1 to account for the solution
                                      removal, then add current solution n to cd assignment list
          else
            I ⌢ n                     Add current solution to cd assignment list
          end
        end
        n=n+1                         Update current solution index
    end
    F \ f_j                           Remove current solution from front
    N = N − 1                         Reduce size of front by 1
    kNN-search(F, I, k)               Update cd assignment for solutions belongong to I
end
```

6.4.4 Main

Algorithm 6.10 and Algorithm 6.11 show the PDE and DEMO algorithms that were implemented for this work. Both methods utilise selection steps involving nondominated sorting and crowding distance pruning. The relationships between the main PDE and DEMO algorithms and Algorithms 6.1-6.9 are flowcharted in Figures 6.8 and 6.9.

The mutation variant of DE chosen for both algorithms was DE/rand/1/Bin [145]. Both algorithms initialise by generating a population of random individuals between the user specified upper b_U and lower b_L parameter bounds.

For the Pareto Differential Evolution approach, a trial population Q_G is generated from a target population P_G through differential mutation and crossover. PDE then adopts the NSGAII selection steps shown in Figure 6.7. In these steps, the target (P_G) and trial (Q_G) populations are appended to each other to create an offspring population O_G, within which all selection occurs. The offspring population is sorted using the nondominated sorting, Algorithm 6.1, into fronts ($\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3 \ldots$) that are sets of solutions with equal dominance ranking. The offspring population is then reduced in size to the length of the initial target population $|P_G|$ such that the most elite ranking solutions are retained. Where truncation splits a front, the solutions in the last included front are further sorted and selected by their crowding distance (Algorithm 6.2) such that the most diverse of the solutions are preserved.

In both sorting steps, domination is used, initially to determine the most elite solutions and then, once crowding distance has been assigned, to determine the most diverse of the elite solutions. Where $\mathbf{g}$ is an array of inequality constraints, Algorithm 6.9 shows how domination is determined between two parameter vectors. Where n is the number of objectives, it can be seen that the parameter vector $\bar{\mathbf{q}}$ dominates the parameter vector $\bar{\mathbf{p}}$ if both produce feasible solutions and if $\mathbf{f}(\bar{\mathbf{q}})$ is less than or equal to $\mathbf{f}(\bar{\mathbf{p}})$ with $f_j(\bar{\mathbf{q}}) < f_j(\bar{\mathbf{p}})$ for at least one index of j, $j \in \{1, \ldots, n\}$. The parameter vector $\bar{\mathbf{q}}$ also dominates where both solutions are feasible and $\bar{\mathbf{q}}$ has a greater crowding distance value at the same nondominated rank. If $\bar{\mathbf{q}}$ is feasible and $\bar{\mathbf{p}}$ is infeasible then $\bar{\mathbf{q}}$ dominates. If both solutions are infeasible, then the individual with the lowest overall constraint violation $\zeta(\bar{\mathbf{z}})$ dominates.

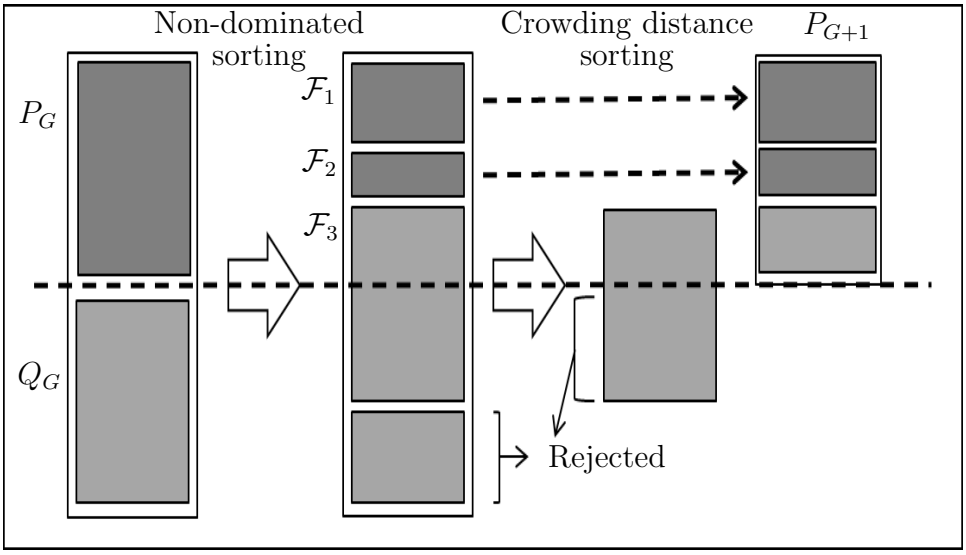

FIGURE 6.7: NSGAII selection steps

Algorithm 6.9 Domination-selection ($\bar{\mathbf{q}} \prec \bar{\mathbf{p}}$)

$$\zeta(\bar{\mathbf{q}}) = \sum_{k}^{|\mathbf{g}|} \max[0, g_k(\bar{\mathbf{q}})]$$

$$\zeta(\bar{\mathbf{p}}) = \sum_{k}^{|\mathbf{g}|} \max[0, g_k(\bar{\mathbf{p}})]$$

$$\bar{\mathbf{q}} \prec \bar{\mathbf{p}} \;\; if \begin{cases} \begin{cases} \zeta(\bar{\mathbf{q}}) \leq 0 \wedge \zeta(\bar{\mathbf{p}}) \leq 0 \\ \wedge \\ \mathbf{f}(\bar{\mathbf{q}}) \leq \mathbf{f}(\bar{\mathbf{p}}) \wedge \exists j : f_j(\bar{\mathbf{q}}) < f_j(\bar{\mathbf{p}}) \\ \vee \\ [\bar{q}_{\text{rank}} = \bar{p}_{\text{rank}}] \wedge [\bar{q}_{\text{cd}} > \bar{p}_{\text{cd}}] \end{cases} \\ \vee \\ \begin{cases} \zeta(\bar{\mathbf{q}}) \leq 0 \\ \wedge \\ \zeta(\bar{\mathbf{p}}) > 0 \end{cases} \\ \vee \\ \begin{cases} \zeta(\bar{\mathbf{q}}) > 0 \\ \wedge \\ \zeta(\bar{\mathbf{q}}) \;\; < \zeta(\bar{\mathbf{p}}) \end{cases} \end{cases}$$

It can be seem from the pseudocode for Algorithm 6.11 and the flowchart in Figure 6.9 the DEMO algorithm closely resembles the PDE algorithm, with both algorithms employing the same mutation, crossover, nondomination and pruning steps. However the DEMO algorithm employs an additional greedy selection step after DE mutation and crossover but prior to the use of the nondominated sorting and pruning steps that it has in common with the PDE.

The greedy selection is used to pre-select the members of the offspring population with the following steps,

$$
\begin{aligned}
&\textbf{if } \bar{\mathbf{q}}_{i,G} \prec \bar{\mathbf{p}}_{i,G} \\
&\qquad O_G \frown \bar{\mathbf{q}}_{i,G} \\
&\quad \textbf{else if } \bar{\mathbf{p}}_{i,G} \prec \bar{\mathbf{q}}_{i,G} \\
&\qquad O_G \frown \bar{\mathbf{p}}_{i,G} \\
&\quad \textbf{else} \\
&\qquad O_G \frown \{\bar{\mathbf{p}}_{i,G}, \bar{\mathbf{q}}_{i,G}\} \\
&\textbf{end}
\end{aligned}
$$

where for each target parameter vector $\bar{\mathbf{p}}_{i,G}$ in the existing population, a potential replacement trial vector $\bar{\mathbf{q}}_{i,G}$ is generated and evaluated. The objectives of the trial and target vectors are then compared using Algorithm 6.9. If the objectives of the target dominate those of the trial, the trial vector is discarded and $\bar{\mathbf{p}}_{iG}$ gets assigned to the offspring population O_G. If the trial dominates the target, the trial $\bar{\mathbf{q}}_{i,G}$ gets assigned to the offspring population. If neither the target nor the trial dominate both vectors are assigned to the offspring population $O_G \frown \{\bar{\mathbf{p}}_{i,G}, \bar{\mathbf{q}}_{i,G}\}$. This results in an offspring population size that lies somewhere between $|P_G|$ and $2 \times |P_G|$. As with the PDE algorithm, the oversized offspring population is then reduced in size to the size of the original population $|P_G|$ by nondominated sorting and crowding distance to produce the population to be carried forward to the next generation. The extra selection step essentially makes the DEMO algorithm a greedier version of the PDE algorithm. Although the PDE algorithm has been used predominantly as the solver in this work, for some scenarios the greedier selection of the DEMO algorithm led to a better convergence to the Pareto front.

In practice, and as is shown in the flowcharts in Figure 6.9 and 6.8, the PDE and DEMO algorithms both had the option of using either the NSGAII pruning or kNN pruning algorithms shown in Algorithm 6.2 and 6.8 respectively. For the psuedo code however, the PDE algorithm (Algorithm 6.10) is shown implemented with the NSGAII pruning and the DEMO algorithm (Algorithm 6.11) is shown implemented with kNN pruning steps.

In general, the kNN algorithm and the cd_{kNN} metric produced better distributed Pareto fronts than the NSGAII pruning algorithm. Comparisons between the two algorithms are covered in detail in Appendix A. A significant reason for this is that the kNN algorithm prunes the Pareto front iteratively, updating the crowding distance measure each time the most crowded solution is removed. The NSGAII crowding distance algorithm on the other hand only assigns crowding distance once, which results in a less accurate pruning of the Pareto front. However, for scenarios having only two objectives, the NSGAII pruning was faster and provided sufficiently diverse Pareto fronts. For problems with three or more objectives the PDE and DEMO algorithms were always used with the kNN pruning method. The differences in the Pareto fronts, resulting from the differences in the two pruning methods for problems with more than two objectives is also covered in greater detail in Appendix A.

Algorithm 6.10 Pareto Differential Evolution (PDE)

$G = 1$ First generation
$P = \emptyset$ Initialise target population
for $i := 1$ **to** NP Create NP real valued vectors
  $\bar{\mathbf{p}}_i = \emptyset$ Initialise target vector
  for $j := 1$ **to** D Each target vector contains D real parameters
    $\bar{x}_{j,i,G} = \text{rand}(0,1)(b_{j,U} - b_{k,L}) + b_{j,L}$ Create parameters within bounds
    $\bar{\mathbf{p}}_i \frown \bar{x}_{j,i,G}$ Add the parameter to the target vector
  end
  $P \frown \bar{\mathbf{p}}_i$ Add the vector to the target populations
end

while $G < G_{max}$ While current generation is less than final generation
  $Q = \emptyset$ Initialise trial population
  for each $\bar{\mathbf{p}}_{i,G} \in P_G$ For each target vector in the population
    $r_1, r_2, r_3 \in \{1, ..., NP\}$ Select 3 random indexes
    $C \in [0, 1]$ Scale factor
    $\bar{\mathbf{v}}_{i,G} = \bar{\mathbf{p}}_{r_1,G} + C(\bar{\mathbf{p}}_{r_2,G} - \bar{\mathbf{p}}_{r_3,G})$ Create mutant vector with Differential mutation
    $CR \in [0, 1]$ DE crossover parameter CR
    $\kappa \in \{1, ..., D\}$ Random parameter index
    $\bar{\mathbf{q}}_i = \emptyset$ Initialise trial vector
    for each $\bar{x}_{j,G} \in \bar{\mathbf{p}}_{i,G}$ For each parameter in the vector
      $r =$rand$(0,1)$
      if $r <= CR \vee j == \kappa$ Crossover between target and the trial vectors
        $\bar{u}_{j,G} = \bar{v}_{j,G}$
      else
        $\bar{u}_{j,G} = \bar{x}_{j,G}$
      end
      $\bar{\mathbf{q}}_{i,G} \frown \bar{u}_{j,G}$ Add parameter to trial vector
    end
    $Q_G \frown \bar{\mathbf{q}}_{i,G}$ Add trial vector to trial population
  end
  $O_G = P_G \frown \{Q_G\}$ Create offspring population from target and trial
  $\mathcal{F} =$ Fast-nondominated-sort(O_G)
  while $|P_{G+1}| + |\mathcal{F}_l| \leq |P_G|$ Until the new target population P_{G+1} is filled
    Crowding-distance-assignment$(\mathcal{F}_l)$ Calculate the crowding distance in $\mathcal{F}_l$
    $P_{G+1} = P_G \frown \mathcal{F}_l$ Include the l^{th} nondominated front in P_{G+1}
    $l = l + 1$ Check the next front for inclusion
  end
  $sort(\mathcal{F}_l, \prec)$ Sort final front in descending order using domination
  $P_{G+1} = P_{G+1} \frown \mathcal{F}_l[1 : (N - |P_{G+1}|)]$ Truncate final front if required
  $G = G + 1$ Increment generation counter
end

Algorithm 6.11 Differential Evolution for Multi-Objective Optimisation (DEMO)

$G = 1$ — First generation
$P = \emptyset$ — Initialise target population
for $i := 1$ **to** NP — Create NP real valued vectors
  $\bar{\mathbf{p}}_i = \emptyset$ — Initialise target vector
  for $j := 1$ — Each target vector contains D real parameters
    $\bar{x}_{j,i,G} = \text{rand}(0,1)(b_{j,U} - b_{k,L}) + b_{j,L}$ — Create parameters within bounds
    $\bar{\mathbf{p}}_i \frown \bar{x}_{j,i,G}$ — Add the parameter to the population vector
  end
  $P \frown \bar{\mathbf{p}}_i$ — Add the vector to the target populations
end

while $G < G_{max}$ — While current generation is less than final generation
  $Q = \emptyset$ — Initialise offspring population
  for each $\bar{\mathbf{p}}_{i,G} \in P_G$ — For each target vector in the population
    $r_1, r_2, r_3 \in \{1, ..., NP\}$ — Select 3 random indexes
    $C \in [0, 1]$ — Scale factor
    $\bar{\mathbf{v}}_{i,G} = \bar{\mathbf{p}}_{r_1,G} + C(\bar{\mathbf{p}}_{r_2,G} - \bar{\mathbf{p}}_{r_3,G})$ — Create mutant vector with Differential mutation
    $CR \in [0, 1]$ — DE crossover parameter CR
    $\kappa \in \{1, ..., D\}$ — Random parameter index
    $\bar{\mathbf{q}}_i = \emptyset$ — Initialise trial vector
    for each $\bar{x}_{j,G} \in \bar{\mathbf{p}}_{i,G}$ — For each parameter in the vector
      $r = \text{rand}(0, 1)$
      if $r <= CR \vee j == \kappa$ — Crossover between target and the trial vectors
        $\bar{u}_{j,G} = \bar{v}_{j,G}$
      else
        $\bar{u}_{j,G} = \bar{x}_{j,G}$
      end
      $\bar{\mathbf{q}}_{i,G} \frown \bar{u}_{j,G}$ — Add parameter to trial vector
    end

    if $\bar{\mathbf{q}}_{i,G} \prec \bar{\mathbf{p}}_{i,G}$ — If trial dominates target
      $O_G \frown \bar{\mathbf{q}}_{i,G}$
    else if $\bar{\mathbf{p}}_{i,G} \prec \bar{\mathbf{q}}_{i,G}$ — If target dominates trial
      $O_G \frown \bar{\mathbf{p}}_{i,G}$
    else — Both are on the same nondominated front
      $O_G \frown \{\bar{\mathbf{p}}_{i,G}, \bar{\mathbf{q}}_{i,G}\}$
    end
  end
  $\mathcal{F} = \text{Fast-nondominated-sort}(O_G)$
  while $|P_{G+1}| + |F_l| \leq |P_G|$ — Until the new population P_{G+1} is filled
    $P_{G+1} = P_G \frown \mathcal{F}_l$ — Include the l^{th} nondominated front in P_{G+1}
    $l = l + 1$ — Check the next front for inclusion
  end
  $cutoff = NP - |P_{G+1}|$
  $\mathcal{F}_l = \text{Pruning}(\mathcal{F}_l, cutoff)$ — Prune the final front
  $P_{G+1} = P_{G+1} \frown F_l$
  $G = G + 1$ — Increment generation counter
end

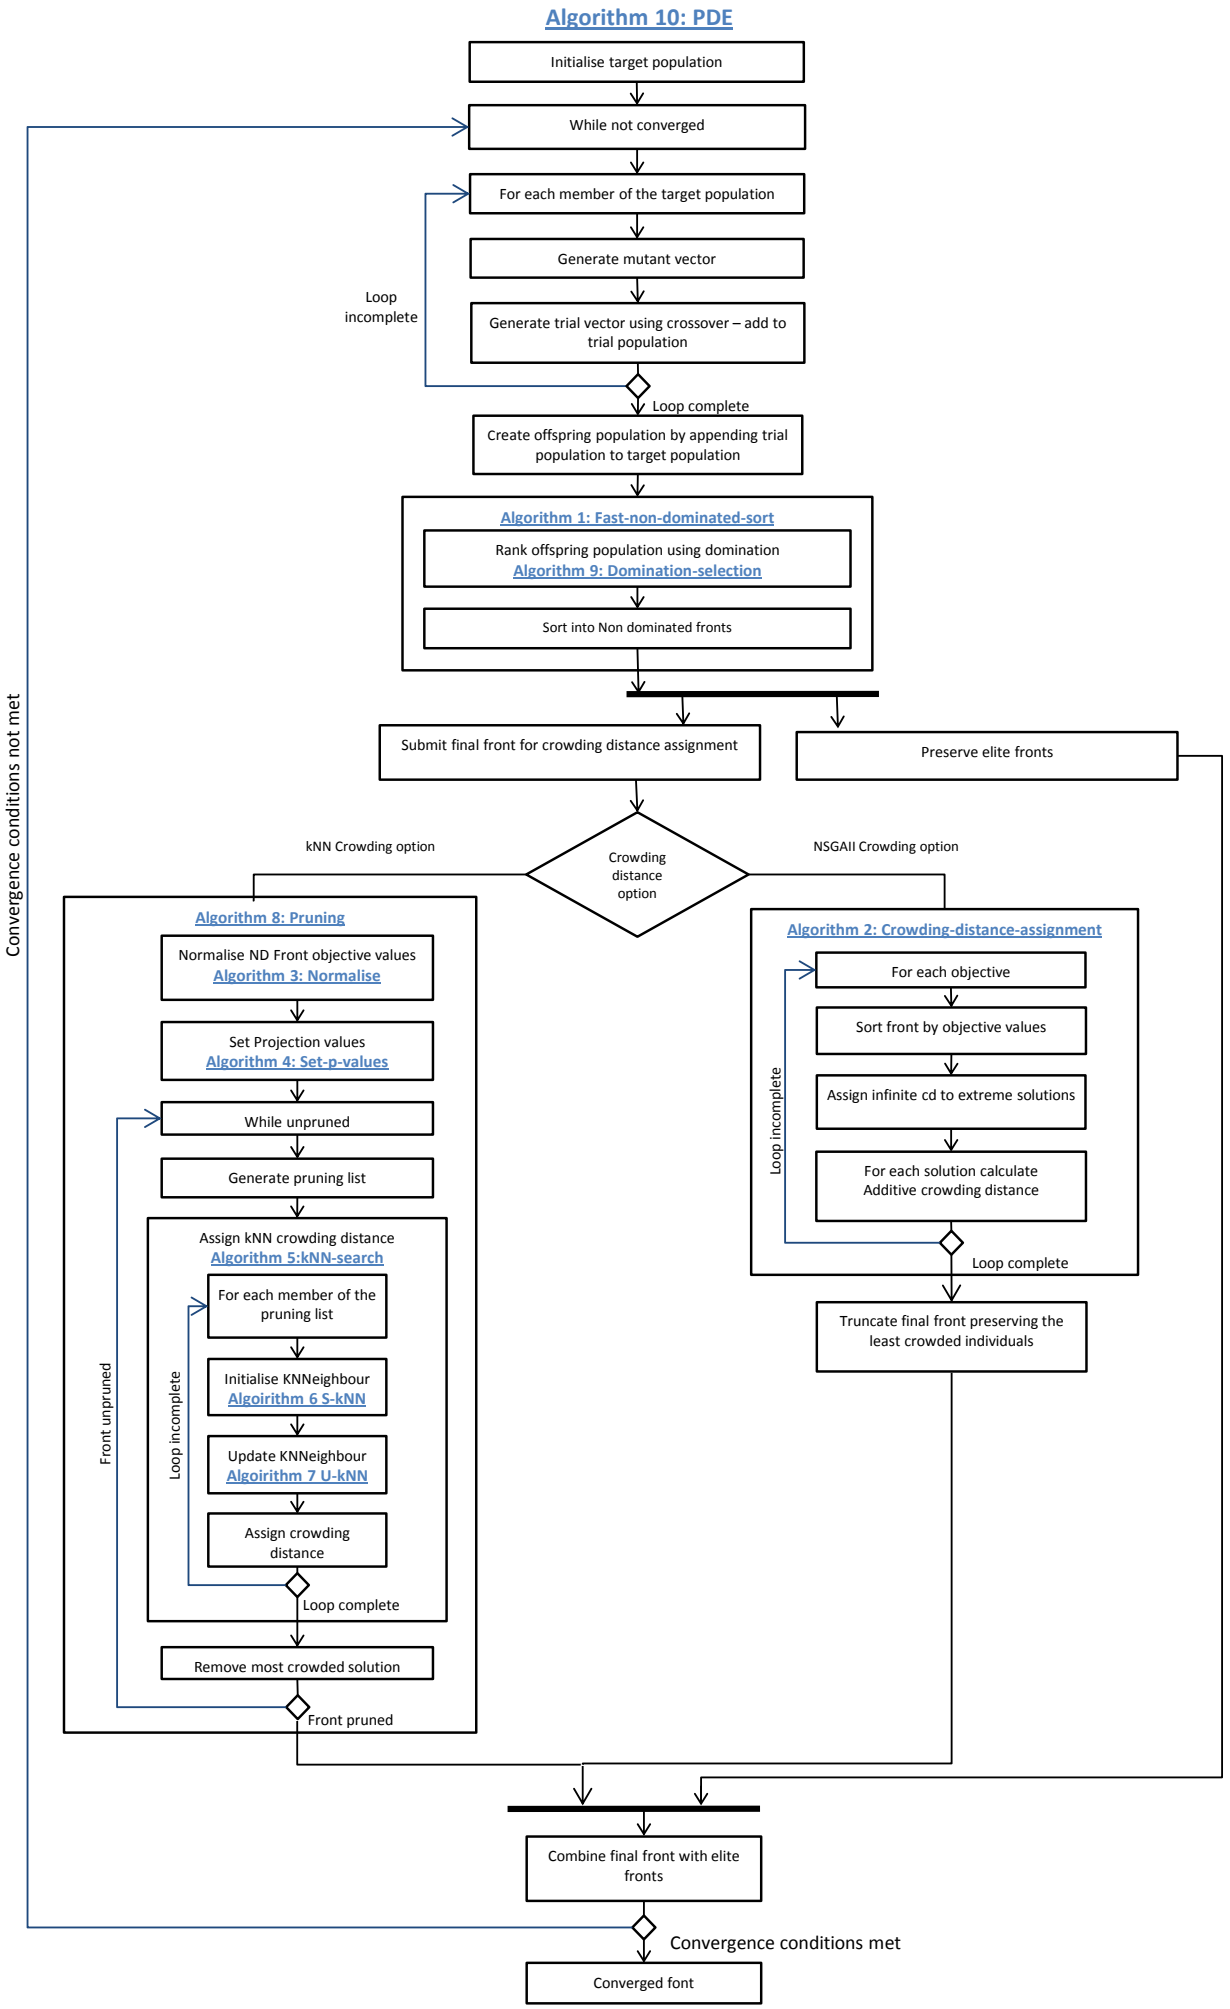

FIGURE 6.8: PDE method overview flowchart

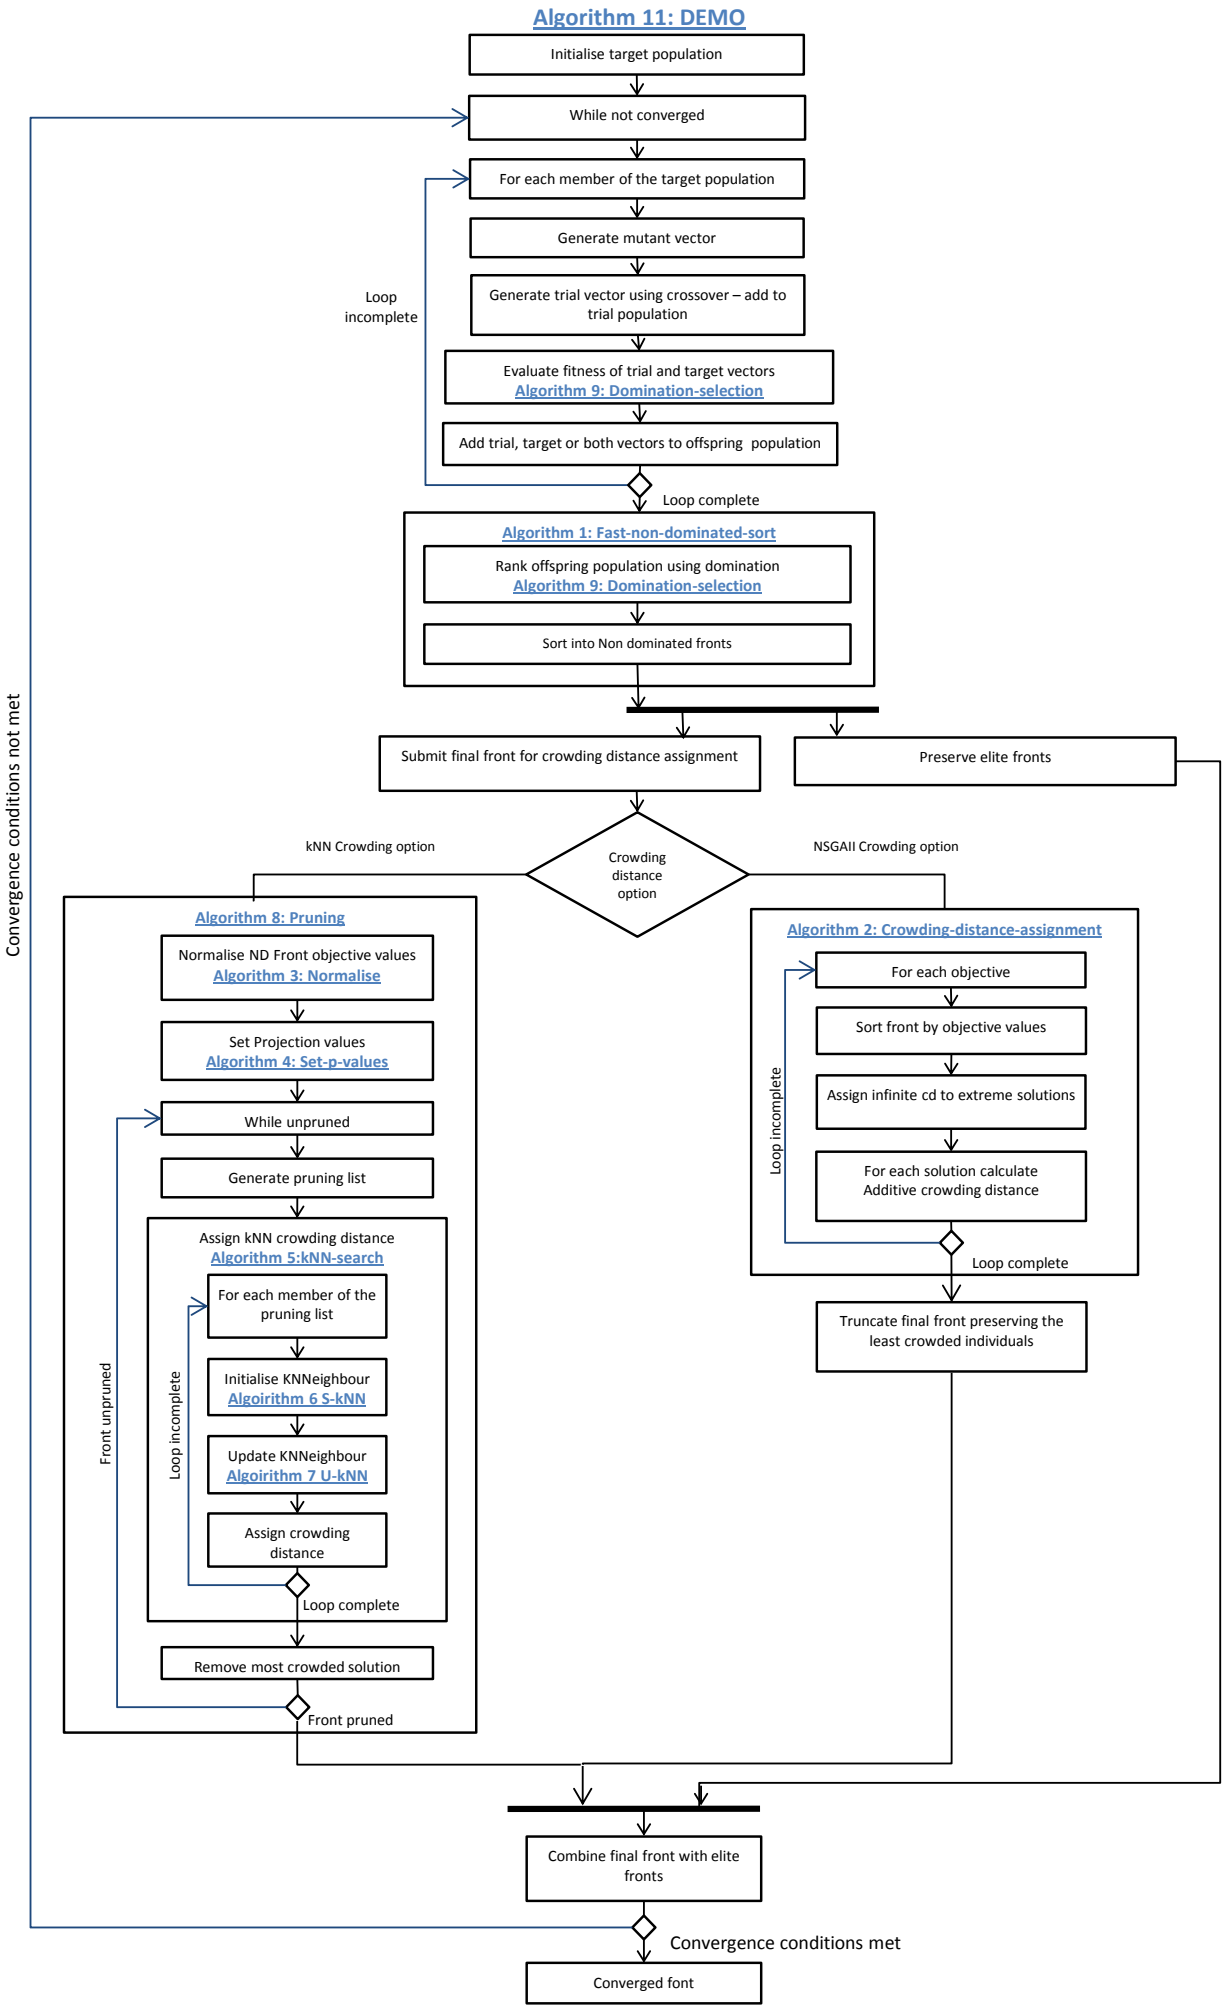

FIGURE 6.9: DEMO method overview flowchart

6.5 Summary: Algorithm and Method Integration

This chapter provides an overview of the Non Linear Programming (NLP) techniques considered for the work in this thesis. The stochastic Differential Evolution (DE) algorithm was ultimately chosen due to its adaptability to the multi-objective procedure optimisation problem and prior success when combined with the IDVD method (IDVD-DE). An additional benefit of the algorithm is that DE only requires two control parameters, making the algorithm relatively straightforward to calibrate.

For the multi-objective problem, the basic DE mutation and crossover mechanisms were supplemented with selection steps involving nondominated sorting and crowding distance pruning. Two algorithms were then used for the case studies herein. These are the PDE and DEMO algorithms. The only difference between the two is that the DEMO algorithm has an additional greedy selection step. For the same parameter settings, the DEMO algorithm was more aggressive in converging, the extra greedy step emphasising elitism over diversity. Generally, the performance of the algorithms was very similar. For case studies where the DEMO algorithn conferred a small advantage, it was used, otherwise the PDE algorithm was used.

Figure 6.10 shows an overview of the IDVD-DE optimisation method as utilised in this work. The scenario inputs are the trajectory boundary values and constraints as defined in Section 5.4. The Differential Evolution based solver, as detailed in Section 6.3, is used to generate initial target and trial populations of optimisation variable vectors. Using the Inverse Dynamics method, detailed in Section 5.4, the optimisation variables are used to generate a set of trajectories. The trajectories are then checked against the dynamics and scenario constraints to determine the feasibility of each trajectory. The feasible trajectories of the target and trial populations are then evaluated in terms of cost using the objective function methods and metrics described in Chapters 3-4. Once the trajectories have been generated, and their constraints and objective values evaluated, the solutions are returned to the DE algorithm where selection occurs for the fittest solutions. If the selected solution set meets the convergence criteria, the optimization loop stops and the results files are generated. If the selected set does not meet the convergence criteria, the set then becomes the target population for the next generation and further trial population are generated so that the IDVD-DE method continues to evolve the solutions until the convergence criteria are met.

In general, the kNN algorithm and the cd_{kNN} metric produced better distributed Pareto fronts than the NSGAII pruning algorithm. Comparisons between the two pruning algorithms are covered in detail in Appendix A. A significant reason for this is that the kNN algorithm prunes the Pareto front iteratively, updating the crowding distance measure each time the most crowded solution is removed. The NSGAII crowding distance algorithm on the other hand only assigns crowding distance once, which results in a less accurate pruning of the Pareto front. However, for scenarios having only two objectives,

the NSGAII pruning was faster and provided sufficiently diverse Pareto fronts. For problems with three or more objectives the PDE and DEMO algorithms were always used with the kNN pruning method. The differences in the Pareto fronts, resulting from the differences in the two pruning methods for problems with more than two objectives is also covered in greater detail in Appendix A.

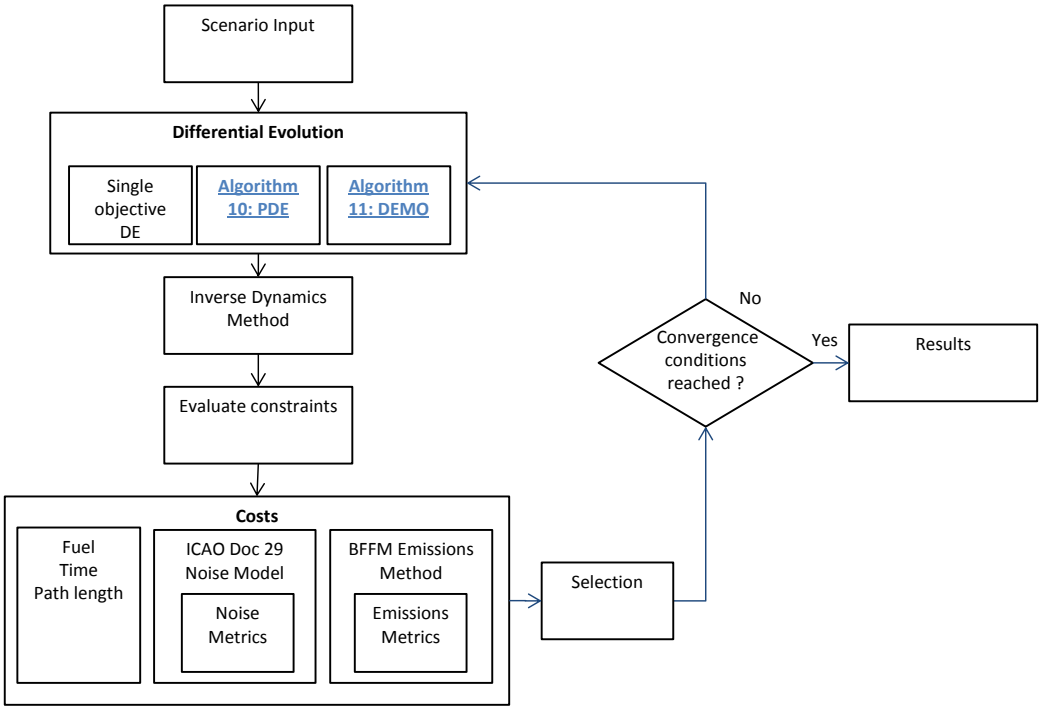

FIGURE 6.10: IDVD-DE Trajectory optimisation method overview

Chapter 7

Environmentally Efficient Trajectory Generation

7.1 Introduction

In this chapter the IDVD-DE trajectory optimisation method is combined with environmental metrics and applied to the solution of both single and multi-objective environmental trajectory optimisation problems. The IDVD-DE approach is first applied to the planning of a noise abatement departure aircraft operating procedure. The planning problem is one with a known solution, which is used for comparison with the IDVD-DE generated solution.

The IDVD-DE approach is further applied to the solution of a multi-environmental-objective trajectory optimisation problem. In this instance it is used to examine the trade-offs in community noise impact relative to emissions of CO2 for a realistic, but not real world, multi-objective case study. This scenario is also used to define a Pareto front analysis approach that is used in all of the remaining multi-objective case studies considered in this thesis.

For both the single and multi-objective problems the convergence characteristics of the IDVD-DE method are examined.

7.2 Sourdine Case Study

Before the IDVD-DE method is applied to more realistic multiobjective scenarios, how well the IDVD-DE method performs at generating environmentally efficient trajectories is first assessed. To do this the IDVD-DE method was used to generate a solution for a problem with a known environmentally efficient trajectory solution. The known

trajectory solution was developed as part of a European commission sponsored project titled Sourdine [67].

The Sourdine project was carried out by an international consortium of companies and research institutions with the aim of developing a series of recommended noise abatement procedures for different classes of aircraft [67]. A noise abatement procedure defines the operational steps taken by the aircraft when arriving or departing to or from an airport in order to minimise noise impact around the airport. Different departure noise abatement options were developed using expert analysis and were then used to generate trajectories that were tested by calculating their noise impact using a modified form of the Integrated Noise Model.

Table 7.1 shows the developed noise abatement procedure for a twin engine medium narrow-bodied aircraft. The aim of the procedure, here termed SDClose, is to minimise the noise impact close to the airport. Specifically, the procedure was shown to minimise EPNL under the centreline at distances 6-13km from brake release (BR) for an Airbus A319 better than any alternative procedure proposed in the study [161].

In this work, the SDClose procedure and its associated trajectory has also been used as a benchmark for comparison with the operating steps proposed by the IDVD-DE method. For the same scenario, it is expected that the IDVD-DE should be able to improve on or replicate the current best available trajectory solution.

Altitude (ft)	Noise Abatement take-off Procedure
1500	-Take-off with flexible thrust (Note 1). - Climb out at V2+10kt (Note 2).
$1500 \leq$ h (ft) $\leq$ 10,000	-Reduce thrust to cutback thrust (Note 3) and pitch (Note 4) to accelerate and retract flaps on schedule. -Very slow power increase: each 750 ft, add 2% N1 and increase pitch so as to maintain CAS till climb power is reached (Note 5) -Then accelerate to 250 kts CAS.

TABLE 7.1: Sourdine case: Sourdine near distance noise abatement departure procedure

Notes

(1) : Possible reduced thrust level calculated from pressure altitude and temperature

(2) : Or maximum pitch angle = 18 degree

(3) : 80% > N1 1.7% gradient One Engine Inoperative (OEI)

(4) : Minimum pitch angle of 10.8 degrees, corresponding to the pitch for a climb at constant speed (V2+10 kt) and N1 < 80% 1.7% gradient OEI

(5) : It takes 5 successive 2% N1 thrust increase steps to go from 80% at 5000ft to climb settings (about 89% at 8000 ft)

Trajectory data for a fast time simulation of an A319 flying the SDClose procedure is provided [161]. The tabulated data provided the parameters for time, distance from brake release, height, speeds in CAS and TAS, climb rate, climb gradient, pitch, angle of attack and aircraft thrust at 285 trajectory points from $t \in [t_0, t_f]$.

As the dynamics model that generated the trajectory is not publicly available, the IDVD-DE method could not be compared directly with the Sourdine trajectory data. To generate valid trajectories for comparison, the Sourdine data was used for reference and the BADA dynamics model was used to generate the baseline trajectory of the Airbus A319 operating the SDClose noise abatement procedure. As The IDVD-DE method uses the BADA drag polars and thrust model as part of the inverse dynamics calculations, valid comparisons could then be made between the baseline and the IDVD-DE optimised trajectories.

In Figures 7.1-7.3 it can be seen that using the BADA dynamics model with the SDClose procedure produces a trajectory very similar to the trajectory generated in the Sourdine study. For the initial climb to 1500 ft, the Sourdine trajectory has a greater available thrust than the BADA generated trajectory (BADA-SDine), yet the climb rate and speeds are broadly the same for both. This indicates that the drag is less for the BADA-SDine trajectory for that phase of flight. The lower BADA drag is also evident for the first acceleration step at 3000 ft, where the BADA-SDine trajectory accelerates quicker than the Sourdine trajectory for approximately the same levels of thrust. For the four steps from 5000ft to 7750 ft, the target height and thrust schedule are specified in the procedure and the pitch is controlled to achieve a continuous linear acceleration across the segments to 7750 ft. For the acceleration segment to 250 kts CAS at 7750 ft, the BADA model again exhibits greater acceleration. However for this segment it seems that the greater acceleration is due to the higher levels of BADA thrust, indicating that the drag for clean configurations may be highly similar in both models. As can be seen from the trajectory results, the BADA dynamics model provides a suitable surrogate for the model used in the original study. As the IDVD-DE method also utilises the BADA model, the BADA simulated procedure can be used as a baseline for comparisons with IDVD-DE optimisation results.

As with the original study, the EPNL under the centreline metric was used to asses the noise impact of the trajectories. To avoid errors inherent in comparing results from two different noise models, the noise results for both the Sourdine and the BADA trajectories were generated by the Integrated Noise Model 7 model [111]. Noise monitoring points were placed from 1000 to 30000 metres from brake release at 1000 m intervals.

For the IDVD-DE optimisation scenario, the final position of the trajectory needed to be extended relative to the BADA trajectory. This was because constraints for the IDVD-DE

simulation prevented instantaneous changes in controls that were possible for the BADA-SDine trajectory. Aerodynamic models, engine model and the performance constraints were otherwise common to both scenarios. As the aim of the procedure is to reduce noise under the centreline from 6-13 km the cost function was chosen to be the Average EPNL recorded at distance 6-13 km.

Comparing the trajectories from the IDVD-DE and the BADA-SDine simulations in Figures 7.1-7.3, it can be seen that that the IDVD-DE thrust profile matches closely that of the BADA-SDine. Both trajectories start with maximum thrust for climb out followed by a thrust reduction over the noise sensitive area. After passing the sensitive area, thrust is increased gradually to maximum climb thrust.

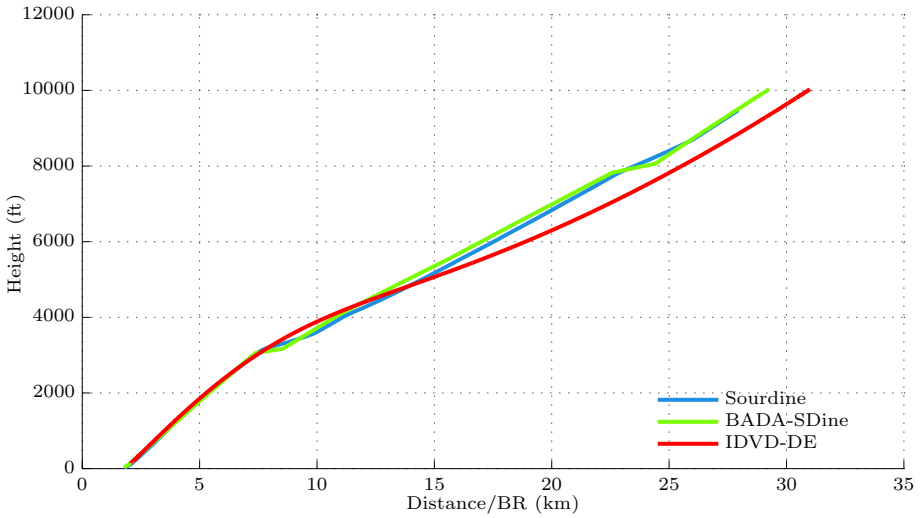

FIGURE 7.1: Sourdine case: Sourdine and IDVD-DE height profile comparisons

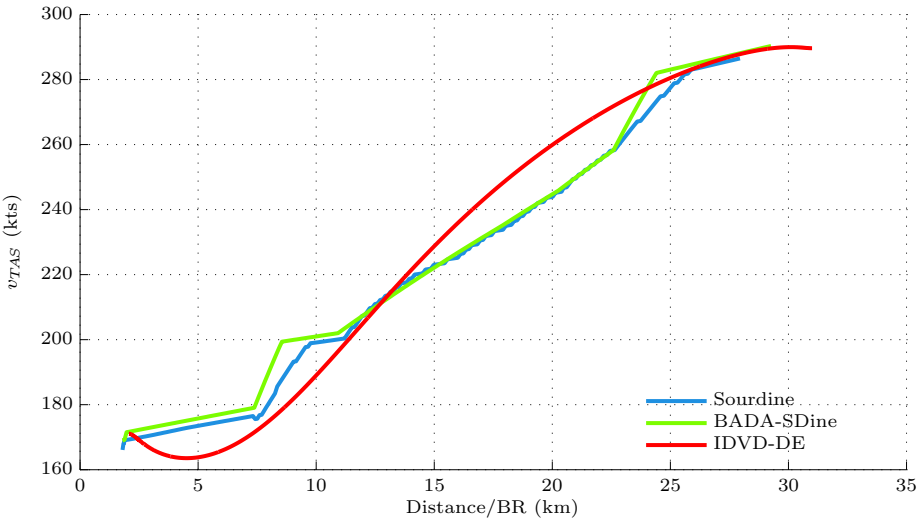

FIGURE 7.2: Sourdine case: Sourdine and IDVD-DE true airspeed profile comparisons

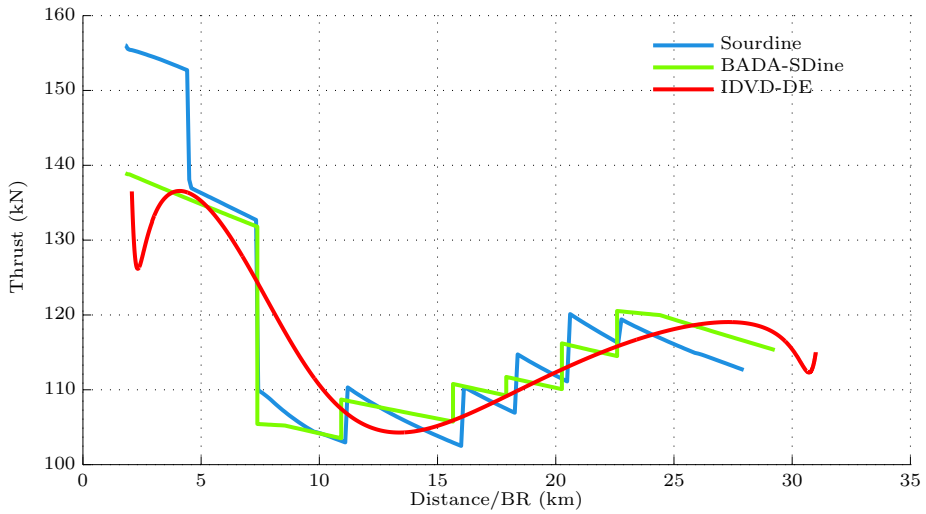

FIGURE 7.3: Sourdine case: Sourdine and IDVD-DE thrust profile comparisons

There are significant differences in the speed profiles generated by the two methods. The IDVD-DEs speed profile over the distances 2-12 km is slower than the speed profile for the BADA-SDine trajectory. For the distances 15-30 km, the opposite is true, and the IDVD-DE speed is higher than the BADA-SDine. As the excess power is approximately the same for both trajectories, it can be seen that the differences between the speed profiles are due to the use of the excess power. For the distances 1-13 km the IDVD-DE trajectory prioritises height gain over speed. After passing the sensitive area the IDVD-DE then prioritises gains in speed to achieve the target speed at the required distance from brake release.

The noise results for the two trajectories are shown in Figure 7.4. It can be seen that the noise results closely match each other. Over the noise sensitive area, the maximum difference between the scenarios was 0.9 EPNdB. The average difference over 15 noise monitoring points composing the noise sensitive area was 0.4 EPNdb. It can also be seen that the different utilisation of excess power by the different methods does not have a significant impact on the noise results.

To test the convergence properties of the IDVD-DE for the problem, the scenario was run 200 times to determine typical variation in the objective value. Figure 7.5 shows a box plot[162] of the results of the convergence testing including the data points. The data points are randomly dispersed to show the results more clearly. The central red bar shows the mean objective value reached, the pink shows the 95% confidence intervals for the mean. The blue box then shows the data within one standard deviation of the mean. It was found that all solutions were within 1% of the lowest objective value found. Therefore the IDVD-DE method showed very good convergence properties for the scenario, especially considering there was averaging involved in the objective function calculation.

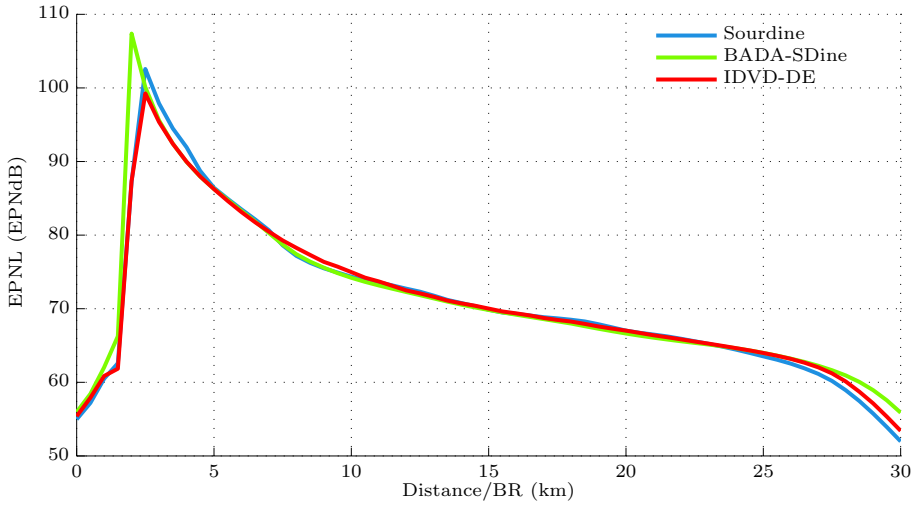

FIGURE 7.4: Sourdine case: Sourdine and IDVD-DE noise under the centreline results

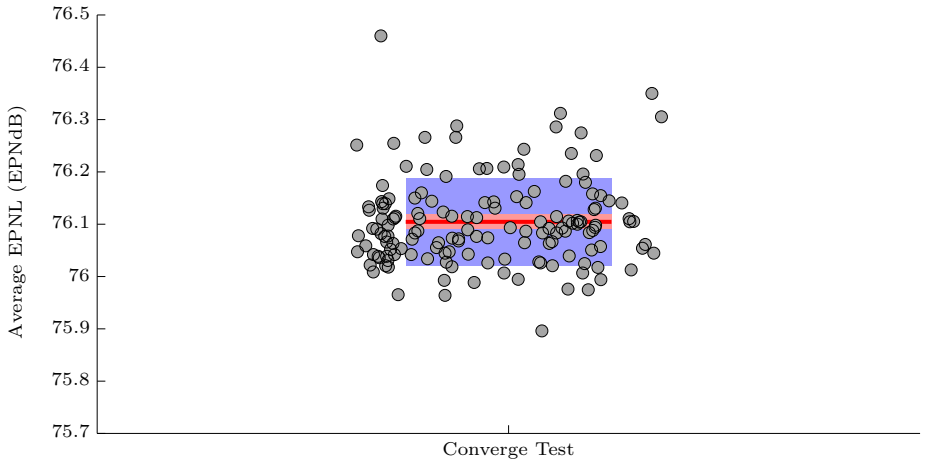

FIGURE 7.5: Sourdine case: IDVD-DE Convergence testing for aircraft departure problem

7.2.1 Summary

It can be seen from the results in this section, for the idealised problem chosen, that the IDVD-DE method did not improve on the noise performance of the noise abatement trajectory produced as part of the Sourdine project. However, the IDVD-DE method did achieve an equivalent performance, consistently converging on a trajectory solution highly similar to the solution developed through expert analysis and extensive testing on the Sourdine project. This provides confidence that the IDVD-DE method can generate environmentally efficient trajectories in line with the best known solutions.

When comparing the two solutions, it should be noted that the Sourdine solution is also a general solution to an idealised problem, and is therefore not tailored to the local conditions at any specific airport. The IDVD-DE method on the other hand, with information on local population and constraints, can generate solutions optimised to local conditions.

The IDVD-DE method, through its suitability for use as a multiobjective trajectory optimisation method, can also consider trade-offs in environmental impacts not considered in the development of the Sourdine noise abatement solutions.

With these observations in mind the IDVD-DE method is now applied to a more realistic multiobjective procedure design problem with specific local conditions and constraints.

7.3 Multiobjective Test Scenario

A departing aircraft scenario was created to demonstrate and test the multi objective trajectory optimisation method. In the scenario, a commercial aircraft is required to climb from an initial climb point below 500 ft at the south west of a large population centre, to an en-route connection point lying at 20,000 ft on the far side of the population centre. The commercial aircraft simulated was the medium narrow-body Airbus A321 aircraft with twin International Aero Engine V2530 engines. The population was artificially created for the scenario, and consisted of 1.5 million people evenly distributed over an area of 45000 hectares. For the scenario, the bank angle ϕ and the minimum climb gradient below 1000 ft were constrained to 0 radians and 12% respectively. For the scenario, the objectives chosen were the greenhouse gas Carbon Dioxide (CO2) and the noise Annoyance Score. The aircraft boundary values for the scenario are shown in Table 7.2.

Time	$x(m)$	$y(m)$	$h(m)$	$v_{\mathtt{t}}(m/s)$	$\gamma_{\mathtt{i}}(rad)$	$\chi_{\mathtt{i}}(rad)$
t_o	9400	2750	30	80	0.13	0.60
t_f	18400	95200	6102	206	0.05	1.56

TABLE 7.2: Test case: Scenario boundary values

Figure 7.6 shows a Pareto front between the minimums of the two objectives. It can be seen from the front that there is a trade-off of approximately 800 kg of CO2 between the most CO2 optimal trajectory and the most noise optimal trajectory. Figure 8.13(a) shows the minimum CO2 trajectory in red, the minimum Annoyance Score trajectory in blue and the intermediate trade-off trajectories in grey.

The aircraft trajectory for the minimum CO2 objective, after climbing out to 1000ft takes a direct route over the population area to the target end point. At 2000ft the aircraft reduces thrust and levels out for acceleration where speed levels are increased quickly relative to the low noise trajectory. Once the acceleration segment is complete the aircraft reintroduces the higher levels of thrust required to climb the aircraft to its target height at the higher flight speed. Overall the lowest CO2 trajectory flies as quickly and directly as possible to the target fix minimising excess track miles, fuel burn and therefore CO2.

The trajectory for the minimum noise objective, initially progresses directly to the east, avoiding over-flying the majority of the population area and therefore minimising the

population exposed to noise. It can be seen from Figure 8.13(a) and from Figure 7.8 that the noise optimised trajectory climbs to a height of 1500 ft where it reduces thrust. Unlike the low fuel burn trajectory the aircraft does not prioritise acceleration but maintains a low thrust and low speed to minimise noise levels while still maintaining speeds over minimum speed levels. On clearing the edge of the population region, higher thrust levels are gradually reintroduced and the aircraft begins a long slower acceleration relative to the low CO2 trajectory. The slow acceleration is supported by the long path that maximises distance away from the population by maintaining a easterly heading for as long as possible before then turning back to the final fix. The trajectory produced mimics closely the Sourdine close-in noise abatement procedure as both involve an initial climb at full thrust followed by acceleration at reduced thrust and a gradual power increase. The Sourdine project was a leading project in the field of noise abatement trajectory operations [67]. Further information on the Sourdine solution can be found in Section 7.2.

Examining the Pareto front in Figure 7.6, it can be seen that at Annoyance Score values below 1×10^{13} and CO2 values above 3800 kg, that the Pareto front flattens out to an almost horizontal line. This shows that for increasingly larger values of CO2 emissions, there is only marginal reductions in the noise Annoyance Score. The reasons for this can be seen from the Pareto front trajectories shown in Figure 8.13(a). It can be seen that there is a noise benefit from flying the aircraft away from the population region. Trajectory paths that arc away from the population maximise the distance between the aircraft and the population, creating larger distances for the noise to attenuate over. However, as the arcs get bigger and bigger eventually the noise benefit delivered by increasing attenuation distances gets smaller and smaller, the Pareto curve plateaus and further reductions in noise are only achieved through larger and larger jumps in CO2 values.

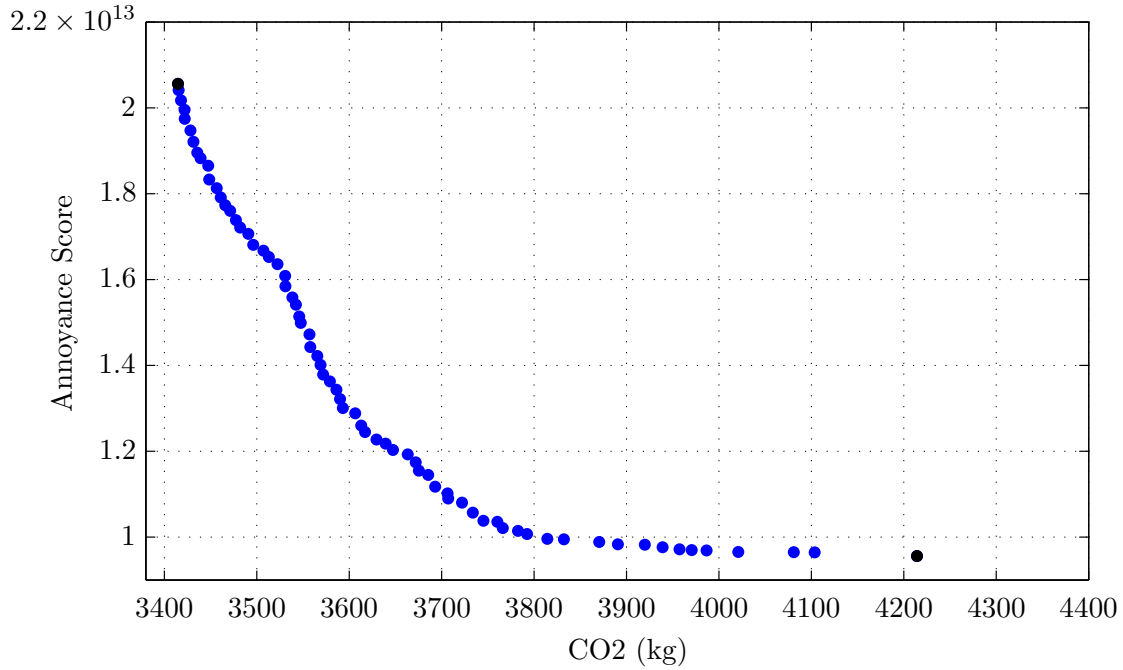

FIGURE 7.6: Test case: Pareto front between CO2 emissions and noise Annoyance Score

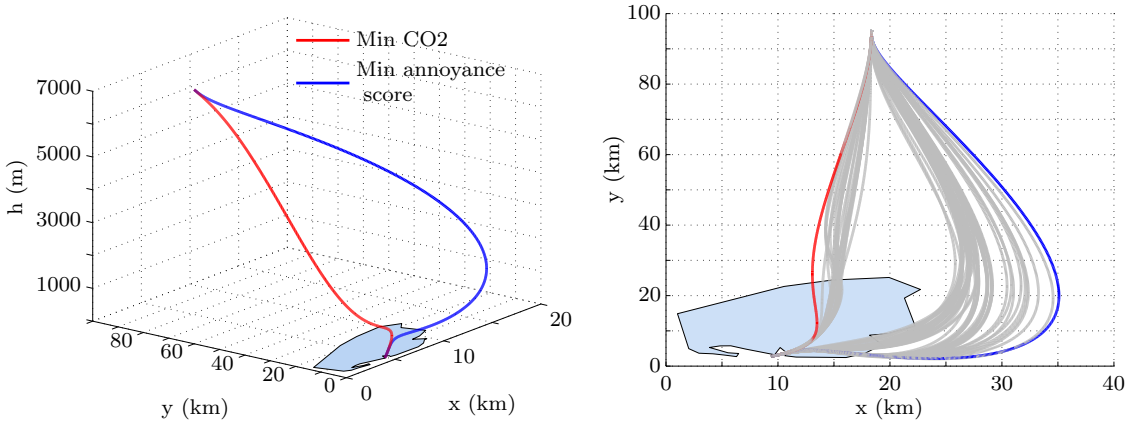

(a) Minima trajectories

(b) All Pareto front trajectories

FIGURE 7.7: Test case: Pareto front trajectories

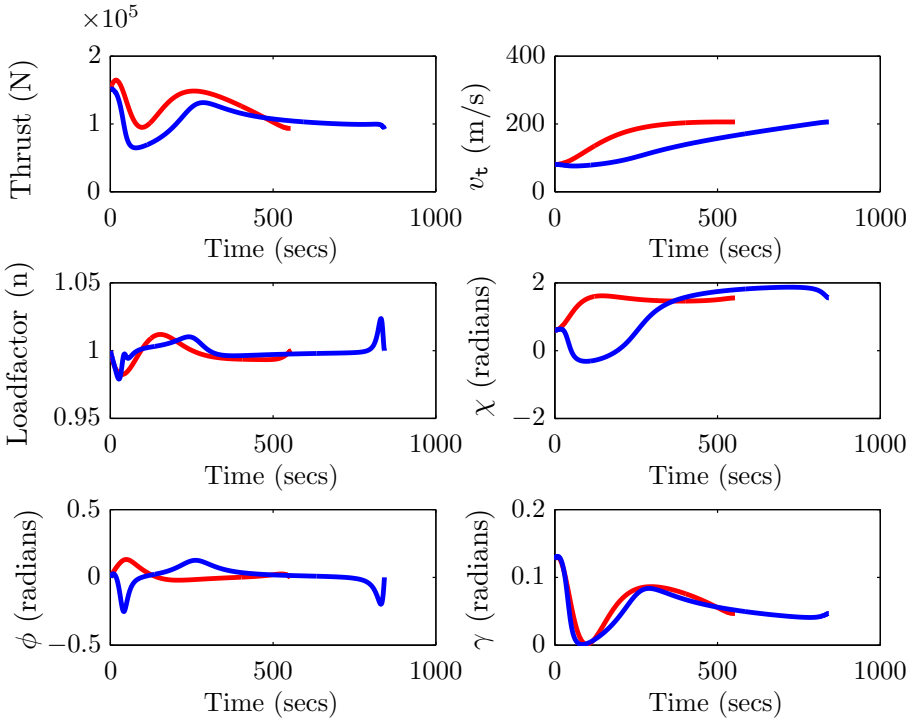

FIGURE 7.8: Test case: Controls and further states for the minima trajectories

7.4 Pareto Front Clustering

Post optimisation, there is a need to consider not only the trajectories that are the minimum for each objective but also the trade-off trajectories whose solutions lie along the Pareto front and offer a balance between the extrema solutions.

To assist in the analysis of the Pareto front, a clustering algorithm was created. The aim of the algorithm was to identify and group closely related solutions. The principle of the algorithm is that closely related solutions have trajectories with similar characteristics, and that this can be used to determine where distinct trajectory behaviour leads to distinct trade-offs between the objectives. In the first instance the algorithm clusters the points on the Pareto front to the extrema solution that each point is closest to. The algorithm will then go on to find the transition points between each cluster, creating transition clusters around those points. The number of transition clusters is user-defined.

Figure 7.9 shows how the algorithm would cluster a 2D pareto front with two conflicting objectives, f1 and f2. The solutions whose distances are closest to the extrema solution A are grouped together as Cluster A. The solutions that are closest to the extrema solution B, are grouped as Cluster B. The distances between each Pareto point and the extrema solutions are calculated using Mahalanobis distance. Mahalanobis distance is a distance measure popular in clustering algorithms as it is insensitive to large differences in the scales between different axes as it accounts for the covariance of multiple objectives [163]. The Mahalanobis distance between two coordinate points $\boldsymbol{x}$ and $\boldsymbol{y}$ is calculated as

$$d_{\mathcal{M}}(\mathbf{x}, \mathbf{y}) = \sqrt{\sum_{N}^{i=1} \frac{(\mathbf{x}_i - \mathbf{y}_i)^2}{sd_i^2}} \tag{7.1}$$

where N is the length of the coordinate vectors and sd_i is the standard deviation of $\mathbf{x}_i$ and $\mathbf{y}_i$ over the sample set.

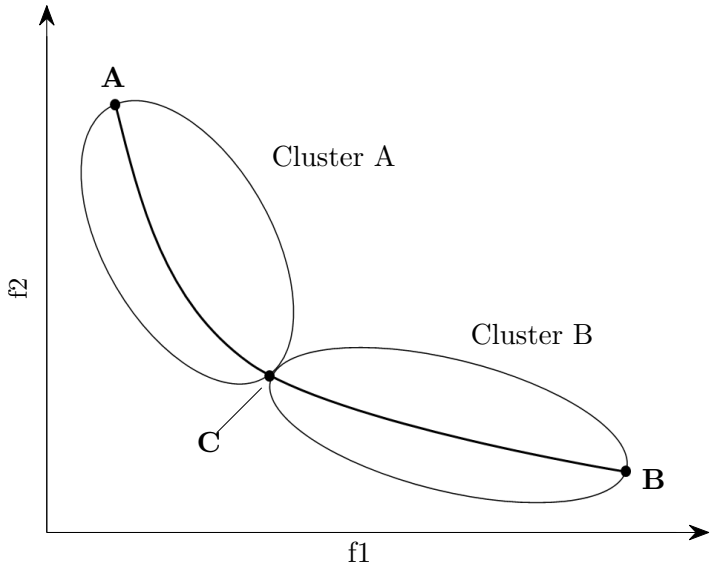

FIGURE 7.9: Pareto front clustering to minima solutions

Although it is useful to cluster trajectories relative to the minima each trajectories objective vector is closest to, It may also be useful to add additional transition clusters identifying transitionary Pareto points where there are notable shifts in the trade-offs between the objectives. The creation of a transition cluster can be seen in Figure 7.10. In Figure 7.10, a new centroid point C is created that is the average of the values of the 2 closest Pareto points between clusters A and B. The Mahalanobis Distance is then calculated between the 3 centroid points A,B,C and all solutions on the Pareto front. The Pareto solutions that have distances closest to each centroid are then clustered to that centroid, altering the extents of Clusters A and B and creating a third cluster, cluster C.

Providing there are sufficient Pareto points, the process of recursive clustering can be continually repeated. In Figure 7.11 the closest points between cluster A and cluster C and then between cluster C and cluster B are identified and 2 new centroid points labeled points D and E are then created. As before, solutions are clustered relative to the centroid that they are closest to. The number of clustering recursions is defined by the analyst. Although, the example is shown for a 2D Pareto front, the method is similarly applied to higher dimensional Pareto fronts.

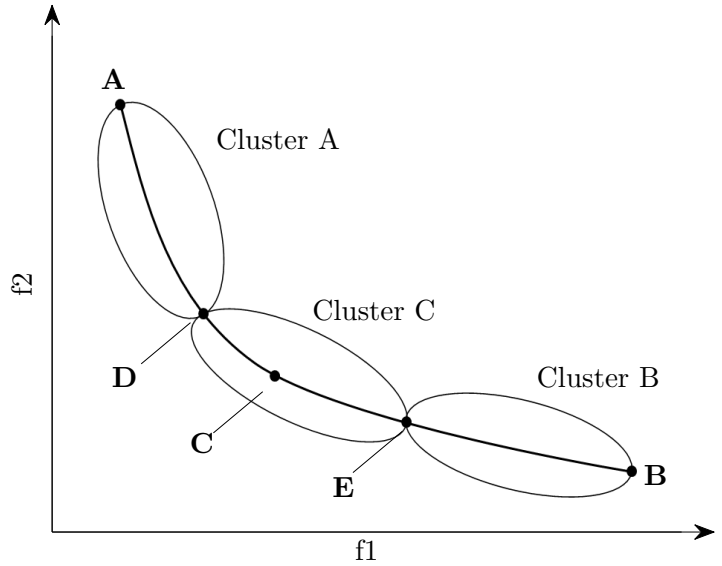

FIGURE 7.10: Pareto front clustering including transition clusters

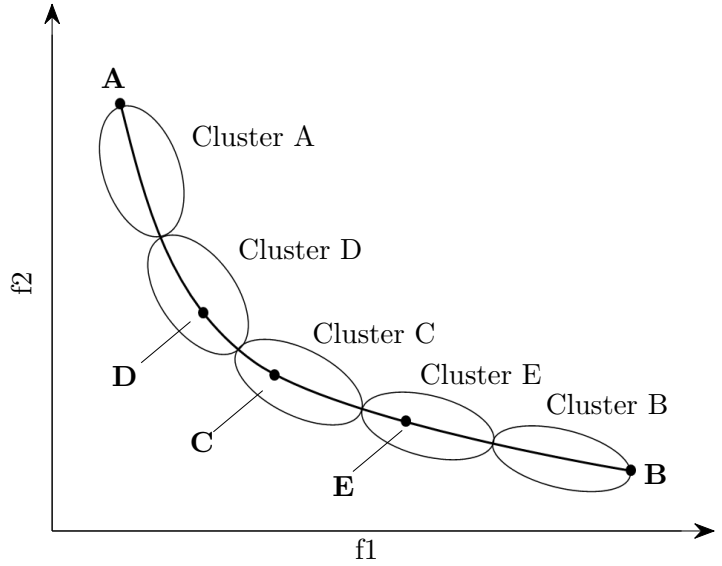

FIGURE 7.11: Pareto front clustering with further transition clusters

7.5 Clustering Example

Applying the clustering algorithm to the scenario in this section allows a more complete analysis of the scenario. Figure 7.12 shows the Pareto front for the departing aircraft scenario with the clustering algorithm applied. The algorithm has partitioned the Pareto front into 3 colour coded clusters. The Pareto points in the red cluster represent the solutions with the lowest fuel burn values, the points in the blue cluster represent the solutions with the lowest noise Annoyance Score values, and the orange cluster contains transition points between the minima clusters. Figures 7.13-7.15 show the trajectory paths, height, speed, thrust, flight path angle, heading angle and roll angle time histories for each of the Pareto front solutions. As can be seen in the figures all the trajectory information is colour coded as per the Pareto front clustering. Breaking down and analysing

the trajectories by cluster.

Red cluster: The trajectories in the red cluster represent the trajectories with the lowest fuel burn values. Trajectories in this cluster perform an early turn to take the shortest paths to the final fix. For the vertical profile, the aircraft climb to approx 3000ft and level off for an acceleration segment allowing all excess power to be committed to acceleration. After the acceleration segment, higher levels of thrust are reintroduced and the climb is resumed at expedited climb rates.

Orange cluster: The orange cluster is an intermediate transition cluster bridging the low noise and low fuel burn clusters. The trajectory paths provide a compromise between the low noise and low fuel burn trajectories by neither flying over, nor far away from the population region, but by flying arc paths that stay close to the edge of the population region thereby minimising noise impact in addition to minimising the excess path mileage that would contribute to excess fuel burn. The trajectories in the orange cluster have on average higher climb rates and speed profiles, similar to those of the low fuel burn cluster. Therefore, it can be seen that the reductions in noise are achieved by routing trajectory paths away from the population, but that the longer paths are flown quickly to minimise fuel burn. As the orange cluster trajectories approach the low noise cluster, the paths are flown slower and thrust levels are reduced as noise starts to be prioritised over fuel burn.

Blue cluster: The blue cluster represents the lowest noise trajectories. The trajectory paths in this cluster initially progress directly to the east, flying adjacent to the population region at low cutback levels of thrust. On clearing the edge of the population region, higher thrust levels are gradually reintroduced. The path arcs for this cluster take aircraft further and further away from the population region. On average the climb rates and speeds are lower than those in the other clusters and the aircraft fly a longer, slower ascent. Extending path arcs away from the population serves to successfully deliver reductions in noise annoyance. However, the Pareto front shows that there is a point of dimishing returns at CO2 values 38500 kg and upwards. At this point, due to noise attenuation, ever larger arcs are required to realise smaller and smaller reductions in noise annoyance.

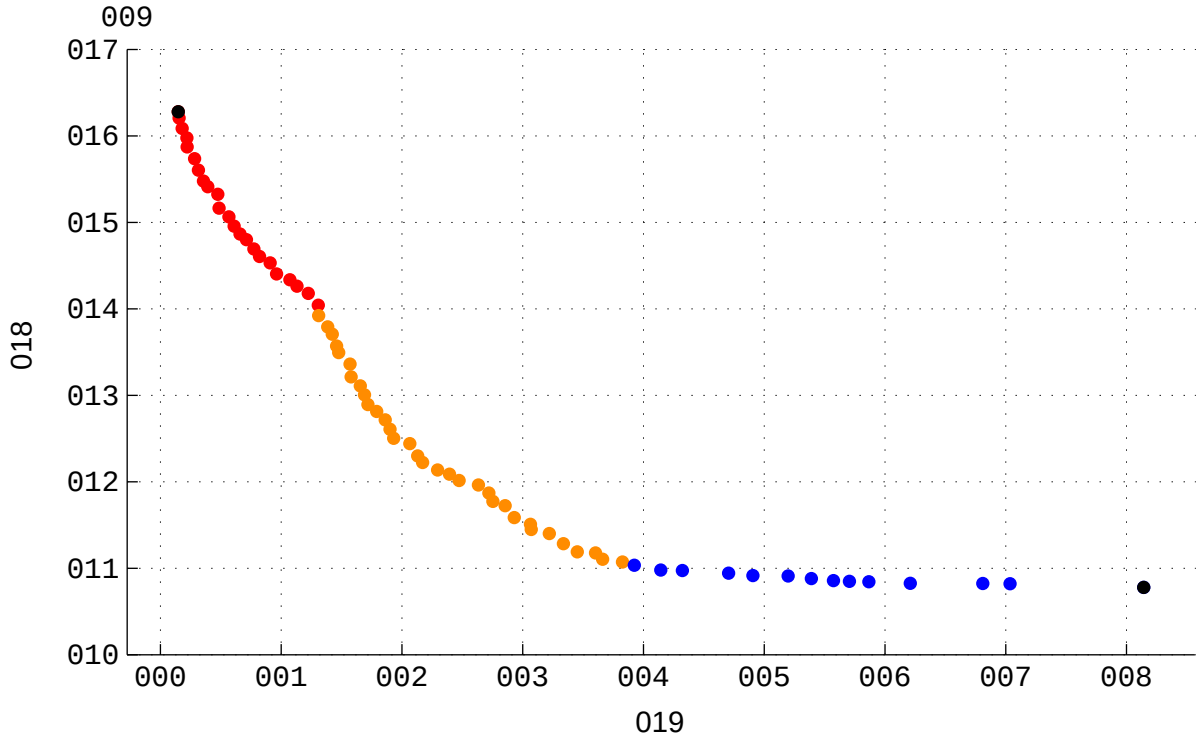

FIGURE 7.12: Test case: Clustered Pareto front between CO2 emissions and Annoyance Score

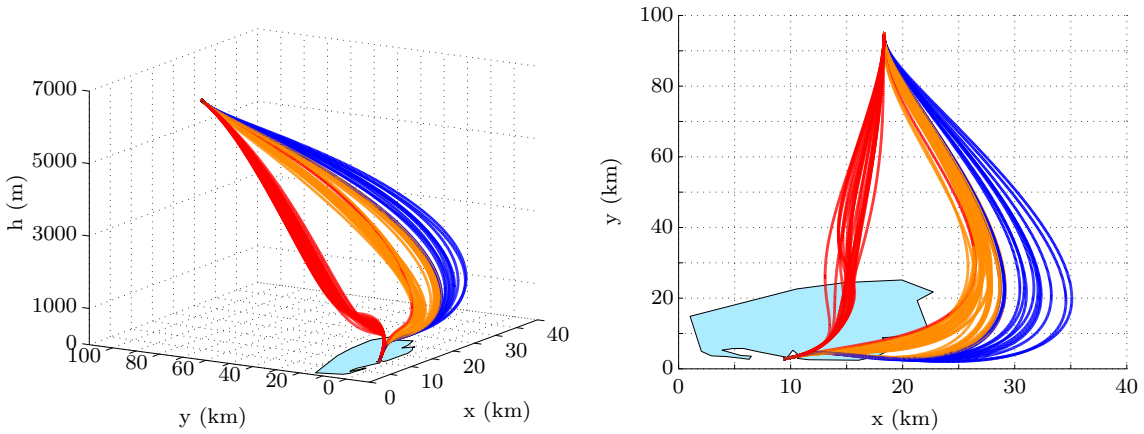

(a) Clustered 3D trajectories

(b) Clustered plan view trajectories

FIGURE 7.13: Test case: Clustered Pareto front trajectories

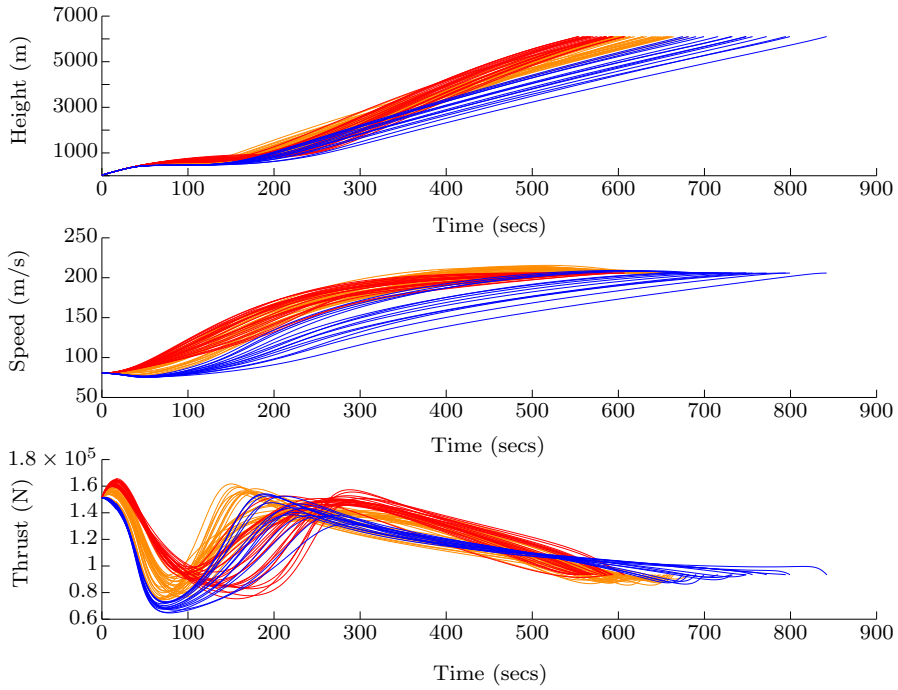

FIGURE 7.14: Test case: Height speed and thrust profiles for the clustered Pareto front

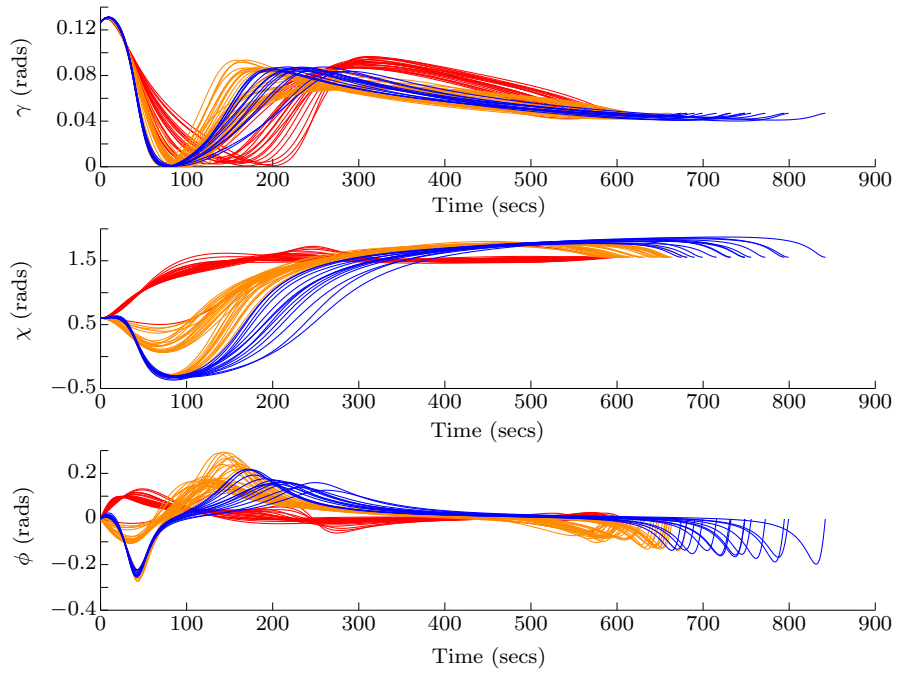

FIGURE 7.15: Test case: γ, χ, ϕ, profiles for the clustered Pareto front

7.6 Convergence Testing

To assess the performance of the evolutionary algorithms when applied to a multi-objective environmental optimisation problem, the departing aircraft scenario was used once more. For the tests, a multiobjective extension to the relative optimality tests performed in [79] was used. First an estimated global Pareto front was generated. This was done by running the departing aircraft scenario optimisation 50 times. This produced 50 Pareto fronts or 3500 Pareto solution points. All the Pareto solution points were added to a single set, where any dominated points were removed, creating a pseudo global Pareto front.

A further 50 scenario simulations were performed and the convergence of each front to the pseudo global front was measured along with the diversity of each front relative to the pseudo global front. In all cases the DEMO algorithm was used with a scale factor of 0.8 and a crossover constant of 0.7. The convergence performance of the IDVD-DE method for the multiobjective scenario can also be visualised in Figures 7.16 and 7.17.

Figure 7.16 shows the global Pareto front for the Annoyance Score and CO2 measures. Plotted relative to this are range bars that show the convergence distribution from the 50 individual runs. The range bars show the min, max and mean Annoyance Score values from the 50 converged fronts at a series of representative CO2 values.

Figure 7.17 shows the convergence density for the 50 individual runs relative to the same global front. For the convergence density plot, the regions with the highest convergence density are shown in red and the regions with the lowest convergence density are shown in dark blue

In general, there is good convergence from the 50 runs to the global front. Analysing Figures 7.16 and 7.17 it can be seen that convergence on the minimum CO2 solution is very good, there is a high density of converged points at the minimum CO2 value and there is very little variation in the corresponding Annoyance Score value relative to the global front.

Progressing along the Pareto front, the region between the CO2 values of 3400 and 3500 kg similarly has a high density of points converged to the global front. Just after the CO2 value of 3500kg, there is a kink in the global front. Referring back to the clustering in Figure 7.12, it can be seen that the change in the slope of the global front occurring in this region approximately coincides with the beginning of the orange transition cluster, where aircraft begin to take longer arc paths around the population centre. At this trajectory behaviour transition region there is a reduction in the convergence density, marked by the orange and yellow convergence density coloring. Over the same region in Figure 7.16 it can be seen that there is a corresponding increase in the spread of the converged Annoyance Score values relative to the global front. The red range bar in this region (at a CO2 value of 3500 kg) shows that there is a maximum difference between the Annoyance

Score values from the individual runs and the Annoyance Score values of the Pareto front of 7 %. However, the mean difference in the Annoyance Scores at the comparable CO2 values is less than 2 %. The reduction in the convergence density and the greater spread in the converged Pareto front values from the individual runs shows that there is some variation in where the trajectory behaviour changes between simulation runs.

Examining the global front over the CO2 range 3550 to 3700 kg shows a second region with high convergence density to the global front. Further along the front, over the CO2 region 3700 to 3900 kg, the global front contains solutions that prioritise noise reduction over CO2. In this region it can be seen that there is both a drop in convergence density and an increase in the spread of the converged annoyance score values relative to the global front. The maximum difference in the Annoyance Score values in this region from the global front is 9 %, although the mean difference remains low at less than 2.5 %.

Along the CO2 value region 3900 to 4300 kg, the global Pareto curve flattens out and there is very little reduction in the Annoyance Score for increased values of CO2 emissions. Over this region in Figure 7.16 it can be seen that the mean converged Annoyance Scores lie with very little variation either on or close to the global front. However, in Figure 7.17 it can also be seen that the density of the converged solutions on the global front in this region is very poor, especially over the range 4000 to 4300 kg CO2, thereby showing that some simulations had difficulty converging on this part of the Pareto front. It can also be seen that the global front is sparse over this same region. Trajectories for this region of the Pareto front are pushed by the optimisation method to fly ever larger arc paths in search of ever smaller improvements in noise annoyance. However, the maximum difference between the Annoyance Score values of the global Pareto front in this region is less than one percent. Therefore solutions that are nondominated in this region are the result of more extreme trajectories that exist in a narrow region of the solution space. As there is such small differences between the Annoyance Score values in this region, the algorithm has a tendency to converge to where the majority of the solutions are, which is to less extreme trajectories with virtually identical Annoyance Score values but with lower CO2 values.

In addition to the visualisations in Figures 7.16 and 7.17, the multi objective performance measures Generational Distance (*GD*), Hypervolume (*HV*), Additive-ϵ ($I_{\epsilon+}$), Spread (*SPD*), Spacing (*SPC*) and Maximum Spread (SPD_M) from [164] were used to assess the convergence and diversity of the 50 individual runs relative to the pseudo global Pareto front. The results are shown in Table 7.3

-	*GD*	*HV*	$I_{\epsilon+}$	*SPD*	*SPC*	SPD_M
Mean	0.0027865	0.564	0.040837	0.46928	0.013814	1.1209
StDev	0.0026106	0.1367	0.02598	0.081403	0.0063941	0.11818

TABLE 7.3: Test case: Scenario convergence and diversity performance measures

The metrics are discussed in detail in Appendix A, but in general *GD* and epsilon provide measures of convergence. The metrics Spread and Maximum Spread provide measures of diversity, and *HV* provides a composite measure of convergence and diversity. A value of $GD = 0$ indicates that all the generated elements are on the Pareto front and low values of *GD* are desired. The *HV* indicator calculates the volume, in the objective space, covered by members of a non-dominated set of solutions and larger values of *HV* are desirable. $I_{\epsilon+}$ is a measure of the smallest distance one would need to translate every solution in an nondominated set so that it dominates the pseudo-optimal global Pareto front.

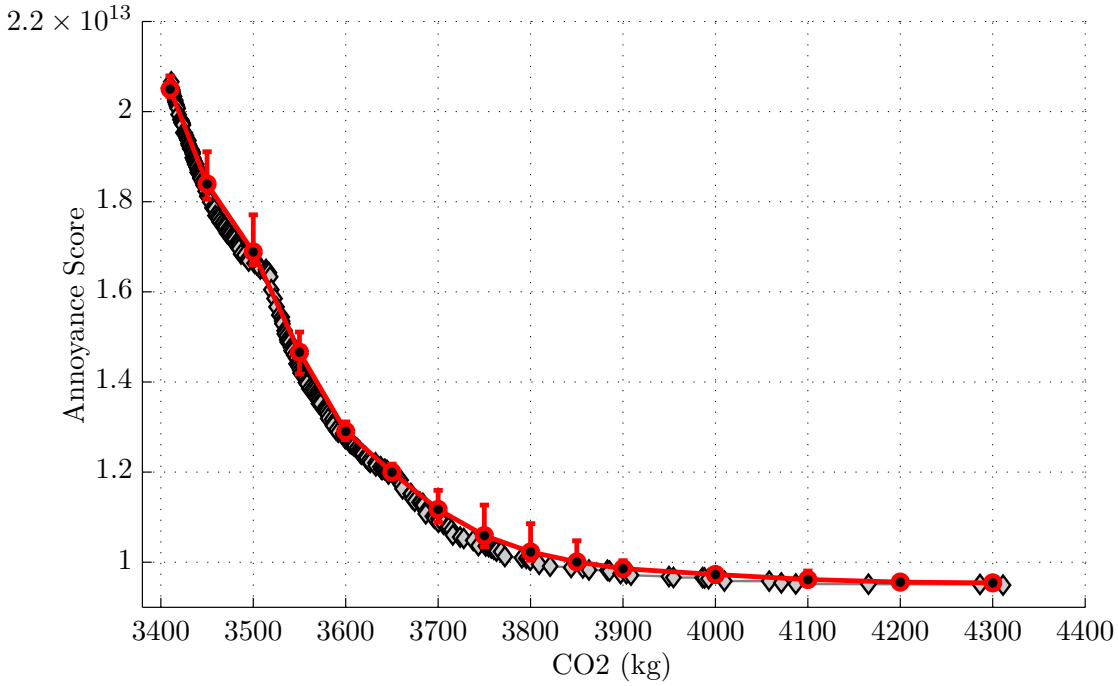

FIGURE 7.16: Test case: Solution convergence distribution relative to the global Pareto front

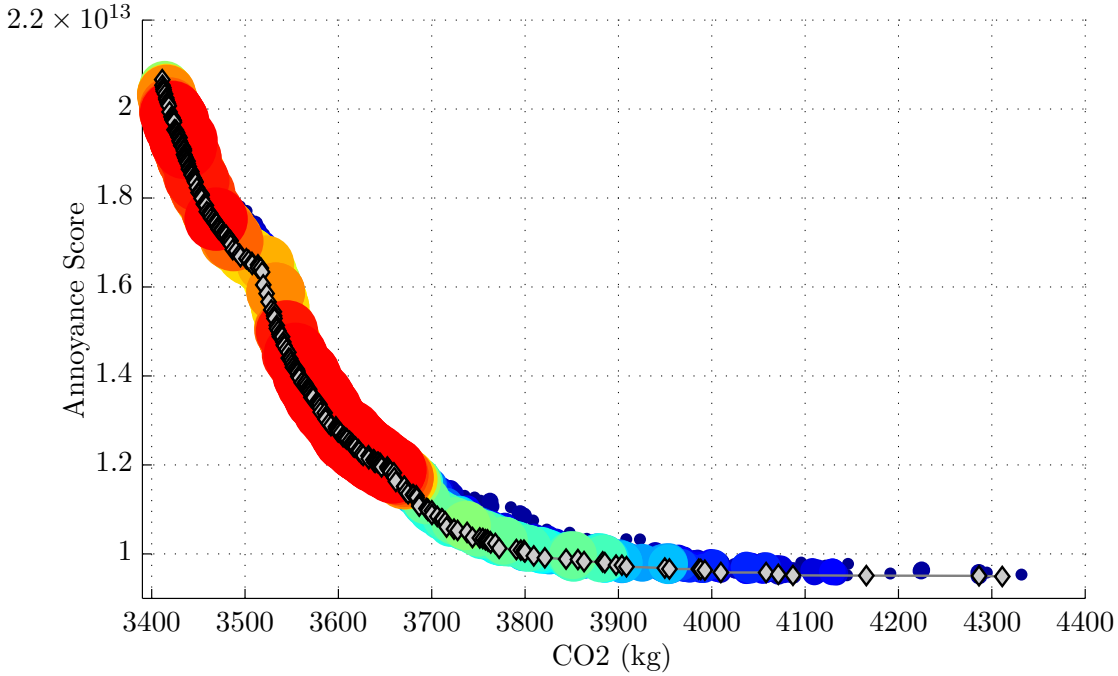

FIGURE 7.17: Test case: Solution convergence density plot relative to the global Pareto front

7.6.1 Summary

In summary the IDVD-DE method performed very well in identifying the correct trade-offs between objectives. The algorithm had some difficulty identifying the lowest noise solution but in all cases identified a low noise solution within 3 % of the global value. Improvements in the convergence are thought to be possible through improvements to the experiment design such as constraints on the extremity of the trajectories, adjustment of the DE settings and more granular population data. However the results were sufficiently encouraging to carry the IDVD-DE method forward for application to further scenarios including more complex and realistic multi-objective scenarios.

Chapter 8

Multi-Objective Environmental Procedure Optimisation

8.1 Introduction

In this chapter IDVD-DE method is applied to two real world environmental procedure optimisation case studies. Previous multi-environmental-objective trajectory optimisation studies in the literature only considered very idealised scenarios. The use of real world case studies in this thesis set demanding requirements on the number and type of environmental objectives that needed to be considered in the trade-off analysis. By applying the IDVD-DE method to real world case studies it can be determined whether the proposed data driven approach could identify procedures or procedure characteristics that provide better trade-offs between the environmental impacts than those proposed by current best practice approaches.

8.2 NEMA Arrival Procedure Definition Case Study

As discussed in Section 2.6.1, Continuous Descent Approaches (CDAs) have been shown to be an effective operational measure for reducing environmental impact from arriving aircraft. As part of the Silent Aircraft Initiative (SAI) research study, low noise arrival procedures were developed and implemented at Nottingham and East Midlands Airport (NEMA) [165]. The goals of the study were:

- to use real and fast time simulation to develop low noise approach procedures,
- to implement the procedures as part of an operational flight trial,
- to asses the operational performance of the procedures,

- to asses the environmental benefits associated with implementing the CDA procedures at NEMA.

As part of the study, two P-RNAV (Precision Area Navigation) CDA approach procedures, NEMAX1A and NEMAX1B, were designed, implemented and subject to flight trials at NEMA. CDA performance was assessed by the amount of level flight below 9000 ft relative to the non trial flights for each individual aircraft type. Environmental impact was assessed by measuring aircraft fuel burn provided by FDR data and by modelling and recording peak noise under the aircraft flight path. The results of the study showed that relative to stepped descent approach trajectories, which were trajectories for non-trial flights with data recorded over the trial period, the trial flights operating the NEMAX procedures increased CDA performance while also reducing fuel burn and peak noise under the aircraft flight path.

In this chapter, the NEMA procedure design study was taken and used to form the basis of a Multiobjective trajectory optimisation study for the IDVD-DE method. Adopting the NEMA study to form the basis of an optimisation scenario offered a number of benefits. The benefits include,

- definition of airspace constraints,
- definition of scenario objectives,
- baseline trajectory results.

Therefore the NEMA scenario presented a case study with a well-defined problem and known solutions. By making the problem the subject of an optimisation case study it was then possible to determine if the IDVD-DE method could identify the known solutions, or suggest alternative or better solutions. In particular a goal of the optimisation study was to determine how knowledge of the entire Pareto front would impact the proposed procedure design.

8.2.1 Setup

Figure 8.1 shows the southerly approach zone for NEMA overlaid on a population density map of NEMA and surrounding population areas. It was the desire of the airport operators that any approach procedures be planned to be within the lateral extents of the approach zone [165]. For the optimisation study, cylindrical path crossing constraints, shown in red in Figure 8.1, were used to approximate the approach zone. The baseline procedure route NEMAX1A and associated waypoints are shown in pink. In the baseline SAI study [165], peak noise was assessed underneath the flight path at the waypoint positions. In the optimisation study noise monitors were also placed an the waypoint positions. However, as the optimization allowed the flight path to vary within the approach

zone, further noise monitors were placed to the left and the right of waypoint monitors such that peak noise under the flight path could be assessed irregardless of where the optimisation method varied the trajectory paths. Figure 8.2 shows the sectorization for the controlled airspace surrounding NEMA. The flight levels available to the approaching aircraft are defined by the vertical extents of the sectors. Figure 8.3 shows the vertical extents of the cylindrical path constraints that were used to approximate the sector flight level limits. For approaching aircraft the standard speed limits are,

- 250 knots IAS below 10000ft,
- 220 knots IAS less than 15 NM from touch down,
- 180 knots IAS less than 12 NM from touch down,
- 160 knots IAS 4 DME (4 nautical miles from touchdown measured by Distance Measuring Equipment).

Figure 8.4 shows the geo-referenced speed constraints for the optimisation study. As can be seen, the 220 knot and the 180 knot constraints are located close to each other, meaning that trajectories in the optimisation study had the option of of meeting the 220 knot constraint earlier or later along the flight path dependant on the how the change in operation effected the objective function. The upper and lower extents of the speed constraints were defined using the trial FDR data from [165]. The boundary values for the IDVD-DE simulation were taken to match the baseline trajectories.

The metrics included in the baseline study were fuel burn and peak noise under the flight path. Both peak noise and fuel burn were therefore used as objectives in the optimisation study. Peak noise under the flight path was represented by calculating the L_{Amax} value at each noise monitor shown in Figure 8.1 and averaging the values over all the noise monitors. The noise Annoyance Score was also included as an extra objective so that overall community noise could be assessed in addition to noise under the flight path.

For the operational trial of the NEMA approach procedures, the two aircraft types most frequently operating the procedure were the Boeing B757-200F and the MD11F, accounting for 94% of the southerly trial flights movements. For trajectory based procedure optimisation studies, it is standard to optimise the procedure for the aircraft type that will most frequently utilise the procedure [69, 106, 107, 108]. However, in this case there were two dominant aircraft types, with the B757-200F accounting for 58% of movement and the MD11F accounting for 36% of movements. Therefore. for the IDVD-DE optimisation study, Pareto fronts and related trajectory information were generated for both the B744 and B752 aircraft types. For the fast time simulations of the procedures, as in the initial study, the MD11F was modelled using a surrogate 747-400 aircraft.

In the baseline study, the approach routes were developed with consideration of a number of wind conditions. The formula used for wind speeds v_w (m/s) was defined in the baseline study as

$$v_w = 0.51444 \left(\frac{h_{ft}}{1000} + 40\text{kts} \right) \tag{8.1}$$

where h_{ft} is the height in feet. Trajectories in the baseline study were simulated for three wind directions, head wind along each leg of the procedure, tail wind along each leg of the procedure and nominal wind direction. For nominal winds, the wind direction was set to be constantly from 270 degrees, which was consistent with the prevailing conditions at NEMA when the airport was operating runway 27 approaches. For the optimisation study, the Pareto fronts were generated using wind speeds defined by (8.1) and the nominal wind direction. Proposed changes to the arrival route, resulting from the analysis to the Pareto front, were then validated for all three wind conditions.

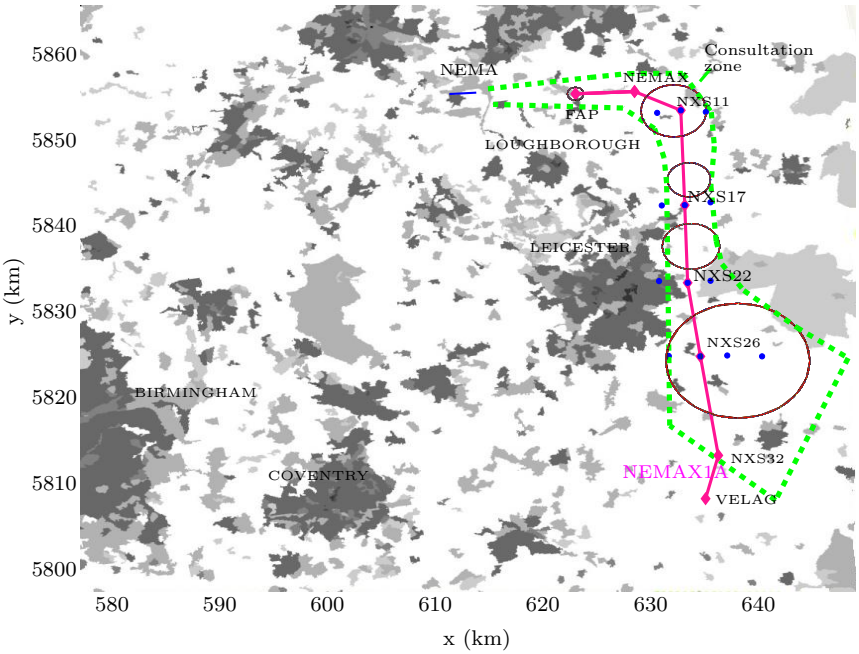

FIGURE 8.1: NEMA consultation zone

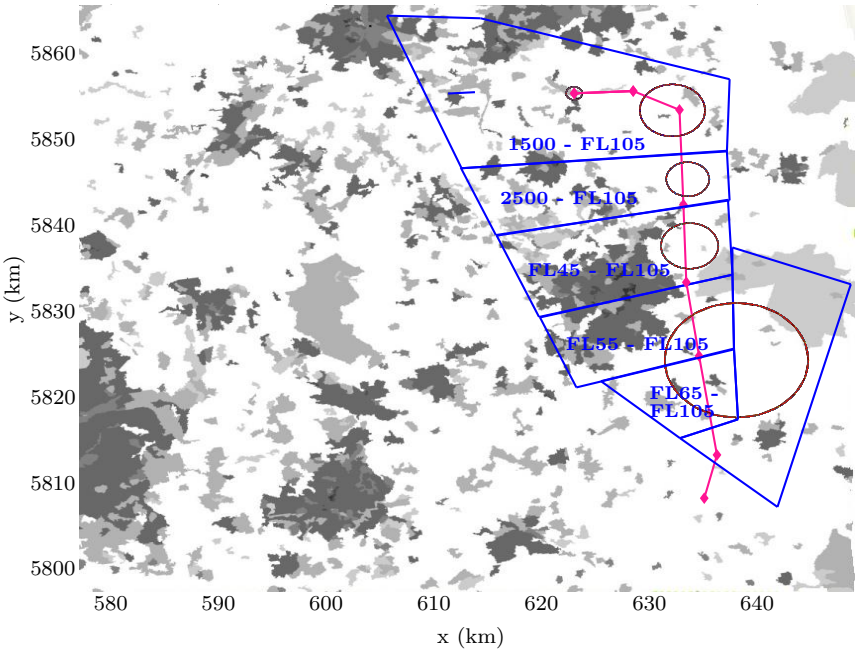

FIGURE 8.2: NEMA sectorisation

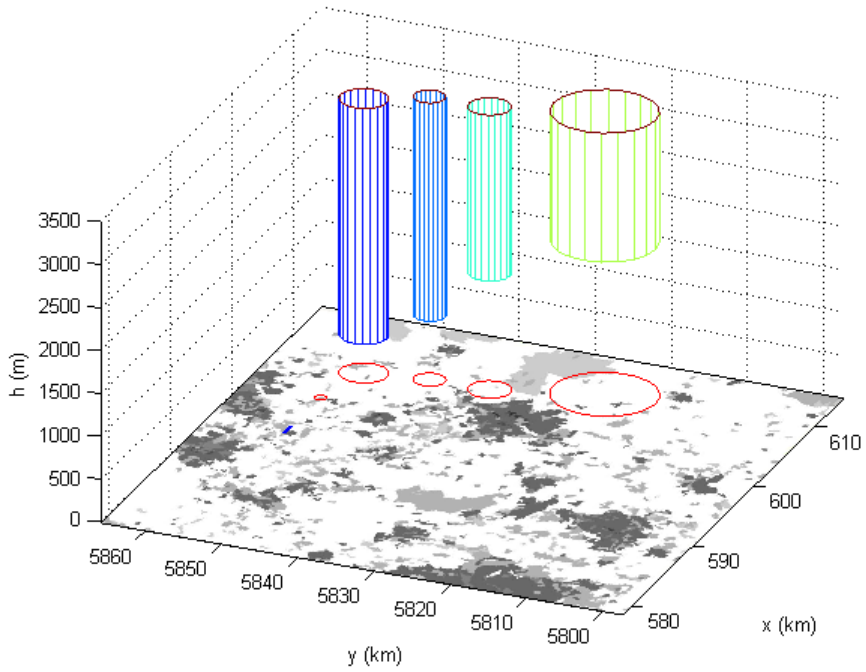

FIGURE 8.3: NEMA case: Optimization height and path constraints

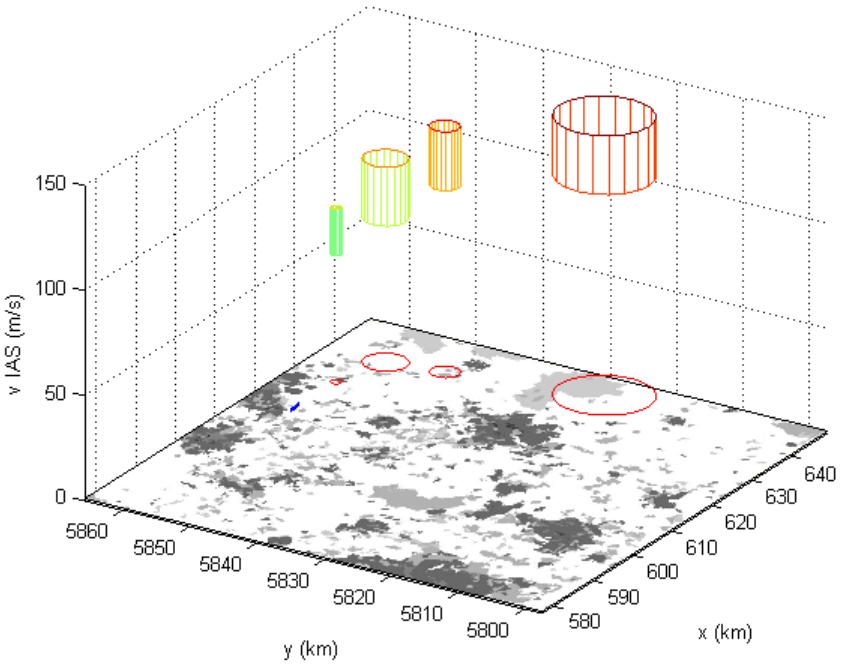

FIGURE 8.4: NEMA case: Optimisation approach speed constraints

8.2.2 B752 Optimisation Scenario Results

Figure 8.5 shows the three objective Pareto front for the B752 aircraft for the Annoyance Score, peak noise under the centreline and fuel burn measures. The Pareto front in Figure 8.5 is clustered into six individual clusters. A minima cluster for each objective and four transition clusters. The dark blue cluster contains the solutions clustered to the minimum L_{Amax} centreline solution. The light blue cluster sits between the low L_{Amax} centreline cluster and the minimum fuel burn cluster and offers trade-offs in the three objectives between the two adjacent clusters. The solutions in the orange cluster have lower Annoyance Scores than the those those in the red cluster, but only at the expense of higher peak noise values. The green cluster is a compromise cluster that doesn't have the lowest of any single objectives but has solutions with low Annoyance Score values that are achieved with relatively small increases in fuel and peak noise from the minima fuel and minima L_{Amax} centreline clusters respectively. The violet cluster shows the solutions that are clustered to the minimum Annoyance Score solution, however these solutions come at increasingly high fuel burn values.

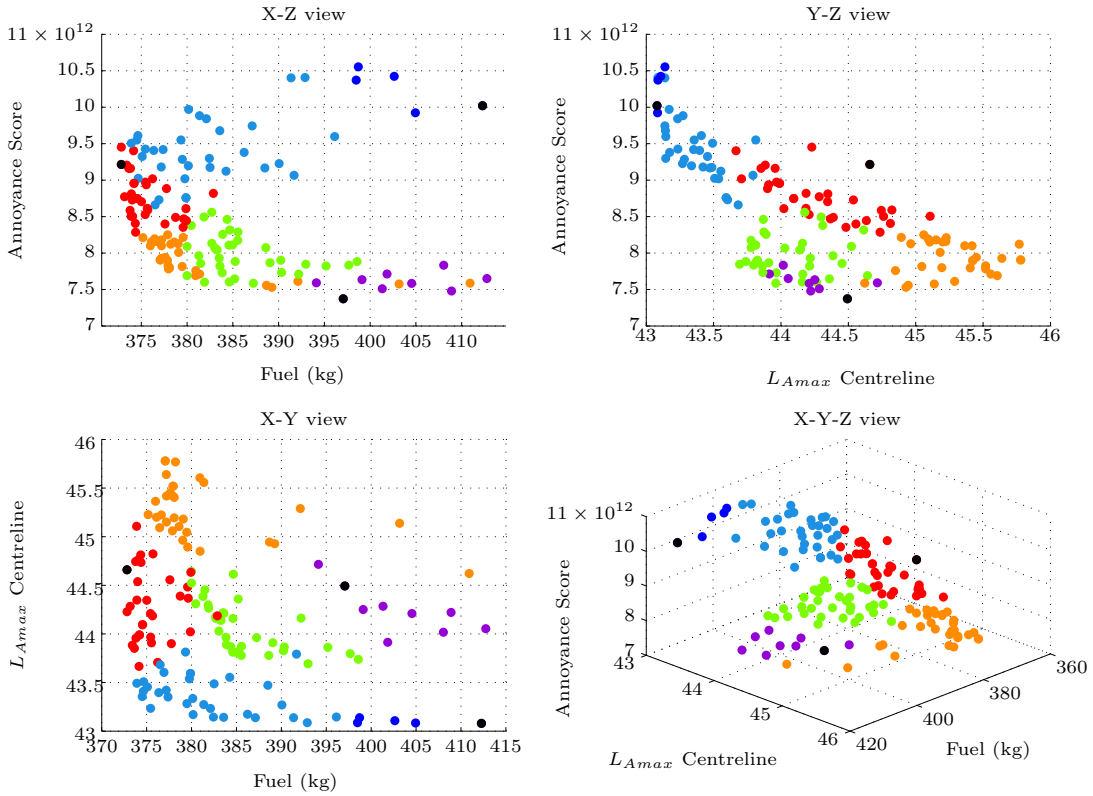

FIGURE 8.5: NEMA case: Four view B752 Pareto front

Figures 8.6 and 8.7 then show the states and controls for the trajectories of solutions with the minima value of each objective. The most obvious differences between the trajectories is in their trajectory paths. The lowest L_{Amax} trajectory is shown in blue and takes a path that pushes against the airspace constraints and maximises the distance from the aircraft to the noise monitors. The minimum fuel burn trajectory takes the shortest path of the three trajectories and also has the shortest flight time. The minimum Annoyance Score

trajectory, shown in violet, flies a slower, longer arcing path that maximises distance from the aircraft path to the city of Leicester and to a lesser extent, Loughborough.

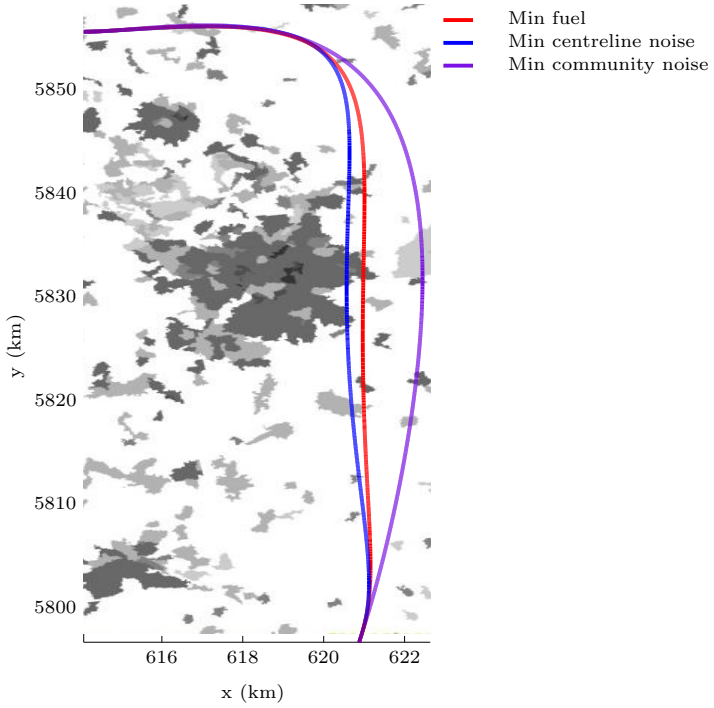

FIGURE 8.6: NEMA case: Minima trajectories for B752 Pareto front

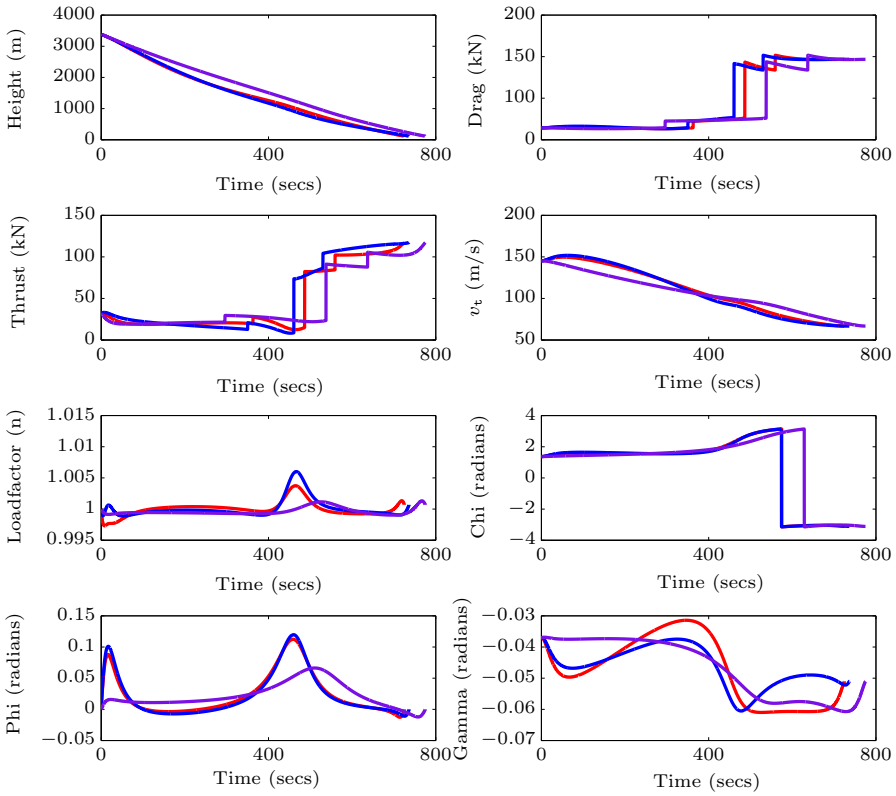

FIGURE 8.7: NEMA case: Minima trajectories for B752 Pareto front

Figures 8.8 to 8.13 show the clustered path, height, speed, thrust, inertial flight angle and energy profiles for the Pareto front solutions. The energy E (MJ) was calculated by summing the aircraft kinetic and potential energies

$$E = mgh + \frac{1}{2}mv_{\mathrm{t}}^2 \tag{8.2}$$

where m is mass, h is height and v_{t} is true airspeed. Examining the Pareto front trajectories relative to their clustering, the following observations were made.

Dark blue cluster
The trajectories in the dark blue cluster have both the highest Annoyance Score and the lowest peak noise under the centreline values. The trajectories in this cluster have the highest Annoyance Score due principally to the close proximity of the trajectory paths to the city of Leicester. The lowest peak noise under the centreline values are achieved by the same trajectory paths as they maximise the distance from the aircraft to the noise monitors.

For the vertical profiles, the trajectories in this cluster have initial speeds that are on average higher than the descent speeds in all the other clusters. Therefore the trajectories initially commence a quicker descent on the base leg. Just prior to the turn onto finals, aircraft increase thrust and flight path angle, reducing the average descent speed of trajectories in this cluster to among the lowest of all the clusters. As the turn radius is proportional to aircraft true airspeed, this helps achieve a tighter turn onto the final leg. Emerging from the turn, the aircraft in this cluster are then lower and slower relative to the trajectories in the other clusters. Therefore trajectories in this cluster are required to introduce extra thrust to increase the flight path angle and merge the onto the constrained descent slope. The use of extra thrust on final approach has the effect of increasing the fuel burn for these trajectories relative to trajectories in other clusters with comparable flight times.

Light blue cluster
The paths for the trajectories in this cluster move slightly away from Leicester at the expense of moving closer to the noise monitors. Therefore the solutions for the trajectories in this cluster have Annoyance Scores that reduce slightly relative to the solutions in the dark blue cluster, while the peak noise under the centreline values begin to increase slightly relative to the dark blue cluster.

The trajectories in this cluster generally have shorter flight times relative to the other clusters, which acts to minimise fuel burn. Many of the trajectories in this cluster have fuel burn values that are indistinguishable from the trajectories in the red cluster, which contains solutions that are clustered to the lowest fuel burn solution

For the vertical profile, the trajectories fly profiles highly similar to those in the dark blue cluster. However, when turning on to final approach, the larger turn radius allows

trajectories to be on average slightly higher and faster requiring less thrust to maintain the constrained flight path angle on final approach.

Red cluster
The trajectories in the red cluster are clustered to the lowest fuel burn solution. The trajectories in this cluster have relatively short flight times and flight paths. The vertical profiles for the trajectories in these clusters are very similar to the vertical profiles in the blue colored clusters in that they have high initial descent speeds on the base leg, followed by an increase in thrust and flight path angle prior to the turn on to finals that reduces the descent speed and descent rate. The larger turn radius allows aircraft in the red cluster to emerge from the turn on average higher and faster than those in the blue clusters. This allows the trajectories in this cluster to on average maintain steeper flight path angles on final approach that acts to lower the average thrust settings relative to the blue clusters. The combination of short flight times and lower average thrust setting on final approach acts to minimise the fuel burn from the trajectories.

Although the solutions in this cluster are clustered to the minimum fuel burn solution, there are a lot of solutions in the light blue cluster and the orange cluster with equivalently low fuel burns. In the light blue cluster however, the trajectories have higher Annoyance Scores due to their proximity to Leicester. In the orange cluster the low fuel burn values come at the expense of higher peak noise under the centreline values.

Orange cluster
The trajectories in the orange cluster have flight paths that, relative to the trajectories in the other clusters, pass closest to the noise monitors, which results in the trajectories in this cluster having the highest peak noise under the centreline values. Apart from the flight paths, the trajectories in this cluster had height, speed and thrust profiles that were largely similar to those in the red cluster.

Green cluster
The solutions in the green and violet clusters have low Annoyance Scores at lower peak noise under the centreline values than the solutions in the orange cluster. The trajectory paths for solutions in the green and violet clusters have long arcing paths away from Leicester that help minimise overall community noise impact without shifting the impact to another community.

The trajectory paths fly to the right of the consultation zone, pushing against the constraints to maximise the distance from the paths to the noise monitors. Of the two clusters, the solutions in the green cluster have on average lower fuel burn than those in the violet. The most significant difference between between the trajectories in the two clusters is that the green clustered trajectories are considerably faster on the base leg. This minimises the flight time for these trajectories helping to minimise total fuel burnt.

Violet cluster
The trajectories in the violet cluster descend slower on the base leg than those in the

green cluster. Prior to the turn onto finals, the thrust is increased and is used to maintain height while increasing the speeds to match those of the trajectories in the green cluster. The slower baseleg and the higher thrust required to match the speed profiles on the final leg both contribute to the cluster having higher fuel burn values for the same flight path as the trajectories in the green cluster. The extra thrust required, did not contribute to a higher Annoyance Score for the cluster and if anything the higher height on the base leg of the approach served to minimise the noise experienced by communities on the ground from trajectories in this cluster.

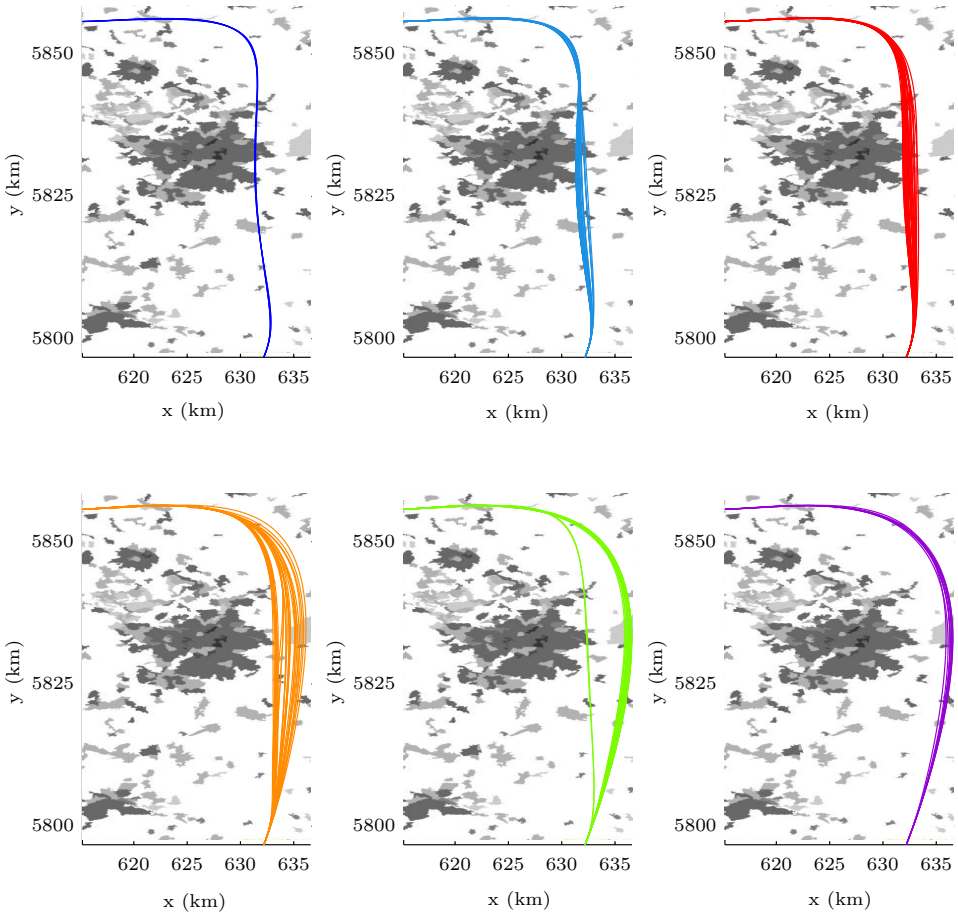

FIGURE 8.8: NEMA case: 752 optimisation clustered trajectory paths

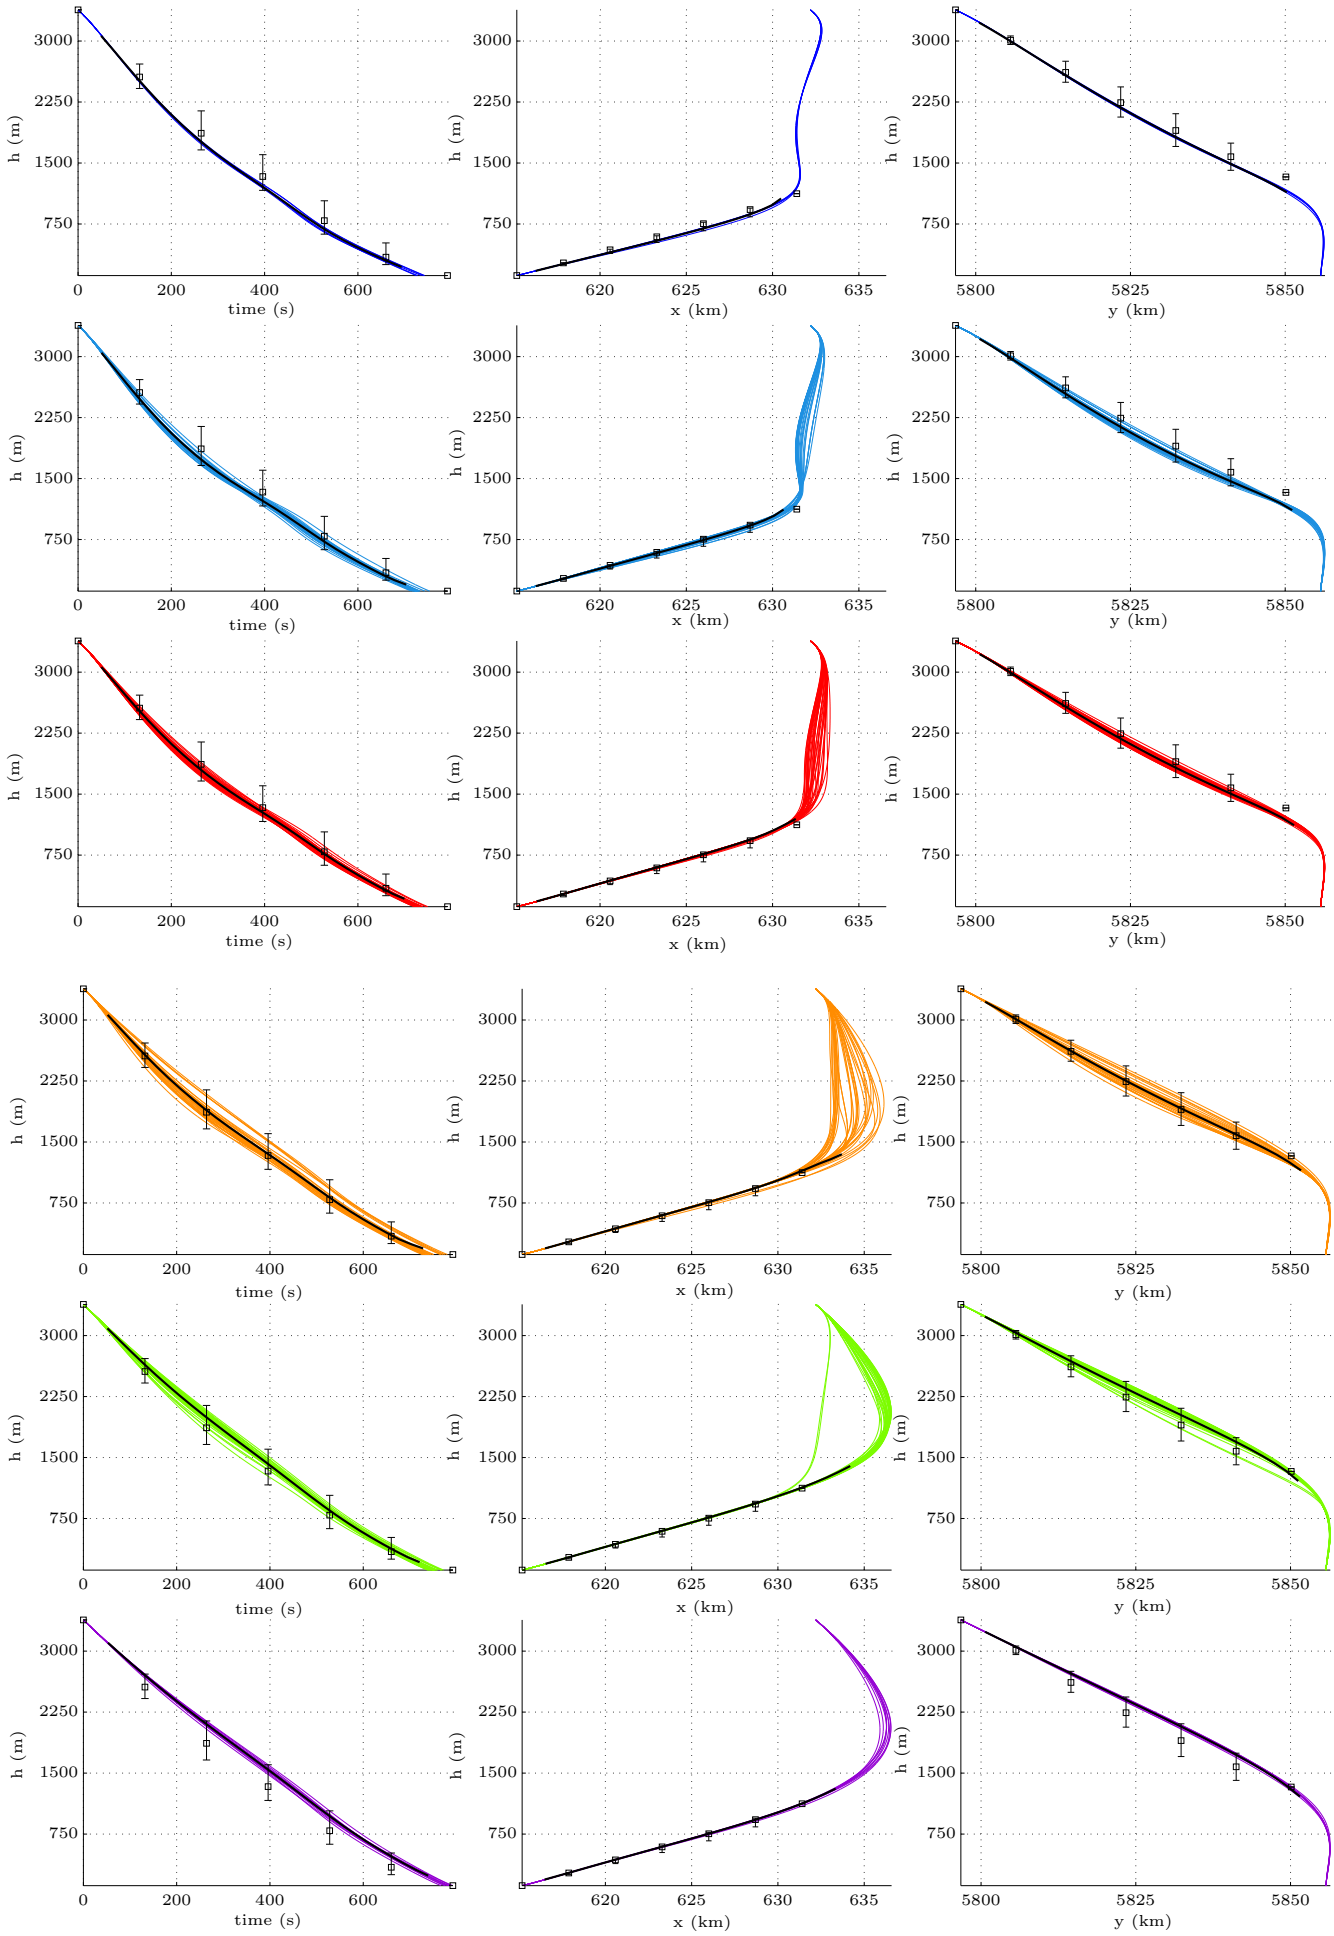

FIGURE 8.9: NEMA case: 752 optimisation clustered height profiles

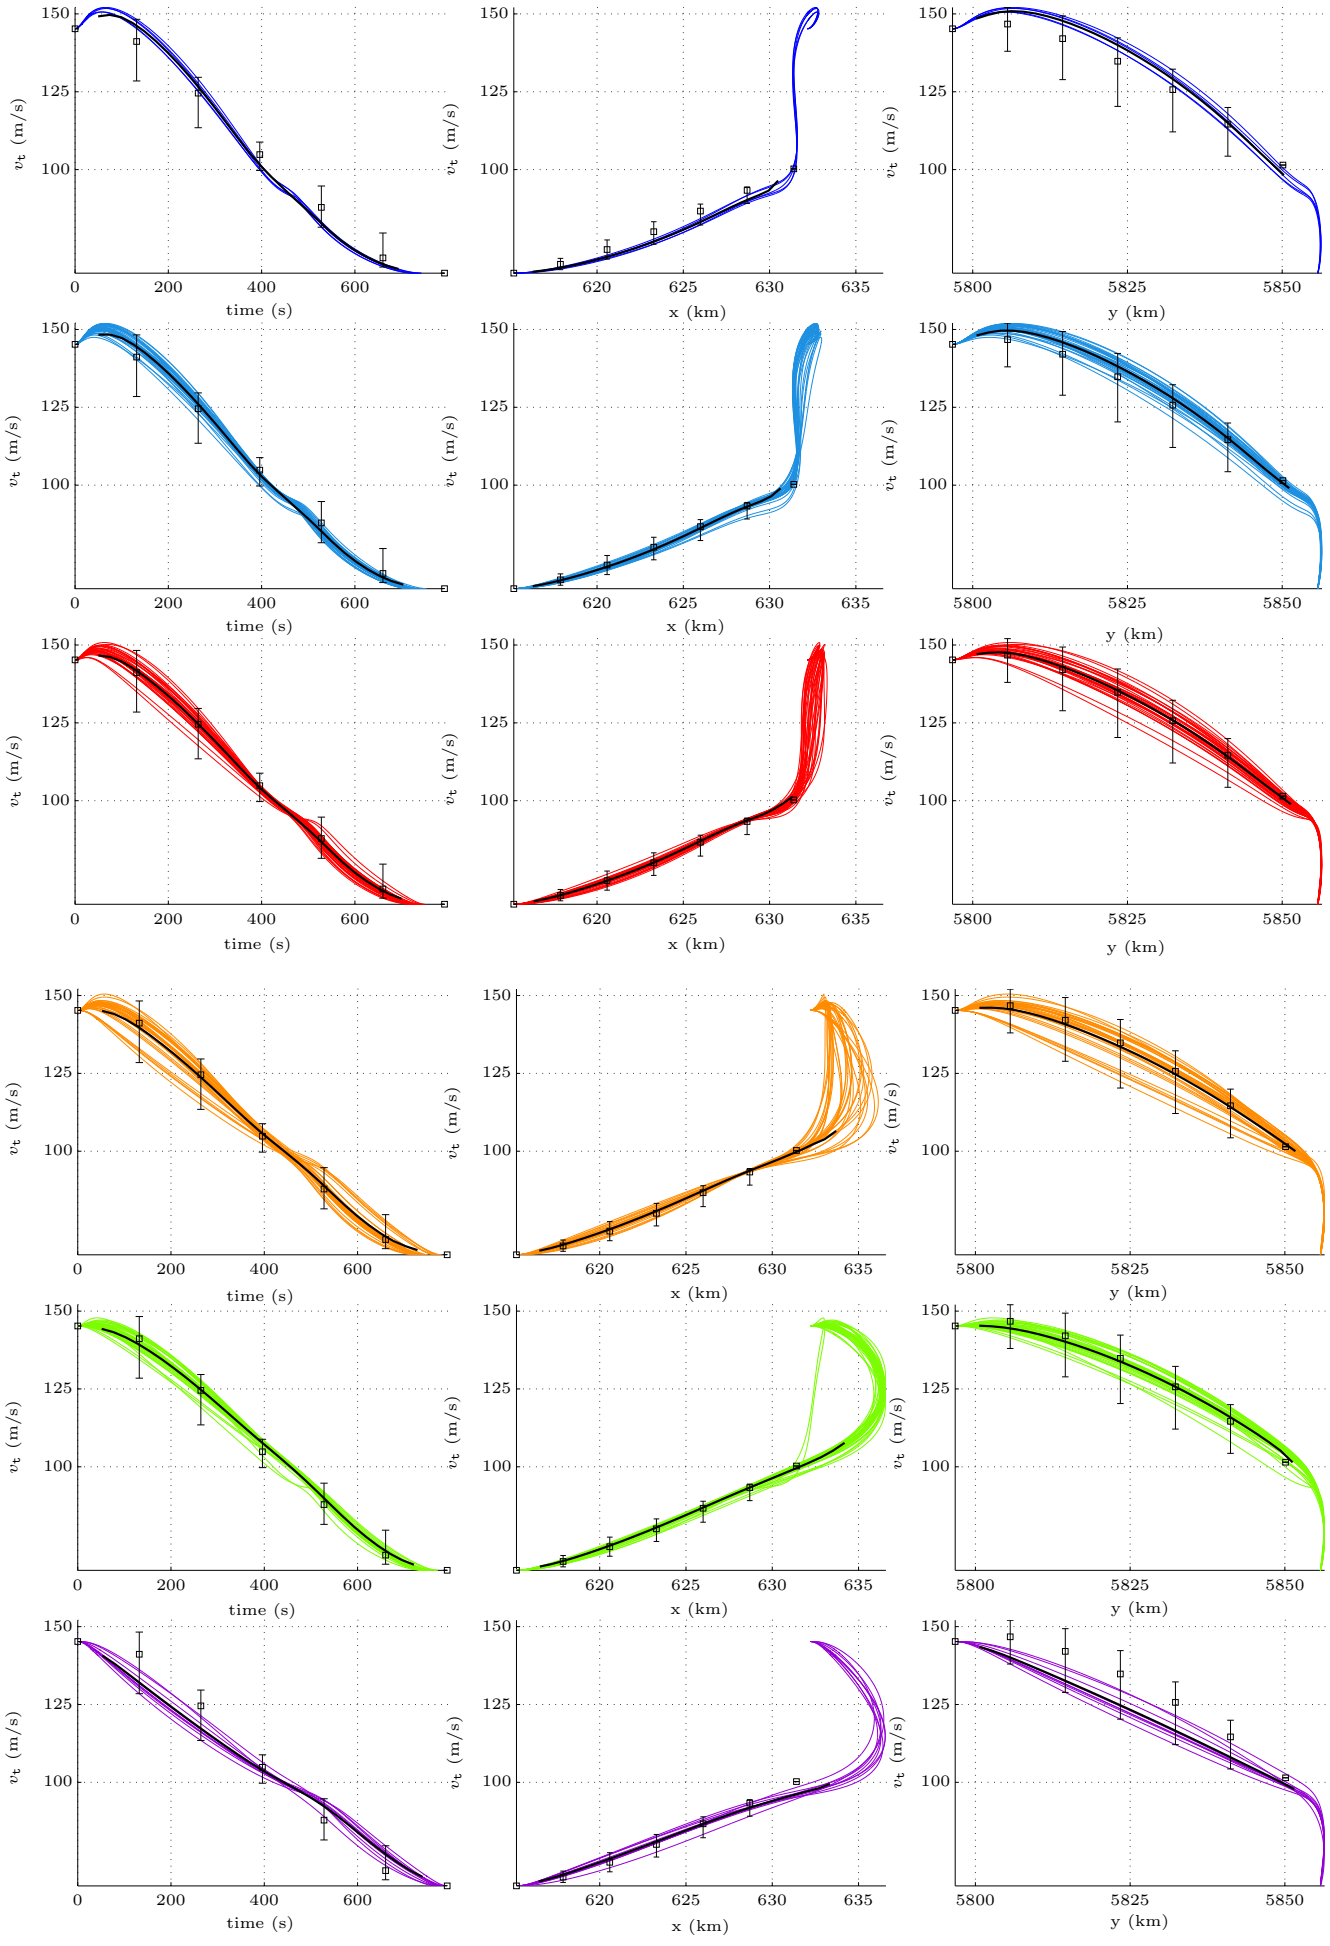

FIGURE 8.10: NEMA case: 752 optimisation clustered speed profiles

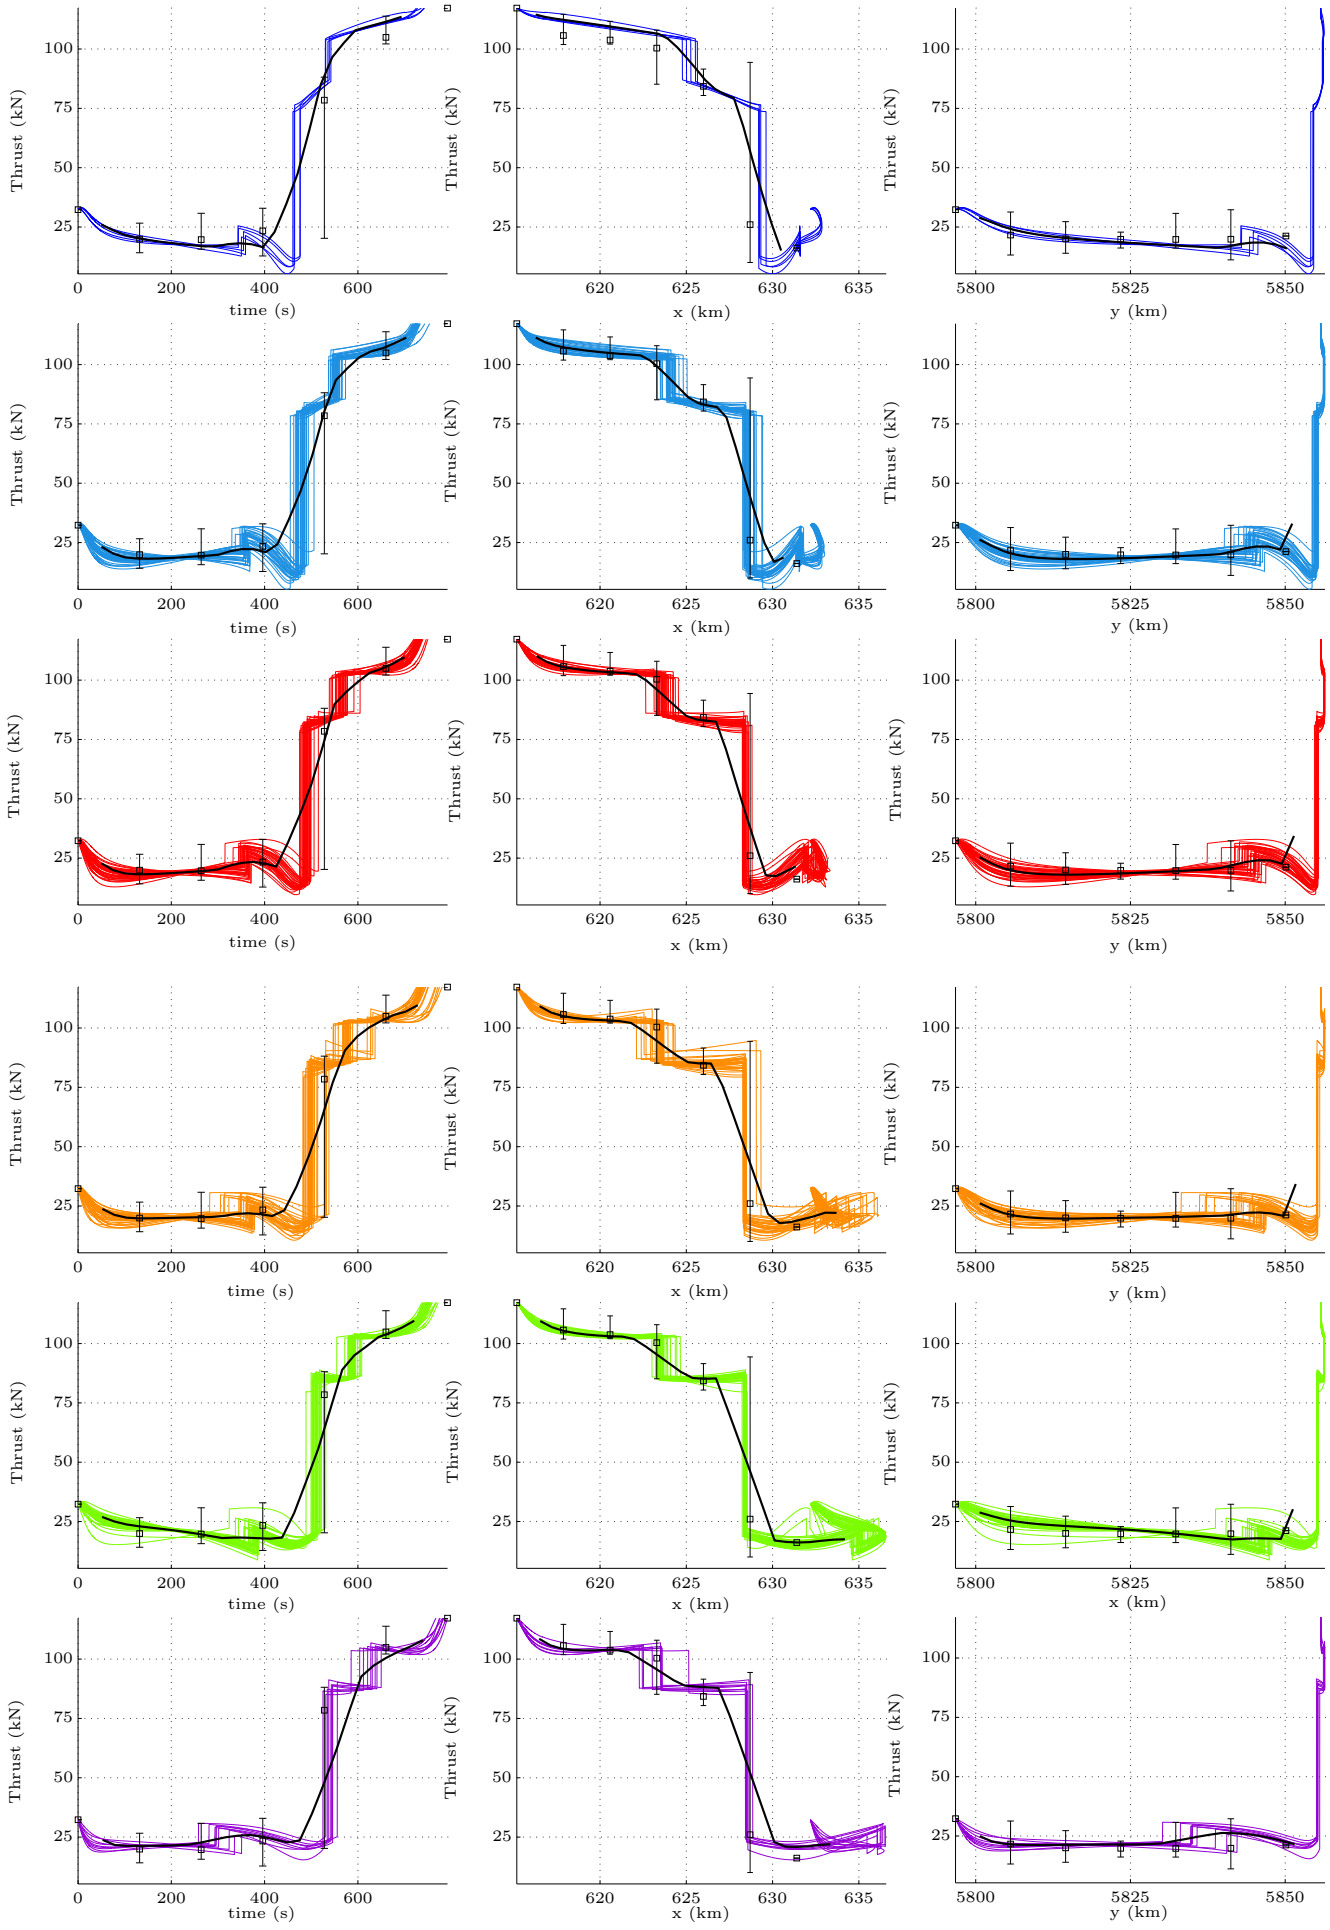

FIGURE 8.11: NEMA case: 752 optimisation clustered thrust profiles

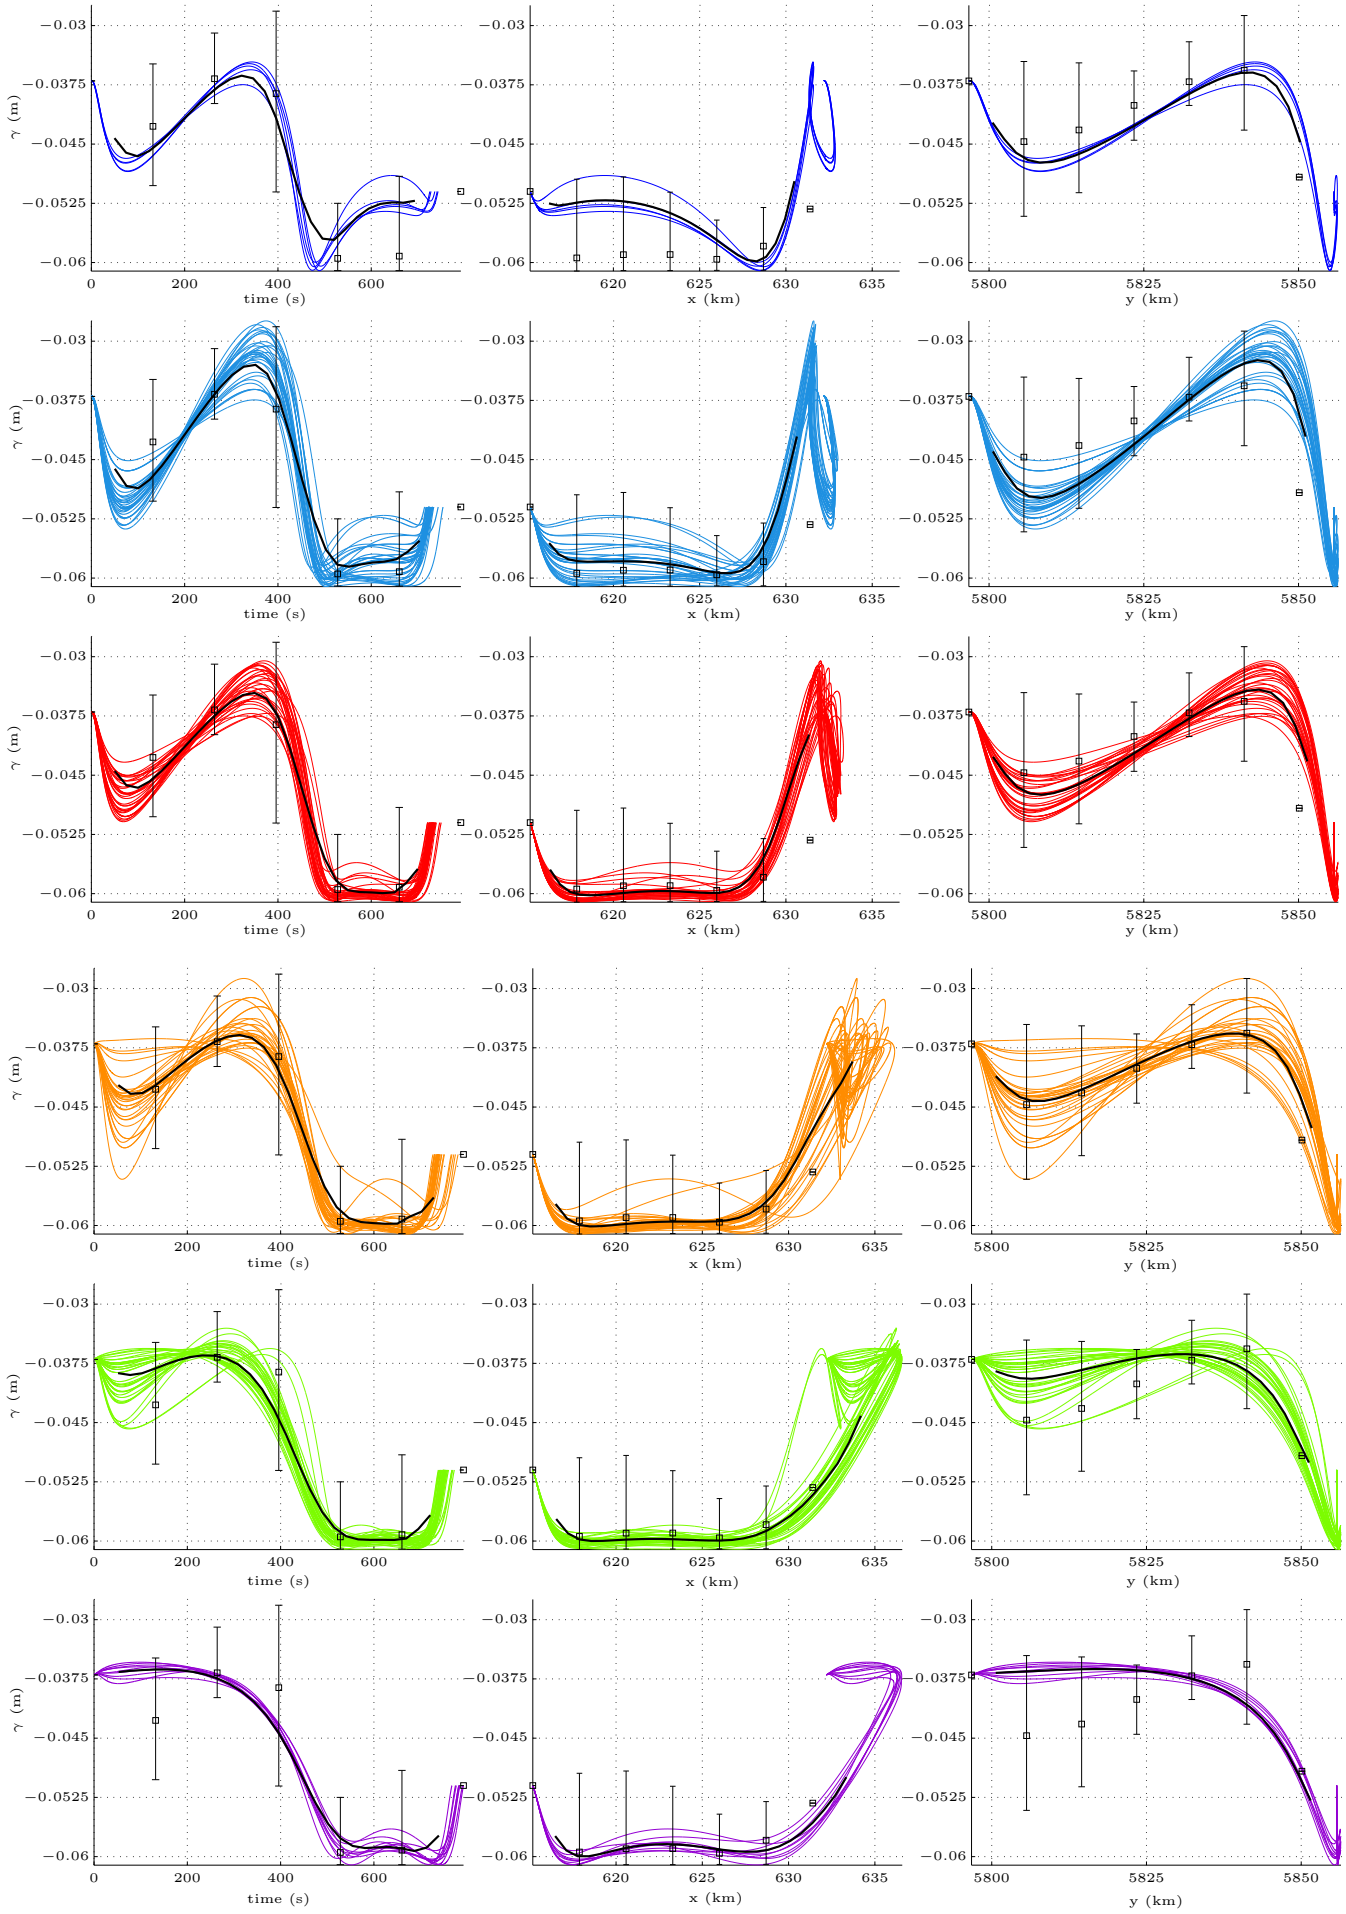

FIGURE 8.12: NEMA case: 752 optimisation clustered γ profiles

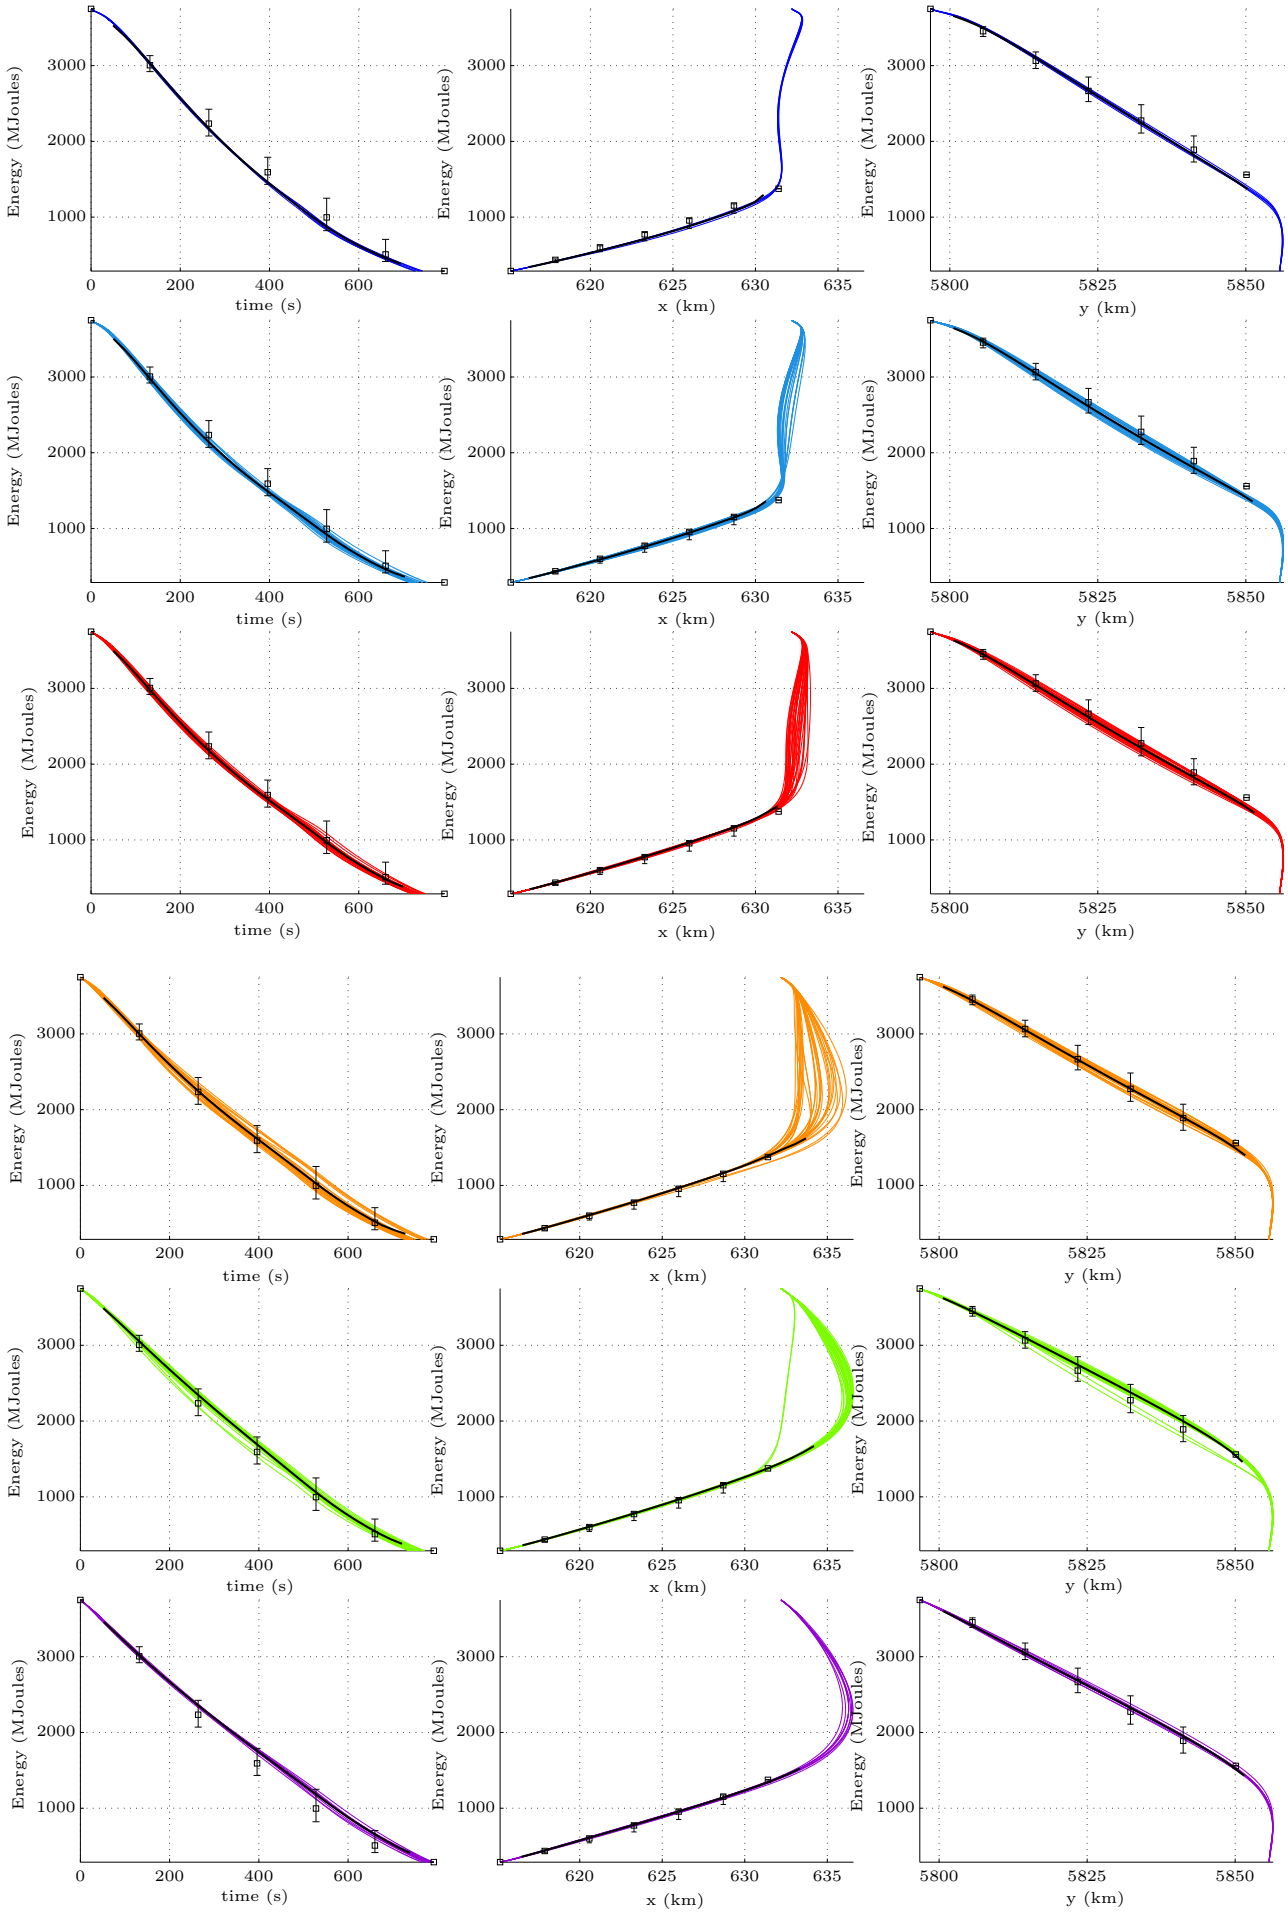

FIGURE 8.13: NEMA case: 752 optimisation clustered energy profiles

8.2.3 Comparison with Baselines

Figure 8.14 shows the height, speed and γ profiles of a Continuous Descent Approach for a Boeing B752 aircraft flying the NEMAX1A procedure. The trajectory was generated by the TSAT model used in the design of the NEMA approach procedures and published as part of the NEMA study report [165]. The descent trajectory was used to define the boundary values and constraints of the optimisation study. The trajectory was additionally used as a baseline for comparison to the trajectories calculated by the IDVD-DE multiobjective optimisation.

The height and speed profiles of the B752 aircraft operating the NEMAX1A descent procedure were digitised from the NEMA trial report and are shown in blue in Figure 8.14. Polynomials were then fitted to the data so that 3D flight paths could be constructed and used to approximate flight path angle and thrust histories. The height, speed and gamma profiles show a CDA descent with alternating constant CAS descent and deceleration segments.

The baseline trajectory starts at 13000 ft (3962 m) and 250 kts (128 m/s) at a distance of 92 km from touchdown. The first step in the descent is a constant CAS descent to 11000 ft (3352 m). The 752 then flies a shallow flight path angle in order to reduce speed to 230 kts (118 m/s). It then commences a long descent on the straight base leg to 5000 ft (1524 m) and 220 kts (113 m/s). Immediately prior to the turn on to final approach, the aircraft reduces speed to 180kts (93 m/s). As is standard, the aircraft crosses the final approach fix at 4DME at 160 kts (82 m/s) and establishes on the ILS, reducing speed to 140 kts (72 m/s) for touchdown. The deceleration steps to 230, 180, 160 and 140 knots are clearly shown by the spikes in the flight path angle, where γ is increased so energy reduction is achieved preferentially through reductions in speed.

It can seen from Figure 8.15 to Figure 8.18 that the trajectories generated by the IDVD-DE method compare favorably to the baseline trajectory. There are differences between the trajectory results and on average it can be seen that the optimisation trajectories have faster descent speeds on both the base and the final legs. This leads to all of the optimization trajectories having shorter flight durations than the baseline trajectory, even where flight path distances are equivalent or greater than the baseline. The IDVD-DE trajectories flew fast to the upper bounds of the 220 kt (113 m/s) speed constraint and reduce speed in order to reduce the turn radius and to meet the 180 kt (93 m/s) constraint. Relative to the pathlength of the trajectories, IDVD-DE trajectories introduce thrust earlier than the baseline. This allows the IDVD-DE trajectories to maintain greater height and speed relative to the baseline just prior to and during the turn to final leg. The IDVD-DE trajectories then turn on to the final leg higher and faster than the baselines. This allows the optimisation trajectories to fly a prolonged steeper flight path angle on final approach so that aircraft energy can be reduced with less need for extra thrust to be introduced. Although the IDVD-DE generated trajectories had, in general, faster speed profiles than the baseline trajectory, the IDVD-DE speed profiles were shown to

be realistic when compared to the FDR speed profiles recorded for from the operational flight trial [165].

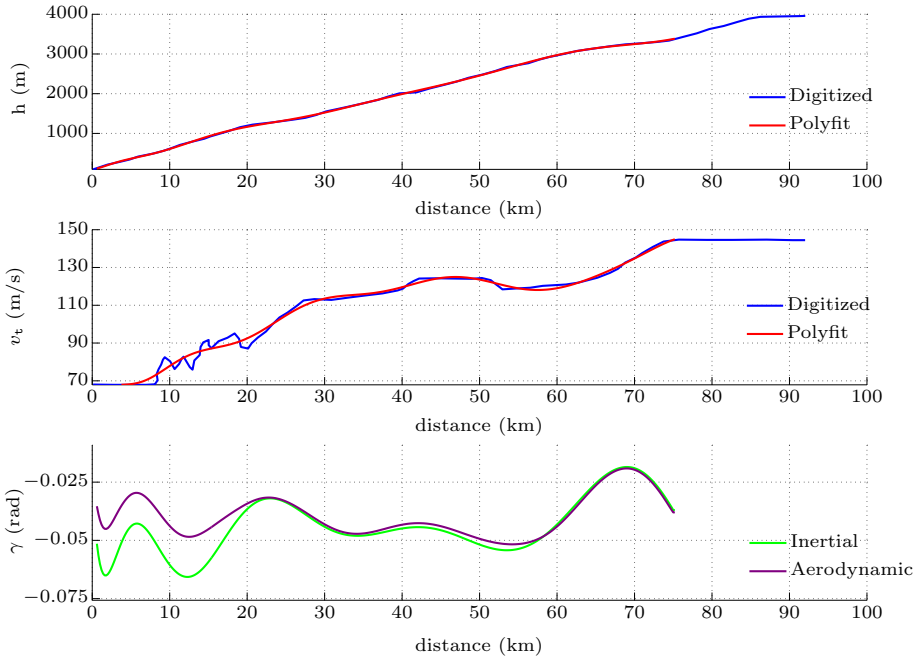

FIGURE 8.14: NEMA case: Reference 757 data for NEMA scenario

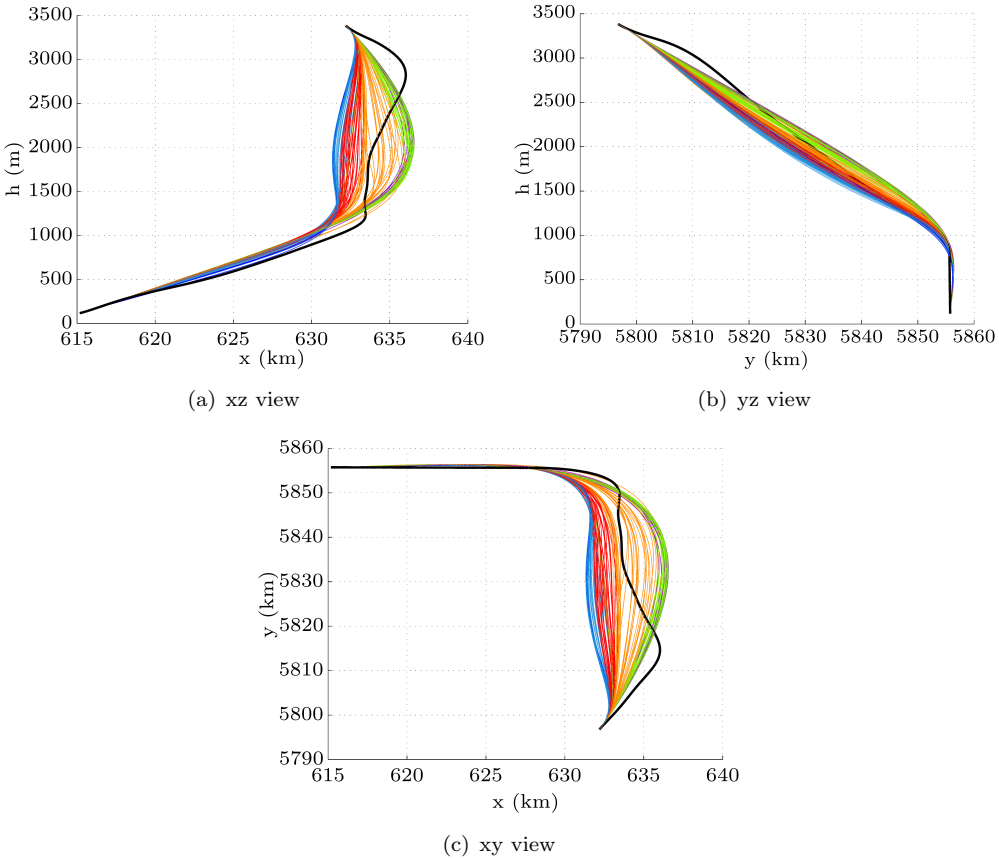

(a) xz view

(b) yz view

(c) xy view

FIGURE 8.15: NEMA case: 752 Pareto front height and path profiles with baseline

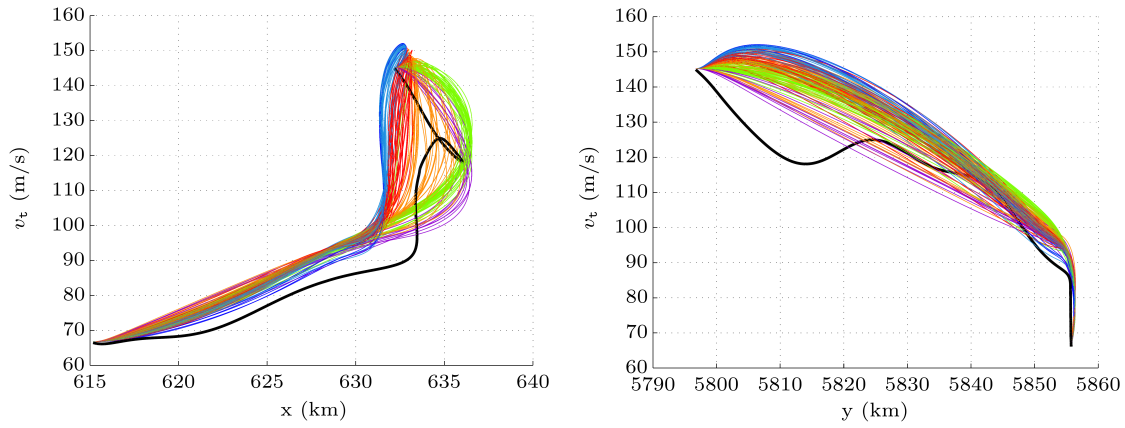

(a) xz view (b) yz view

FIGURE 8.16: NEMA case: 752 Pareto front speed profiles with baseline

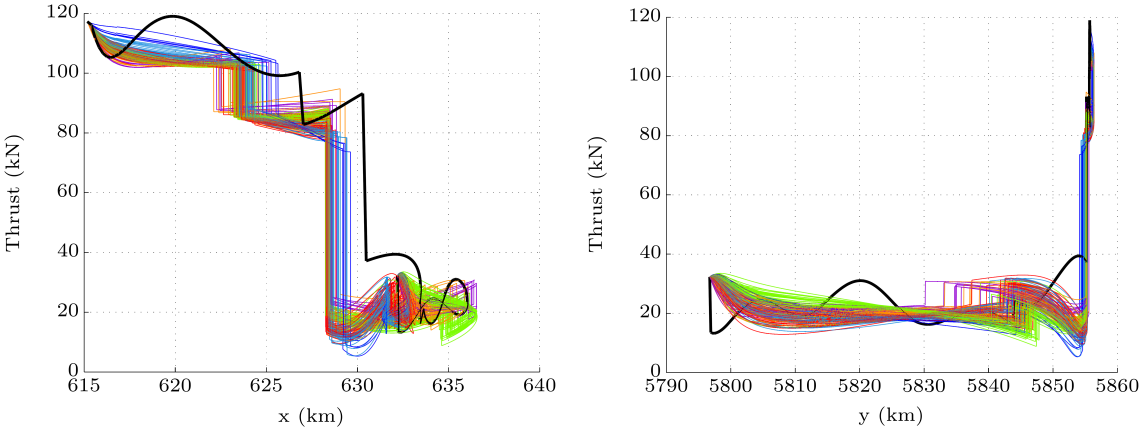

(a) xz view (b) yz view

FIGURE 8.17: NEMA case: 752 Pareto front thrust profiles with baseline

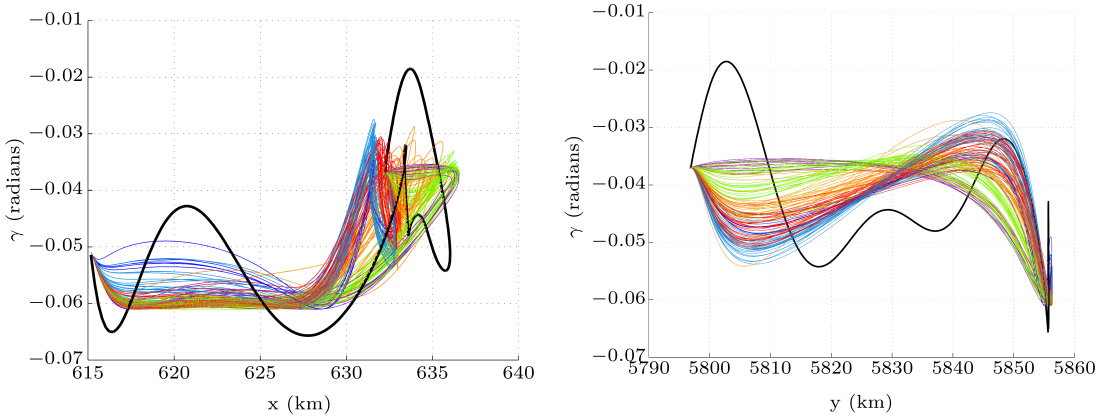

(a) xz view (b) yz view

FIGURE 8.18: NEMA case: 752 Pareto front γ profiles with baseline

8.2.4 B744 Optimisation Scenario Results

The second most route dominant aircraft operating in the trial was the MD11 aircraft. Any procedure design would ideally accommodate the efficient operation of both the dominant aircraft types. Therefore a second Pareto front was generated for the MD11 aircraft (modelled here, as in the SAI study, as a B744) so that an approach procedure could be developed with consideration of the operating characteristics of both aircraft types.

The Pareto front in Figure 8.19 is clustered into five individual clusters. As with the B752 Pareto front in Figure 8.5, there are three minima clusters. Unlike the B752 Pareto front, there are two instead of three transition clusters, as the extra cluster was not required to differentiate the trajectory behaviour.

It can be seen from the Pareto fronts for both the B752 and the B744 aircraft, that the relationships between the trajectories and the objectives are similar in the blue clusters for both aircraft. Therefore, as with the B752, the dark blue cluster for the B744 simulation contained the solutions with the highest Annoyance Scores and lower peak noise values. This was due to the solutions in that cluster having trajectory paths close to Leicester but away from the noise monitors

The red cluster in Figure 8.5 contains the solutions that are clustered to the lowest fuel burn cluster. The trajectories in this cluster for the B744 aircraft have longer arcing paths than the trajectories in the equivalent cluster of the B752 optimisation. The shallower turn created by the longer path allows the larger aircraft to fly faster into and out of the turn from the base leg onto the final leg. This minimises flight time and allows the aircraft energy to be held high at the start of final approach so that the reduction in height and speed can be managed in a manner that minimises the introduction of thrust on final approach. For the B744 optimisation, the trajectories in the red cluster fly the most directly over the noise monitors and therefore have the highest peak noise values.

Examining the B744 Pareto front and related trajectories, it can be seen the solutions in the orange and green clusters have virtually identical trajectory paths. However, the trajectories in the green cluster have on average higher fuel burn values than those in the orange cluster. The trajectories in the green cluster fly a slower base leg than those in the orange cluster, lengthening the flight time and increasing the relative fuel burn. The trajectories in the green cluster are however higher on the base leg, which helped contribute to lower Annoyance Scores for some solutions in that cluster. Trajectories in the green cluster also introduce thrust earlier than in any of the other clusters. This is used prior to the turn onto finals to keep both height and speed high such to minimise the use of thrust on final approach.

It can be seen from the Pareto front in Figure 8.19 that the Annoyance Scores for the red, orange and green clusters are similarly distributed between the values 3.4 to 4×10^{13}. The B744 is a large 4 engine aircraft, making it a particularly noisy aircraft, and it can be seen from the plot that, for this aircraft, once the flight path was moved away from Leicester, that this noisiness makes the noise Annoyance Score relatively insensitive to further changes in the paths and vertical profiles. The other metrics do however remain sensitive to the vertical and horizontal profiles so that similar Annoyance Scores values were achieved at different peak noise and fuel burn values.

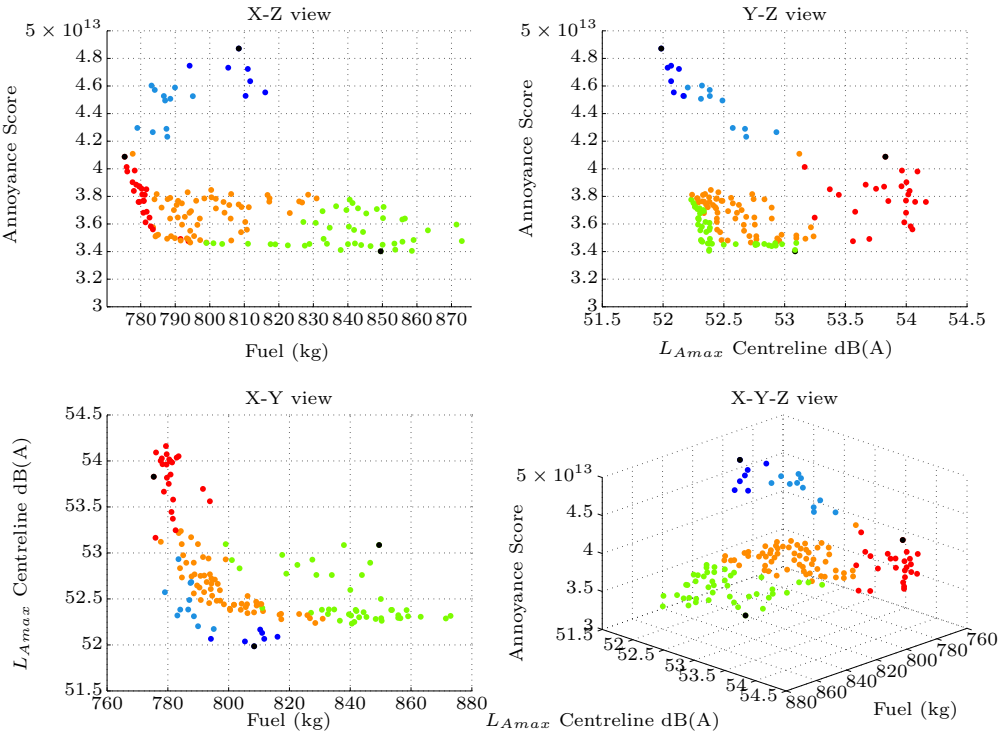

FIGURE 8.19: NEMA case: Four view B744 simulation Pareto front

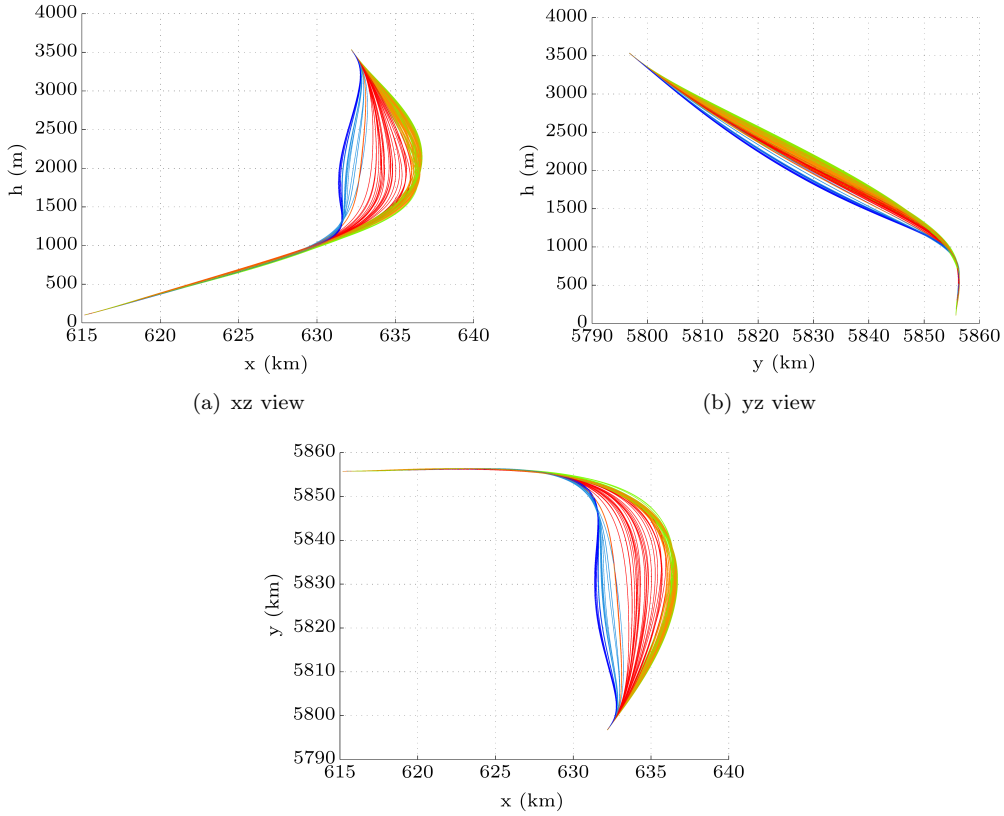

(a) xz view

(b) yz view

(c) xy view

FIGURE 8.20: NEMA case: B744 Pareto front height and path profiles

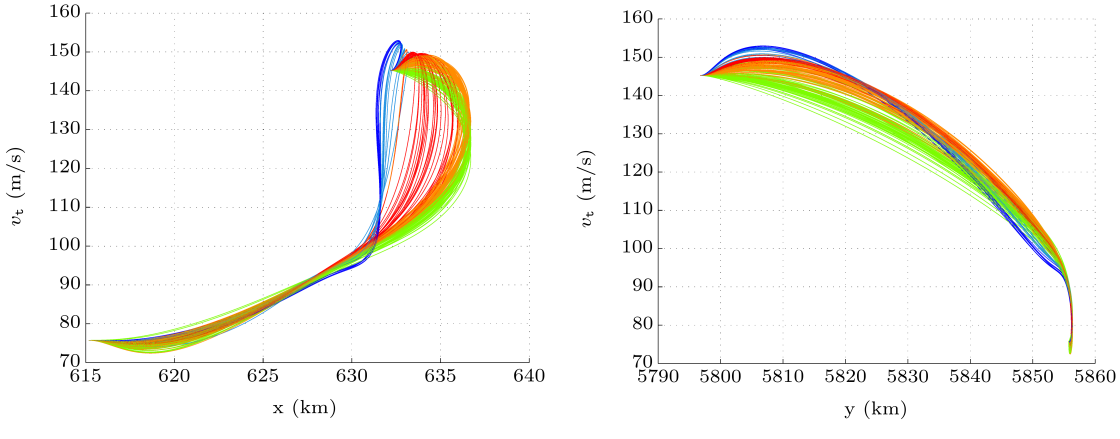

(a) xz view

(b) yz view

FIGURE 8.21: NEMA case: B744 Pareto front speed profiles

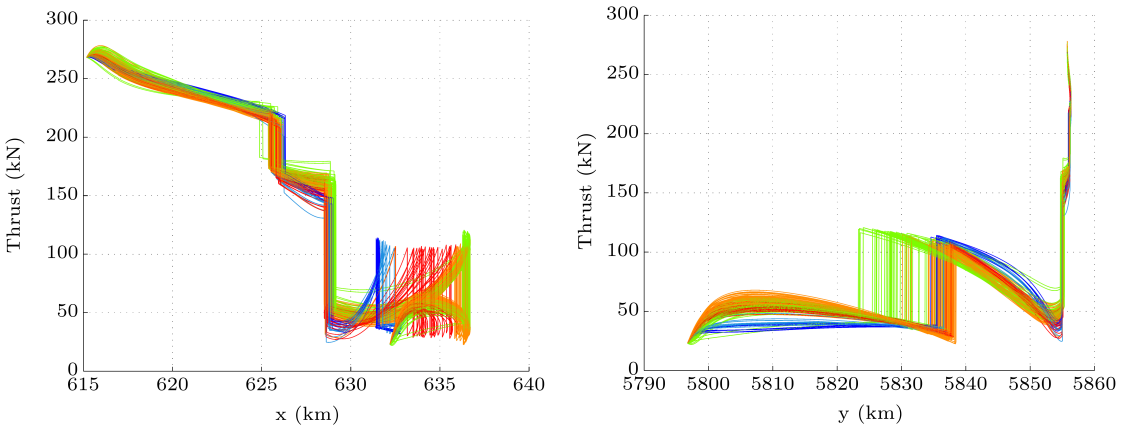

(a) xz view (b) yz view

FIGURE 8.22: NEMA case: B744 Pareto front thrust profiles

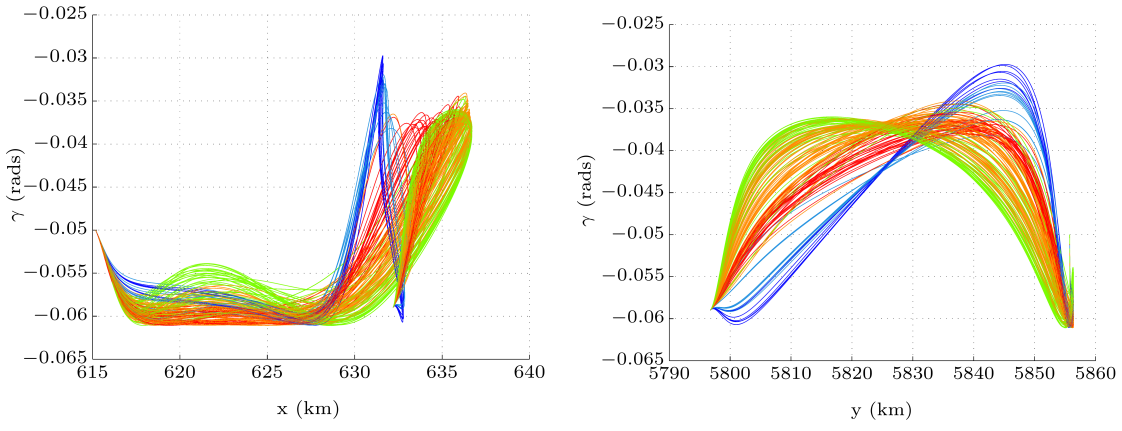

(a) xz view (b) yz view

FIGURE 8.23: NEMA case: B744 Pareto front γ profiles

8.2.5 Proposed Procedure Alteration

By analysing the results of the B752 and B744 optimisation scenarios in Sections 8.2.2 and 8.2.4 a number of observations could be made.

First, moving flight paths away from Leicester had the effect of minimising overall community noise impact. Also, allowing aircraft to turn onto the final approach leg higher and faster, i.e. at a higher energy state, reduced the need to introduce energy in the form of thrust on final approach to stabilise the aircraft, establish and then maintain a descending 3 degree slope.

Therefore it was decided to define a new approach procedure that, relative to the NEMAX1A baseline, moved aircraft paths away from Leicester. The new procedure was also chosen to facilitate a higher and faster turn onto final approach. This was achieved by defining a procedure route, that, relative to the NEMAX1A baseline, had a larger turn radius between the base and final legs. The shallower turn onto finals also meant that the procedure could in general be flown faster, as from the results, it could be seen that minimising flight time was a factor in minimising trajectory fuel burn.

The new proposed procedure route is shown in blue in Figure 8.24. The alteration to the baseline NEMAX1A procedure route was developed from solutions in the green cluster for the B752 optimisation scenario and the yellow cluster from the B744 optimisation results. The trajectories of these solutions corresponded well with the desired changes in the procedure and accordingly the solutions offered lower levels of community noise annoyance and peak noise under the flight path, while not incurring excessive fuel burn penalties.

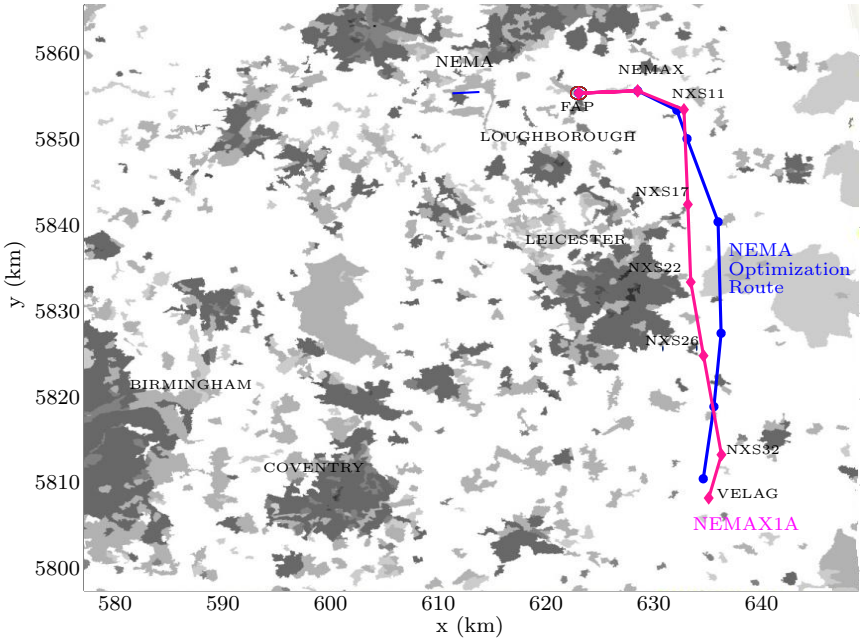

FIGURE 8.24: NEMA procedure alteration

The changes to the procedure were designed to enable particular types of vertical profiles. However to determine if this had been successfully achieved and to determine what the remaining environmental trade-offs were for aircraft operating the route, further simulations for the B757 and B744 aircraft were conducted in Sections 8.2.6 and 8.2.7.

In these optimisation scenarios, the aircraft were constrained to fly within 1 nautical mile (RNP1) of the new proposed route centreline. The sectorization height and consultation zone constraints remained in place from Section 8.2.1. The speed constraints in Section 8.2.1 were also retained as they they reflected standard operating constraints but still allowed considerable flexibility of the speed profile along the trajectory path.

The peak noise under the centreline metric was not considered a useful metric in this study and was at this point dropped from the analysis. The peak noise under the centreline metric was used in the baseline SAI study and was therefore carried over for use in the optimisation study. However the metric did not provide a good indicator for any sort of community noise impact. Because of the averaging involved in the calculation of the measure, there were a lot of very different trajectories with equivalent noise under the centreline performance. This made if difficult for the algorithm to converge on solutions with tangible differences in noise under the centreline values. This seemed to contribute a noisiness to the distribution of solutions on the Pareto fronts, making the trade-offs between the measures more difficult to analyse. The metric simply had the effect of pushing the trajectories away from the noise monitoring points. Also, the choice of the altered procedure path acted to minimise the peak noise measure and therefore it was not considered necessary to further optimize for this metric. It is proposed that, in general, the peak noise measure only be used in further optimisation studies where the change in noise values can be tied to noise impact at specific places or communities.

8.2.6 B752 Constrained Path Simulation

Figure 8.25 shows the clustered Pareto front between fuel burn and Annoyance Score for B752 aircraft constrained to be within 1 nautical mile of the new procedure path shown in Figure 8.24. The related Parallel Coordinate plot is shown in Figure 8.26. The Parallel Coordinate plot shows the relationships between the optimisation variables and the objective measures.

It can be seen from the Parallel Coordinate plot in Figure 8.26 that the path optimisation variables are virtually identical for all clusters. Changes in the trajectories were achieved almost exclusively through changes in the speed profile optimisation variables. Therefore, it can be seen in Figures 8.27-8.30 that the paths and height profiles are the same for all the Pareto front trajectories and that the differences in the objective values are driven by the differences in the trajectory speed profiles and the thrust required to achieve the differing speed schedules.

Red cluster
The red cluster contains the trajectories that have the lowest fuel burn values. These trajectories have the highest speeds on the base leg and then the lowest speeds on the final leg. The high speed on the base leg help minimise the flight time and therefore total fuel burnt. The red cluster trajectories initially have the highest levels of trust, which was required to maintain the same descent height and angle as the other clusters only with a higher descent speed. As the trajectories turn on to finals, the height is maintained at lower thrust levels by trading off speed for height. Therefore the trajectories in this cluster have lower speeds for most of the finals leg. Extra thrust is introduced late on final approach to maintain the flight path angle. This requires extra fuel burn at this stage of the descent, but by having a fast base leg descent and later trading some of that speed off for height, allows the trajectories in the red cluster to have less total fuel burnt than the trajectories in the other clusters.

Green cluster
The initial thrust for the trajectories in the green cluster is slightly less than those in the red cluster. Therefore the initial descent speed is lower than the red cluster. Approaching the turn onto the final leg the trajectories have less speed that can be traded off to retain height. Therefore extra thrust is introduced relative to the red cluster to maintain the descent height and angle. Similarly, on final approach more thrust is required than in the red cluster to maintain the descent slope at the higher speeds. Once joining the ILS however, the higher speed is then traded off for height and used to maintain the descent slope lowering the thrust used on the ILS. In the green cluster, the speed profiles for the trajectories on the base leg are slower than those in the red cluster. On finals the opposite case is true with the green clustered trajectories being faster than the red cluster. This speed schedule leads to less thrust being utilised in two distinct regions of the descent. The first region is as the aircraft approach Leicester and the second regions is as the aircraft is close to the ground descending on the ILS. This had the impact of

reducing noise on regions approaching Leicester and close to the airport, which then had the impact of reducing the noise Annoyance Score.

Blue cluster

The trajectories in the blue clusters were the slowest on the base leg and fastest on the final leg. This further continued the trend established by the trajectories in the green cluster where the thrust is minimised approaching Leicester and also when the aircraft is descending to the airport. This thrust schedule had the impact of reducing noise relative to the trajectories in the other clusters. The two low thrust regions at the start and end of the descent required that there be an intermediate region of high thrust introducing energy into the system where the aircraft approach and emerge from the turn on to finals. The two low thrust regions have the impact of reducing the noise Annoyance Score. However, the necessary intermediate high thrust region combined with the longer flight time from the slower base leg has the impact of increasing the fuel burn relative to the other clusters.

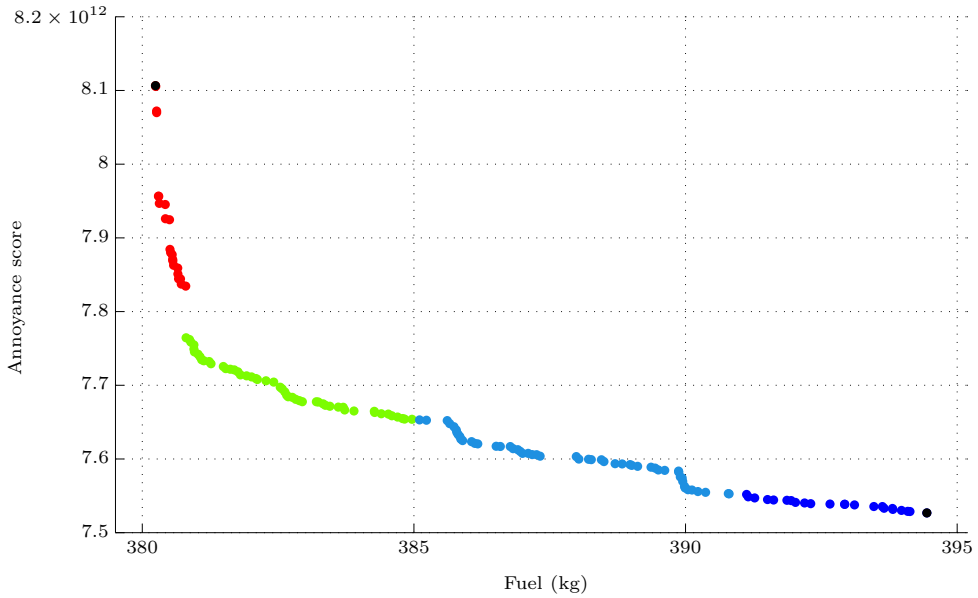

FIGURE 8.25: NEMA case: B752 constrained path Pareto front

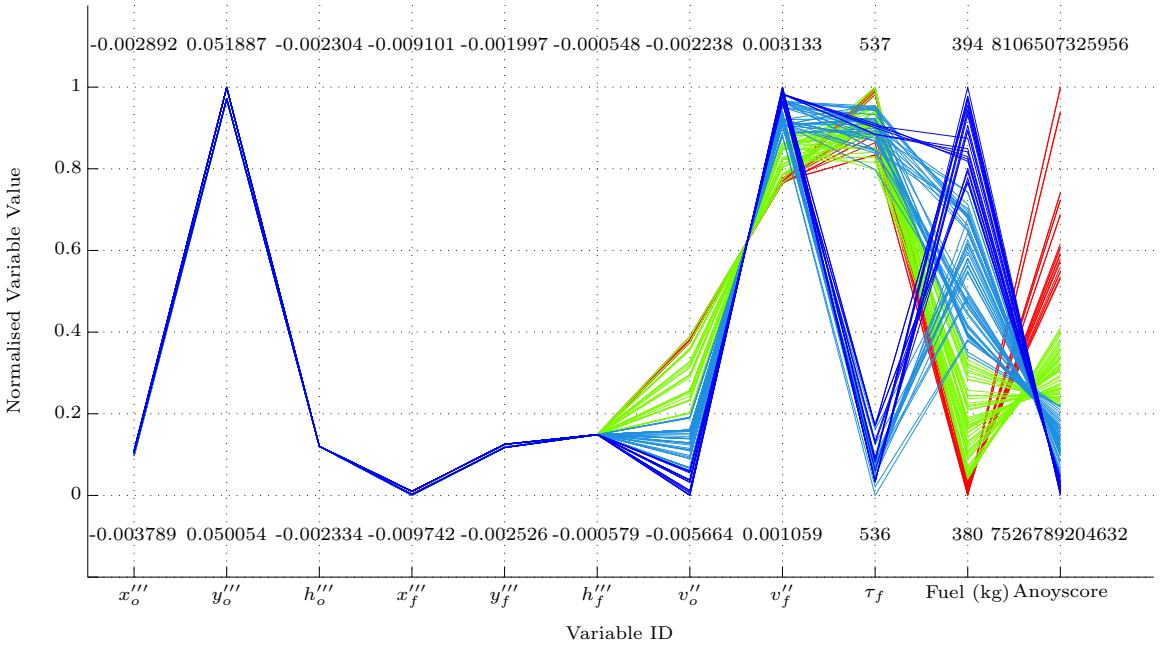

FIGURE 8.26: NEMA case: B752 constrained path Pareto front Parallel coordinate plot

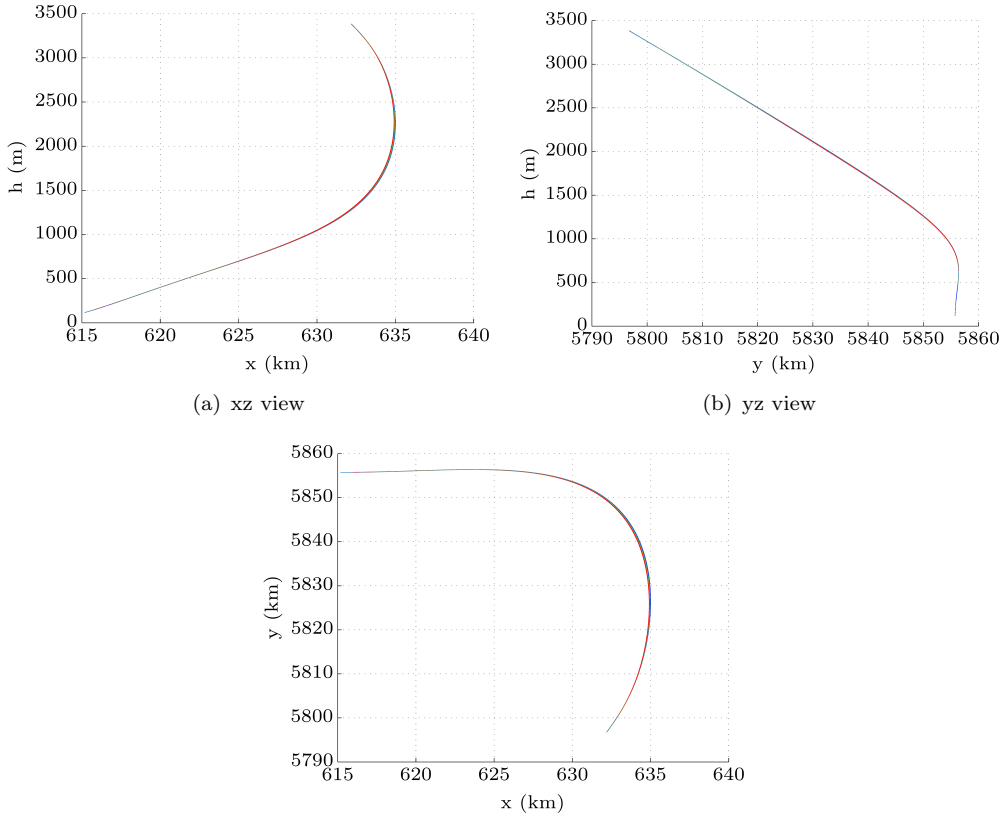

(a) xz view

(b) yz view

(c) xy view

FIGURE 8.27: NEMA case: B752 constrained path Pareto front height and path profiles

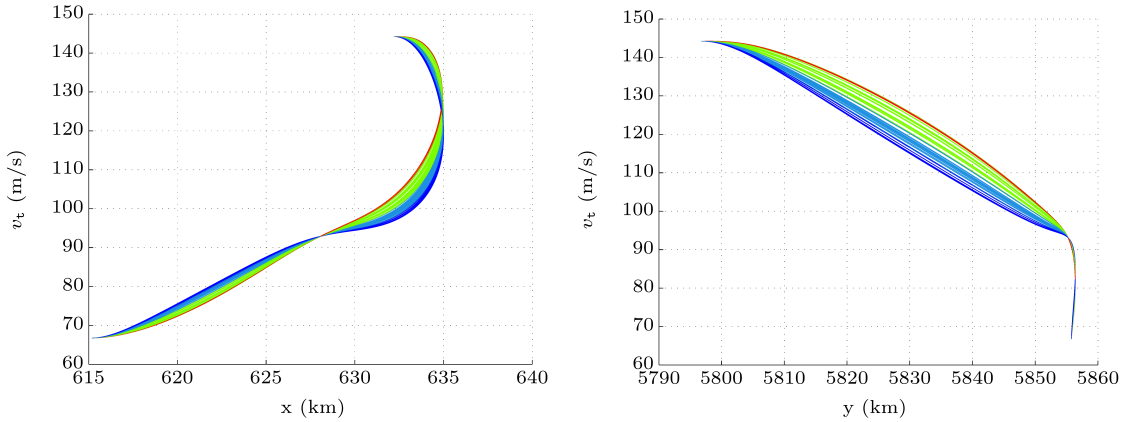

(a) xz view (b) yz view

FIGURE 8.28: NEMA case: B752 constrained path Pareto speed profiles

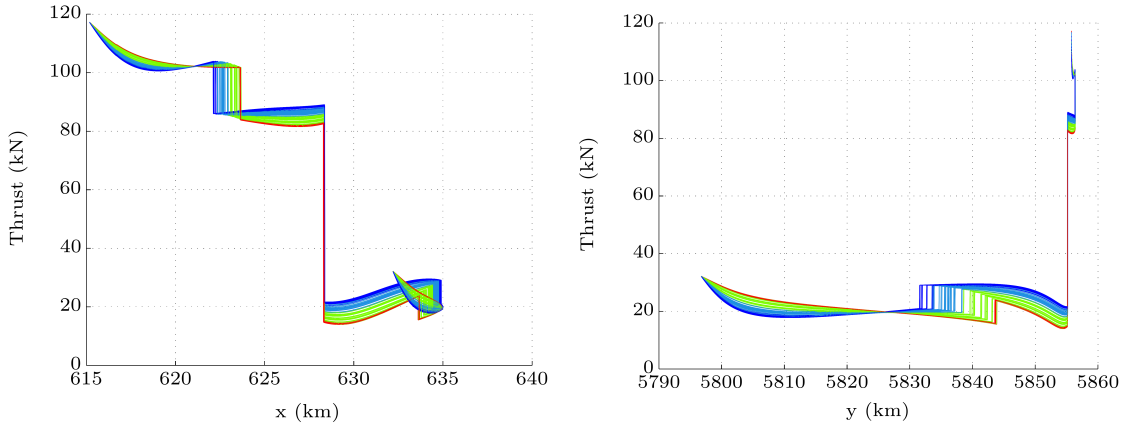

(a) xz view (b) yz view

FIGURE 8.29: NEMA case: B752 constrained path Pareto front thrust profiles

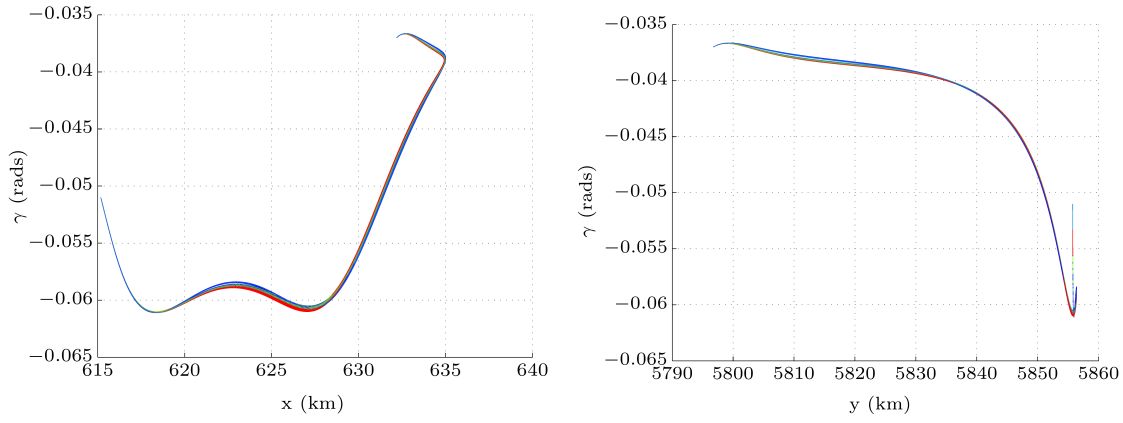

(a) xz view (b) yz view

FIGURE 8.30: NEMA case: B752 constrained path Pareto front γ profiles

8.2.7 B744 Constrained Path Simulation

Figure 8.31 shows the Pareto front between the fuel burn and Annoyance Score measures for the B744 aircraft on the constrained path. The front is split into two clusters. The red cluster contains the trajectories with the lower fuel burn values, while the blue cluster contains the trajectories with lower Annoyance Score values. The principle difference between the two is the trajectories in the red cluster fly a slightly shorter path to minimise flight time and fuel burn and the trajectories in the green cluster fly a slightly longer path that requires more fuel burn but maximises distance to populations therefore marginally reducing noise. As with the B752 Pareto front in Section 8.2.6, the trade-offs in environmental measures on the proposed procedure path were minimal. The values of the Annoyance Score and fuel burn measures on both Pareto fronts were also low relative to the range of values on the initial Pareto fronts in Section 8.2.2 and Section 8.2.4 respectively.

Figures 8.32 and 8.35 show the 3D height speed γ and thrust profiles. It can be seen from this that the height and speeds for the two clusters are almost indistinguishable, with the trajectories in the blue cluster using slightly more thrust and therefore fuel to complete the longer path.

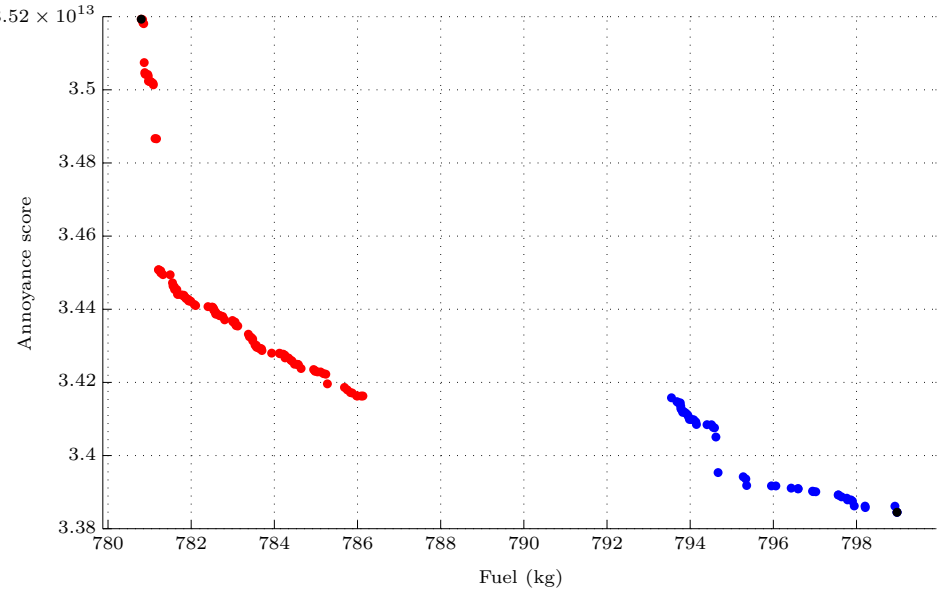

FIGURE 8.31: NEMA case: B744 constrained path simulation Pareto front

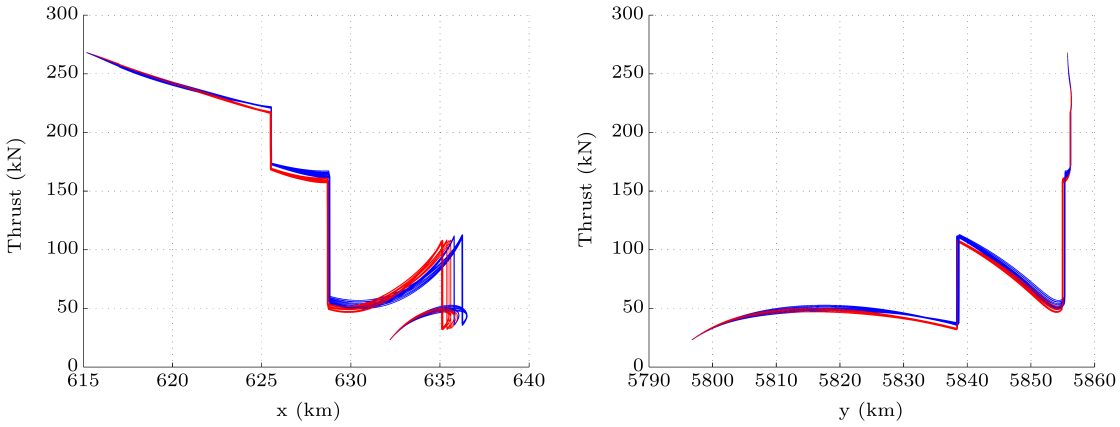

(a) xz view (b) yz view

FIGURE 8.32: NEMA case: B744 constrained path Pareto front thrust profiles

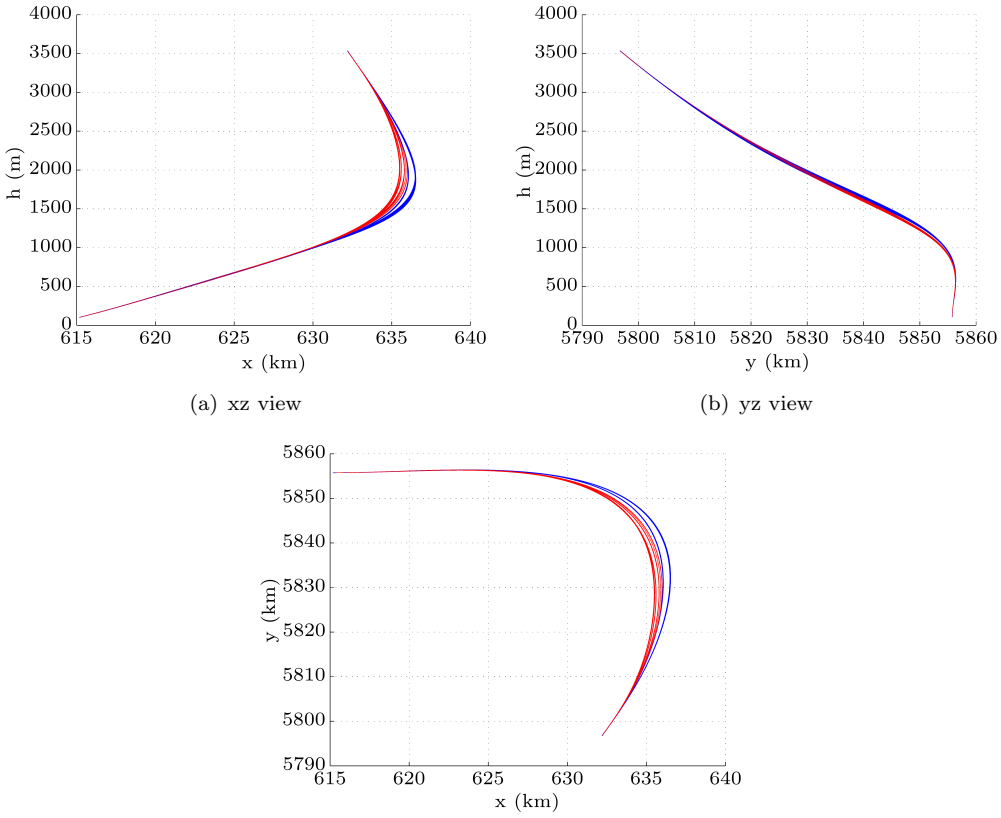

(a) xz view (b) yz view

(c) xy view

FIGURE 8.33: NEMA case: B744 constrained path Pareto front height and path profiles

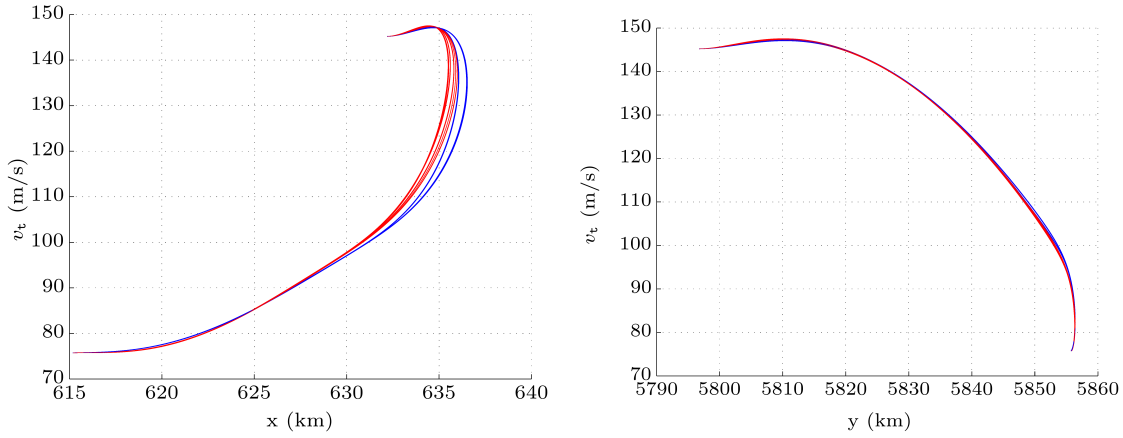

(a) xz view (b) yz view

FIGURE 8.34: NEMA case: B744 constrained path Pareto front speed profiles

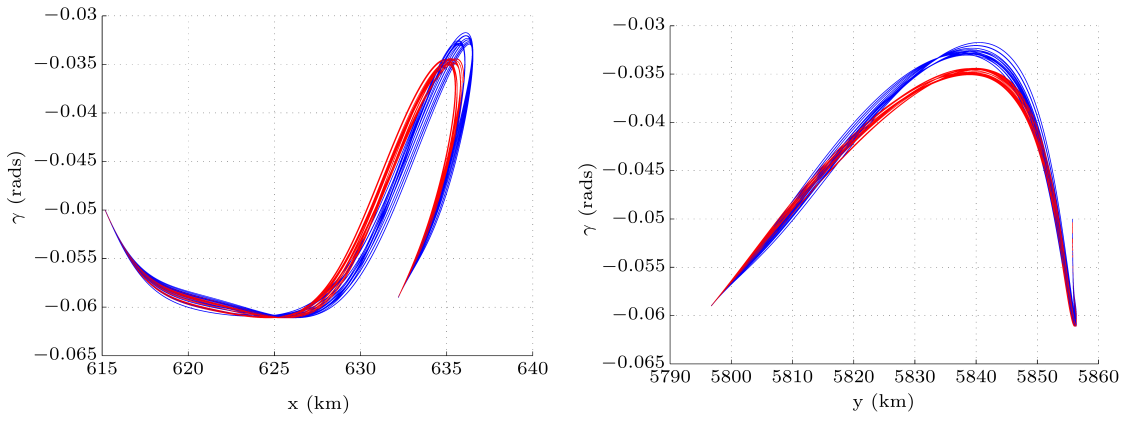

(a) xz view (b) yz view

FIGURE 8.35: NEMA case: B744 constrained path Pareto front γ profiles

8.2.8 Summary

For the NEMA case study, the IDVD-DE method was used to identify approach procedure characteristics that lead to a good balance between noise and fuel reduction goals. These characteristics involved maximising the distance between the route path and the city of Leicester and increasing the radius of the procedure turn onto final approach.

All the Pareto front solutions identified were Continuous Descent Approaches with no level segments. Although CDAs reduce both noise and fuel burn relative to stepped descents, the results showed that the objectives may still not be complementary and that there are trade-offs to be made between the objectives. The results showed that, even while flying a CDA, the management of energy on descent still had a significant impact on the levels of fuel burnt and the levels of noise annoyance experienced by the community. The IDVD-DE optimization also showed that while generic CDA guidance is useful, that to achieve the best trade-offs between the objectives, the height speed and thrust profiles of the CDAs should be tailored to local constraints and population distributions.

Pareto results were generated for the two most dominant aircraft on the route. In general, the two sets of results were complementary and a change to the baseline approach route was proposed with consideration of both these sets. The optimisation of trajectories along the proposed route change converged on a Pareto set of solutions that had low fuel burn and Annoyance Score values relative to the range of values in the original Pareto fronts. The trade-offs between the Pareto trajectories were also significantly reduced. However, the optimisation converged on a subset of solutions that defined a Pareto set of ideal ways of operating the aircraft relative to the objectives, without consideration of how the sensitive the Pareto solutions were to potential changes in the approach procedure. The route may be flown in ways other than those highlighted by the Pareto front trajectories and the sensitivity of the objectives to these potential changes in operating procedure is worth further investigation.

The IDVD-DE generated trajectories compared favorably to the baseline trajectories and to the FDR recorded trajectories presented in the baseline study [165]. The IDVD-DE trajectory results had, in general, faster descent speeds than the trajectories generated by the baseline simulations. This helped reduce flight time and therefore fuel burn. It could also be seen on final approach, that some of the higher speed was traded-off to maintain the descent flight path angle, reducing the levels of thrust required on finals.

Examining the IDVD-DE trajectory results relative to current CDA guidance. The current guidance recommends deploying high lift devices as late as possible to delay the introduction of thrust on final approach. Although high lift devices allow flight at lower flight speeds, the extra lift comes with a lot of extra drag and extra thrust must be used to maintain the descent. However, the IDVD-DE results show that the early introduction of low levels of thrust in advance of the deployment of high lift devices may be used to maintain steeper flight path angles at higher speeds. Therefore when the aircraft does

deploy high lift devices the aircraft is already in a high energy state and requires less thrust response to maintain the descent. It is suggested by these results that the early introduction of thrust and the subsequent reduced thrust response required with the deployment of high lift devices can act to provide an overall reduction in noise and fuel burn relative to current CDA guidance.

The peak noise under the centreline metric was used in the baseline study and was therefore carried over for use in the optimisation study. The inclusion of the measure effected the domination of solutions and therefore the shape of the Pareto front. However, the metric did not provide a good indicator for any sort of community noise impact. The metric simply had the effect of pushing the trajectories away from the noise monitoring points. It is proposed that, in general, the peak noise measure only be used in further optimisation studies where the change in noise values can be tied to noise impact at specific places or communities.

The Pareto results shown in this chapter were generated using the wind speeds from (8.1) and the nominal wind direction defined in Section 8.2.1. The IDVD-DE method was also used to test the flyability of the proposed procedure alteration in peak head and tail wind conditions by generating trajectories along the route for those wind directions. No issues were found in terms of flyability, however the impact of high head and tail winds on the Pareto front and the environmental trade-offs was not investigated. It is thought that the impact of wind direction on the Pareto front trade-offs is a subject worthy of further study beyond the analysis here.

8.3 Luton Airport Departure Procedure Definition Case Study

Terminal Control (TC) North is an arrangement of airspace sectors north of London, UK. As part of a planned airspace redesign of the TC North airspace, the United Kingdom Air Navigation Service Provider (ANSP) NATS proposed environmentally optimising a number of the SID and STAR designs. One of the SIDs chosen, was the Luton Olney SID, a departure route that runs North from Luton Airport [166].

The Olney SID and the proposed changes to it in the TC-North consultation were taken to form an optimisation case study. The environmental impacts of the current and the proposed designs were used as baselines for comparison with solutions proposed by the IDVD-DE method. The aim was to see if the IDVD-DE method could provide better trade-offs between the environmental impacts than those found in the initial study.

8.3.1 Baseline

The current (D0) and proposed (D1) centrelines for the SIDs are shown in Figure 8.36. The principle aim of the change to the proposed route was to reduce noise impact by moving traffic away from the densely populated communities of Milton Keynes (MK) and Leighton Buzzard (LB). As can be seen in Figure 8.36 the route centreline of the current route has been moved such that it no longer passes directly over the population centres. However, to achieve this, the proposed route has been extended, increasing the fuel burn and emissions cost of flying the SID.

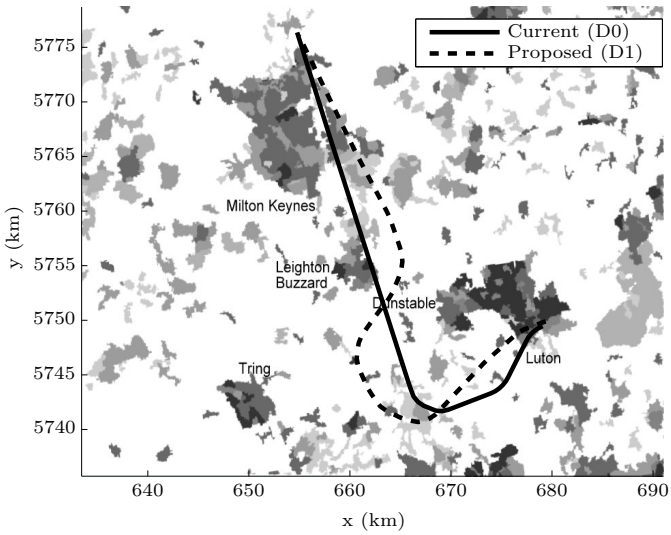

FIGURE 8.36: Current and proposed SIDs from the TC North Airspace Change Proposal

The metrics used by NATS in their redesign were fuel burn/CO2 (as calculated by the BADA model) and the number of people overflown. The number of people overflown was used as a surrogate for noise impact. In this case study a single event simulation was conducted using an Airbus A320, the dominant aircraft on the route. To better asses noise

impact, this case study has adopted 3 noise measures, L_{Amax} at Milton Keynes (MK), L_{Amax} at Leighton Buzzard (LB) and the total Annoyance Score. L_{Amax} was taken as an average L_{Amax} from a grid of points covering each population centre. To balance out the focus on localised noise impact at MK and LB, the Annoyance Score measure was also used to provide an overall community noise impact value. As with the IDVD-DE method, the noise impacts for the baselines were calculated by INM 7 with population data taken from the European Environment Agency database [167]. The fuel burn was calculated using the BADA fuel model.

Table 8.1, shows the results from the baseline scenarios. As can be seen the proposed SID change achieved its stated goal of reducing the noise impact on large communities. However, this goal was only achieved at the expense of greater fuel burn and greater overall community noise impact as indicated by the fuel burn and Annoyance Score results. Figures 8.37(a) and 8.37(b) show the SEL dB(A) footprints for the two baseline solutions. Figure 8.37 shows the contribution each SEL footprint contour level makes to the total Annoyance Score value. It can be seen that the 80 SEL dB(A) contour and the population within the contour make the largest contribution to the total Annoyance Score value. In terms of contribution, the 70 SEL dB(A) contour is the second largest contributer to the Annoyance Score, followed then by the 90 SEL dB(A) footprint contour. It can be seen from Figure 8.37 that, for this scenario, the 70, 80 and 90 SEL dB(A) contours levels dominate the calculation of the total Annoyance Score and that the population distribution provides a natural weighting to each SEL noise level.

Scenario	Annoyance Score	L_{Amax} MK (dBA)	L_{Amax} LB (dBA)	Fuel (kg)
Current (DO)	2.6820×10^{12}	43.4	45.7	644
Proposed (D1)	4.1145×10^{12}	39.1	36.1	677

TABLE 8.1: Luton case: Environmental impact results for baseline procedures

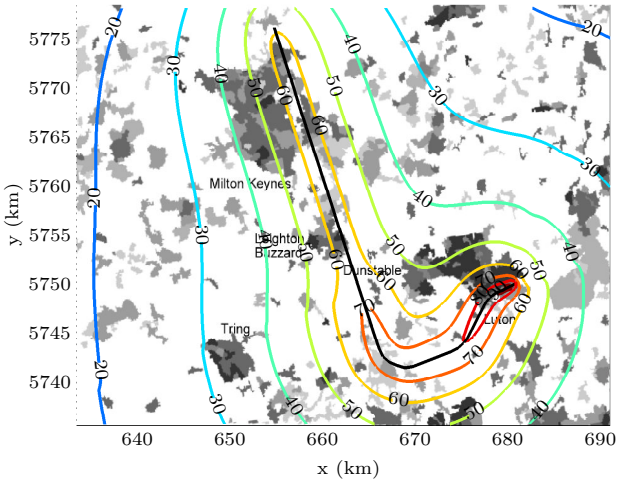

(a) Luton case: D0 trajectory SEL dB(A) footprint

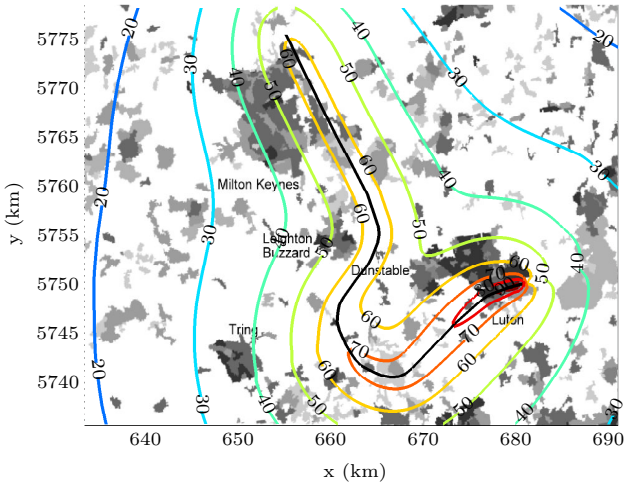

(b) Luton scenario D1 trajectory SEL dB(A) footprint

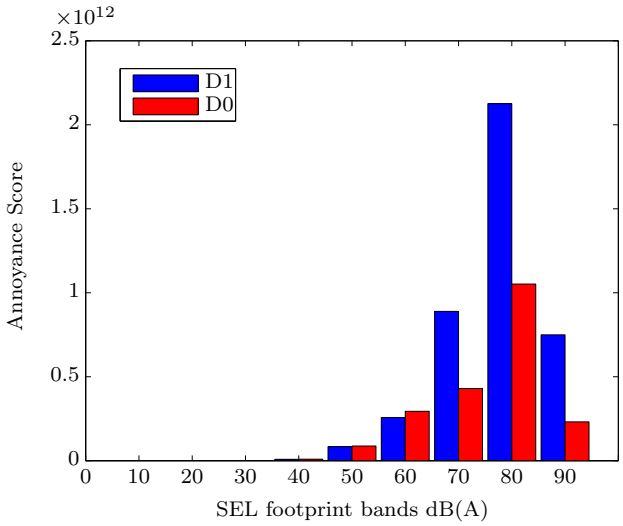

FIGURE 8.37: Luton case: Contribution of SEL levels to the Annoyance Score

8.3.2 Optimisation Results

To determine if the IDVD-DE method could propose a better balance of solutions, a multi-objective trajectory optimisation was run. Figure 8.38 shows the four objective Pareto front results produced by the IDVD-DE simulation. The x axis shows the L_{Amax} (dBA) for Milton Keynes, the y axis shows the L_{Amax} (dBA) values for Leighton Buzzard and the z axis shows the Annoyance Score. The 4th axis represents fuel burn, where the size of each data point of the Pareto front reflects the magnitude of the fuel burnt (kg). The range of point sizes and related fuel burn values are shown to the right of the plot. Figure 8.39 shows the trajectory paths for the Pareto front solutions in Figure 8.38. The red trajectory is the lowest fuel burn trajectory, the orange trajectory is the trajectory with the lowest average peak noise at Milton Keynes, the light blue is the trajectory with the lowest average peak noise at Leighton Buzzard and the green is the trajectory with the lowest Annoyance Score. It can seen from the trajectory paths in Figure 8.39 and from the Pareto front in Figure 8.38 that the four objectives are not complementary and that there are significant differences between the minimum trajectories for each objective. Especially notable is that there are tradeoffs between the three noise objectives and that minimising for noise at large population centres does not automatically guarantee reduced overall community noise impact.

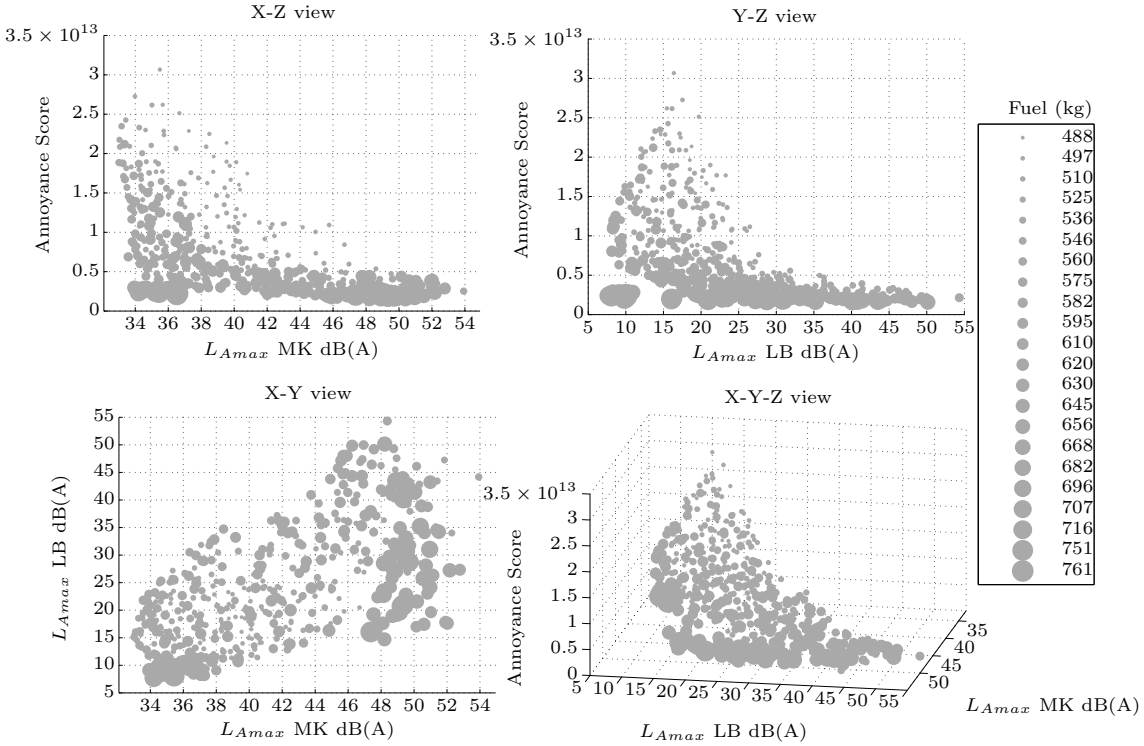

FIGURE 8.38: Luton case: Four view 4D Pareto front for the Annoyance Score, L_{Amax} MK, L_{Amax} LB and fuel burn measures

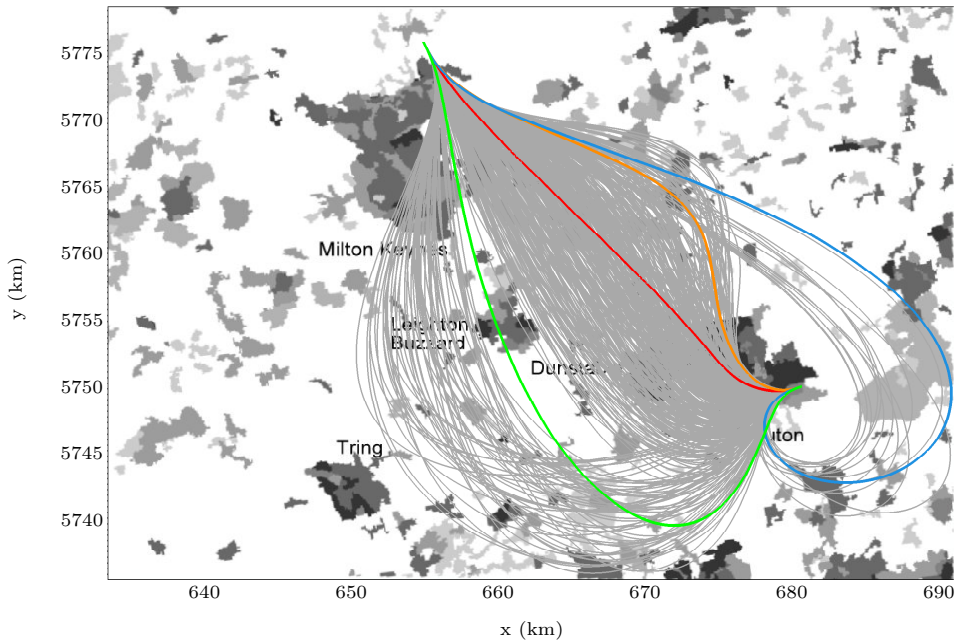

FIGURE 8.39: Luton case: Pareto front trajectories

Evident from Figure 8.38 is that there was a large number of available solutions with Annoyance Scores less than or equal to the existing baseline D0 procedure. Therefore, as an initial cutoff, it was decided that the Pareto front solutions would be constrained to have Annoyance Scores less than $3 \cdot 10^{12}$. This ensured that the solutions identified by the optimisation would at worst only marginally increase the overall community noise impact from the current day noise impact.

The solutions on the Pareto front are clustered according to the extrema solution each point is closest to. The clusters were calculated by measuring the Mahalanobis distance from the extrema solutions to every other solution on the front. Therefore the Pareto front was segmented into 4 individual clusters representing trajectories closest to the lowest fuel burn, Annoyance Score, L_{Amax} at Leighton Buzzard and L_{Amax} at Milton Keynes extrema points respectively.

Figure 8.41 shows the horizontal paths for the Pareto front trajectories, with the trajectories for each cluster plotted independently. Figures 8.42, 8.43 and 8.44 show the clustered height, speed and thrust profiles for the Pareto front solutions. The average height, speed and thrust profile, independently calculated for each cluster using a moving average, is plotted in black on each chart. Range bars are included to show the mean and the extents of data for each profile set.

Low Fuel Burn Cluster (red cluster)
The low fuel burn cluster represents the trajectories with the lowest fuel burn values. The trajectories in this cluster have a path that includes an early turn to the final fix. In this cluster, there is an early focus on maximising speed at the expense of height. Once

250 kts CAS has been reached there is typically a fast climb to the final fix. Although accelerating at low altitudes is expensive in terms of instantaneous fuel burn, it was found that minimising total flight time had a greater effect in minimising total fuel burned.

Low LB L_{Amax} Cluster (light blue cluster)
The low LB L_{Amax} cluster represents the trajectories with the lowest maximum noise impact at Leighton Buzzard. The trajectories in this cluster take a long path that circumvents LB and maximises the slant distance between the aircraft and the community. The longer path distances resulted in longer shallower climbs. Thrust cut-back occurs early and shallow climbs allow low thrust levels to be maintained until close to the final fix. As all the trajectories in this cluster pass directly over Milton Keynes the peak noise levels at Milton Keynes are consequently high for this clusters trajectories.

Low MK L_{Amax} Cluster (orange cluster)
The low MK L_{Amax} cluster represents the trajectories with the lowest maximum noise impact at Milton Keynes. In this cluster the trajectories maximise the distance from MK by turning as early as possible and emphasising height gain at relatively high thrust levels in order to maximise slant distance between trajectories and Milton Keynes. After the initial climb, gains in speed are emphasised and the cluster has an average speed profile that is close to the average of the profile set.

Low Annoyance Score Cluster (green cluster)
The low Annoyance Score cluster represents the trajectories that have the lowest Annoyance Score. Trajectories in this cluster exhibit close to average height, speed and thrust profiles. However, the trajectories in the cluster have two important properties. They all avoid overflying Leighton Buzzard, passing either to the left or right of the population centre. They also avoid overflying the Dunstable community adjacent to Luton. A subset of the cluster's trajectories do overfly central Milton Keynes. However, these trajectories more than the others in the cluster, maximise their distance to Leighton Buzzard and Dunstable. This shows that in order to minimise Annoyance Score, the key in this study was to primarily minimise the higher ($\geq$ 70 dB(A)) SEL values occurring near takeoff and secondarily, the lower SEL values ($<$ 70 dB(A)) impacting communities further away from the departure runway.

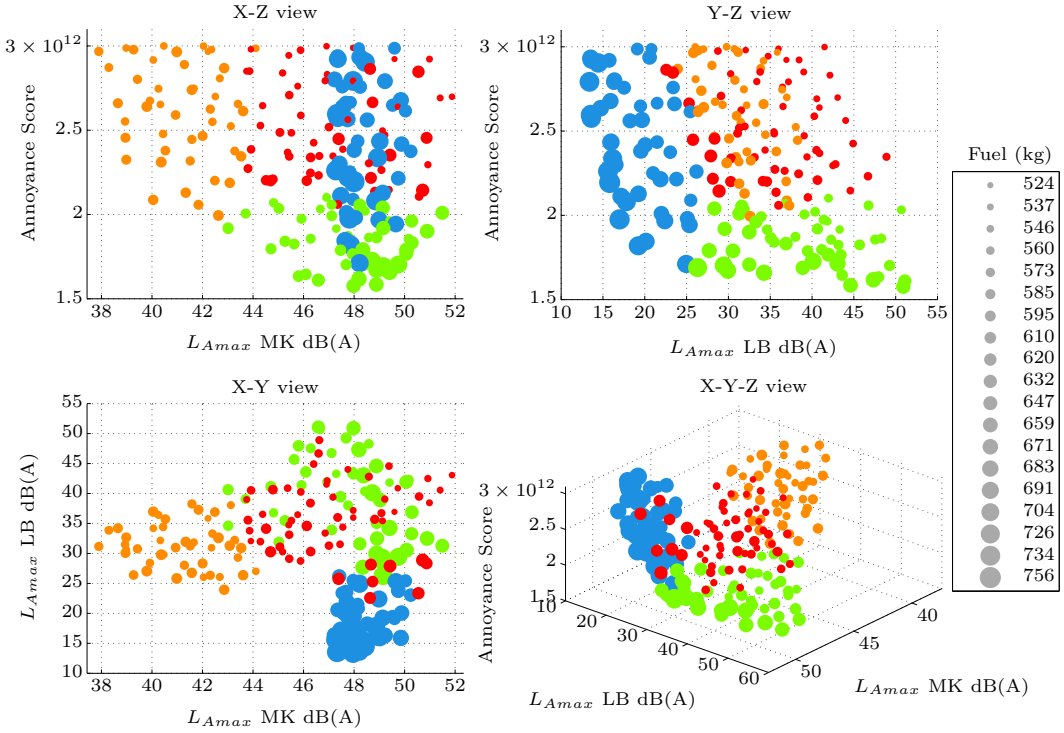

FIGURE 8.40: Luton case: Clustered Four view 4D Pareto front for the Annoyance Score, L_{Amax} MK, L_{Amax} LB and fuel burn measures

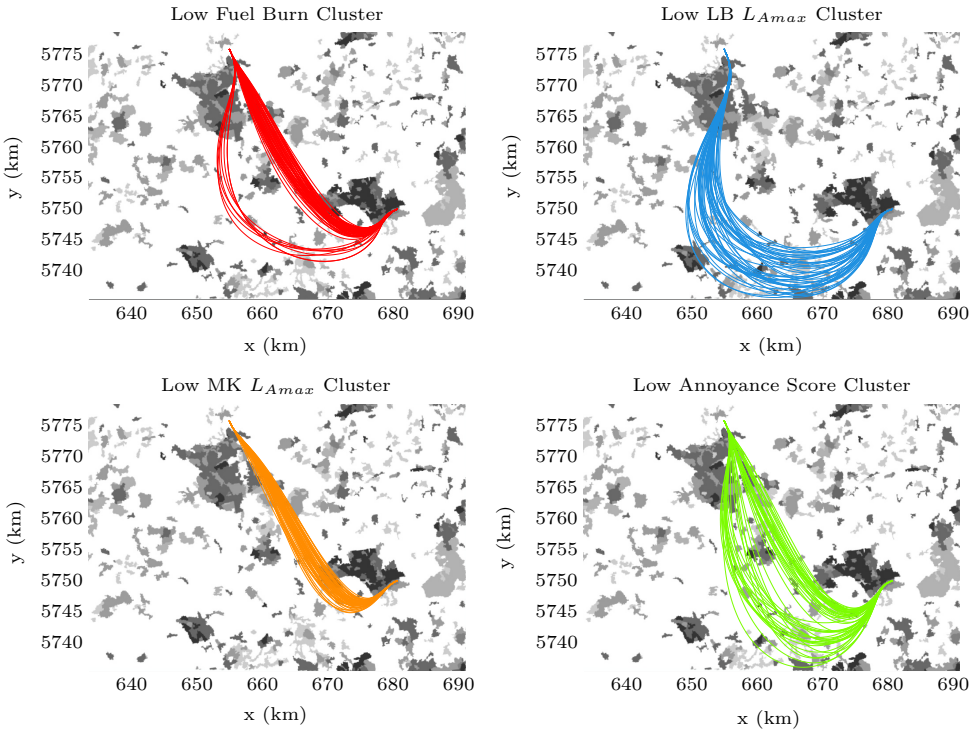

FIGURE 8.41: Luton case: Clustered Pareto front trajectory paths

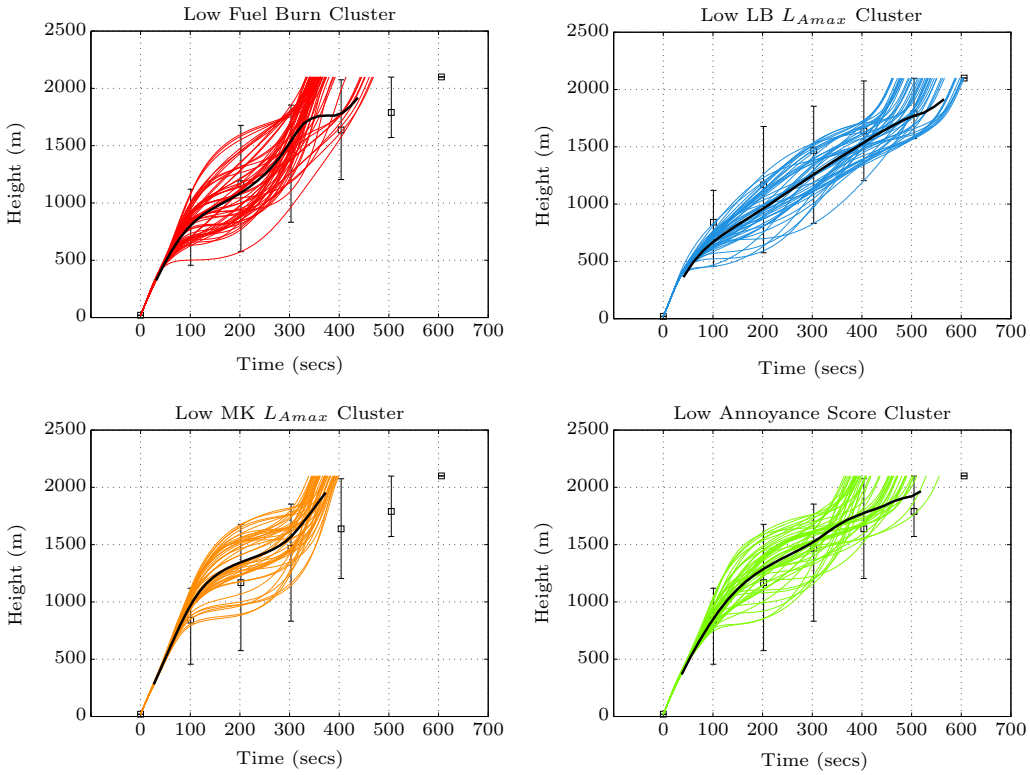

FIGURE 8.42: Luton case: Clustered Pareto front trajectory height profiles

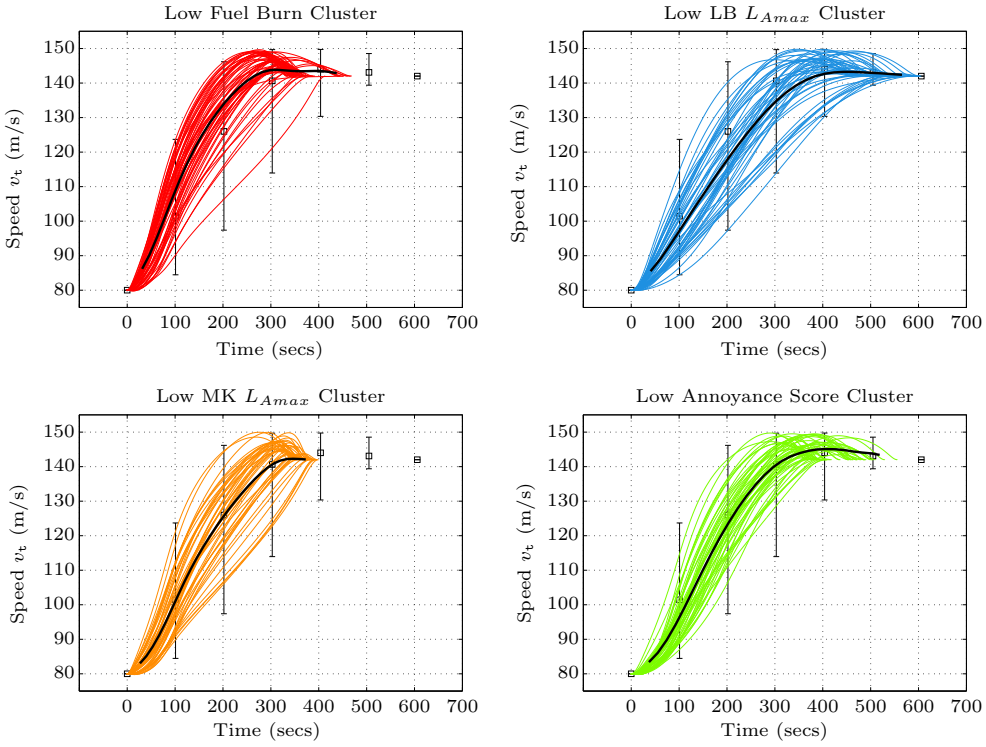

FIGURE 8.43: Luton case: Clustered Pareto front trajectory speed profiles

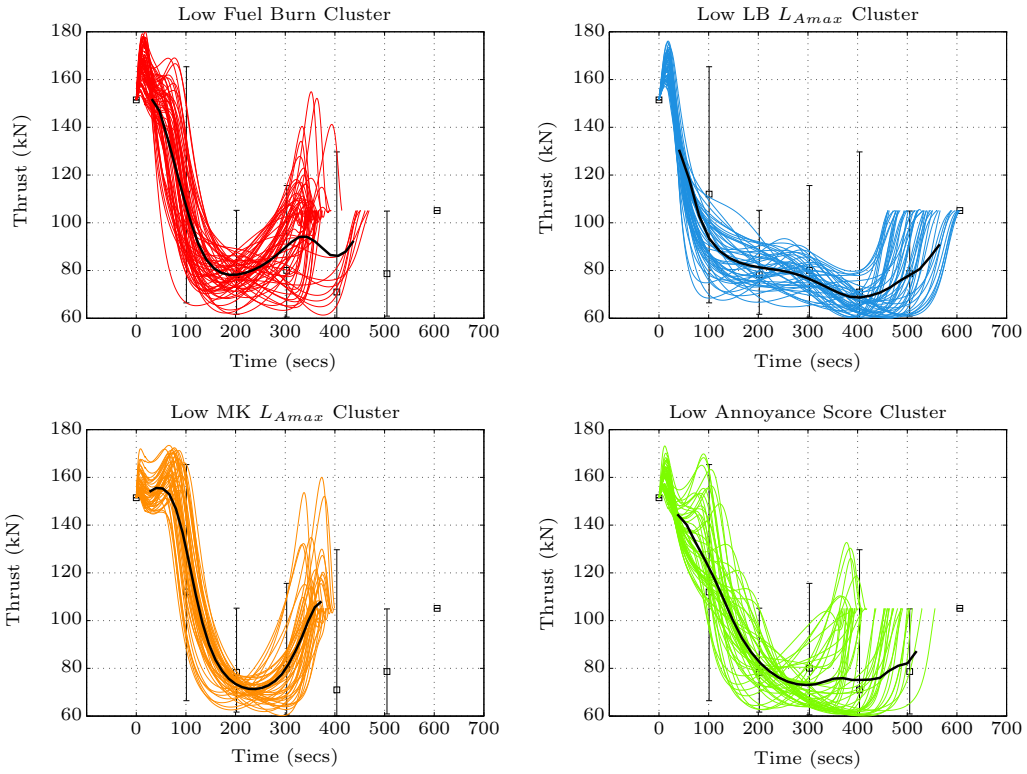

FIGURE 8.44: Luton case: Clustered Pareto front trajectory thrust profiles

8.3.3 Air Quality

The emissions of NOx, HC and CO were not considered as design objectives in the baseline study but are included here in supplement to the Pareto objectives in Figure 8.40. LAQ emissions impact is often only considered to 3000 ft above the airfield as this height defines the start and end of the LTO cycle (see Section 3.3). However, the LTO cycle has generally been a driver for improved engine technology rather than as a driver for improved operations. An ICAO CAEP review of emissions research has also highlighted that LAQ emissions have potential health impacts worthy of consideration above 3000 ft. Therefore, for the results and analysis presented here, the emissions of HC, CO and NOx are for the entire procedure.

Figure 8.45 shows the box plots for the trajectory fuel burn and the emissions of NOx, HC and CO clustered as per the Pareto front clustering. The plots show the actual data values plotted over a box showing the mean and a single standard deviation of the mean.

From Figure 8.45 it can be seen that the lowest fuel burn cluster also has the lowest average emissions values for all of the emissions types. Therefore minimising total trajectory fuel burn clearly has a significant role to play in minimizing emissions. The clustered results for the emissions HC and CO in Figure 8.45 exhibit a direct relationship to the clustered total fuel burn values despite the indices of HC and CO both varying inversely to fuel flow. This is due to other factors influencing the emissions results, principally flight time and thrust settings.

CO and HC emissions are formed more prominently at lower combustion efficiencies occurring at lower thrust levels. Therefore, interpolating off the bilinear curves in Figure 3.2, lower thrust levels and related fuel flows lead to higher emissions indices of HC and CO. Therefore trajectories with the highest total values of HC and CO are those that have the lowest thrust settings coupled with the longest flight duration times. As flight duration time was also strong indicator of fuel burn, there is a correlation between the clustered fuel burn results and HC/CO emissions results in Figure 8.45. Cutback thrust and cutback duration were additional factors that impacted the comparisons between clusters having trajectories with similar flight duration times.

The influence of contributing factors other than fuel flow is especially true in the case of the clustered NOx results, which despite there being a linear relationship between EINOx and fuel flow, do not correlate well with the clustered fuel burn values. NOx is formed as a by-product of fuel burn more prominently at higher combustion temperatures related to higher thrust settings. The high thrust levels on climb out for trajectories in the MK L_{Amax} and Annoyance Score clusters result in more NOx values being interpolated off the higher end of the Log-Log Fuel Flow versus EINOx line in Figure 3.2. This leads to the high total levels of NOx for these clusters relative to the average flight duration time for trajectories in the clusters.

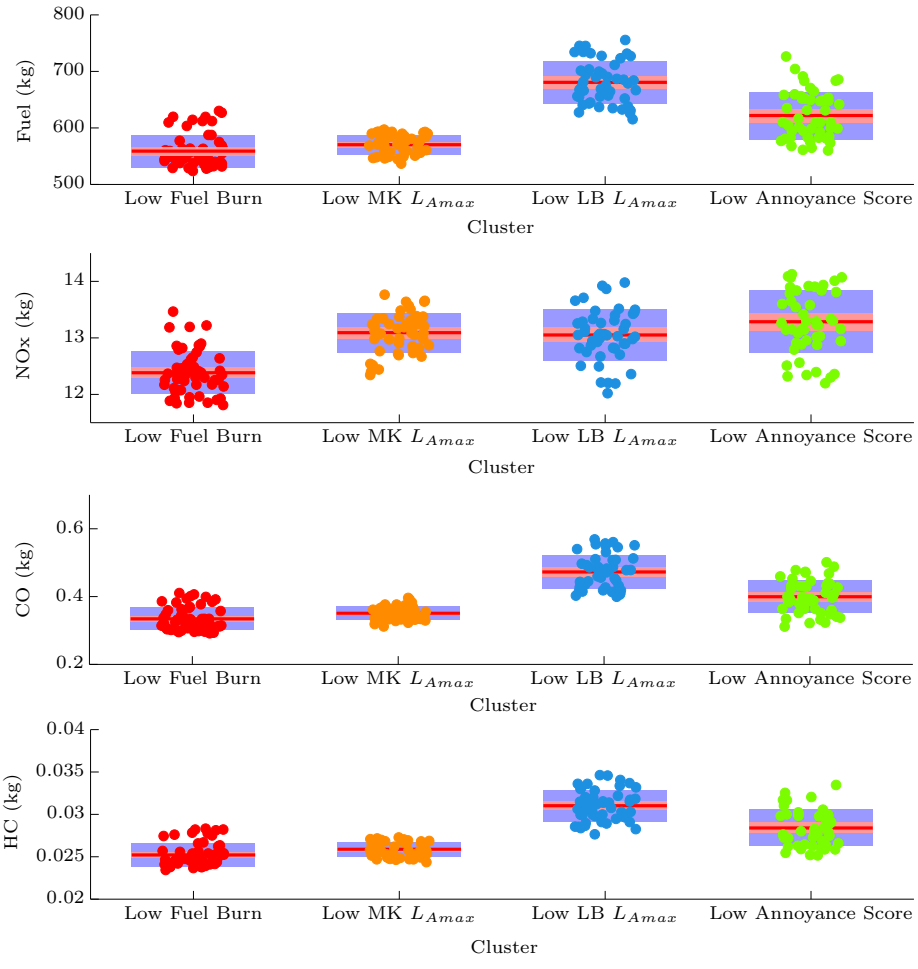

FIGURE 8.45: Luton case: Clustered Pareto front air quality results

8.3.4 Selection

Using the trajectory and Pareto front information in Figures 8.40 - 8.44, it was found that all of the baseline noise results could be achieved at lower fuel burn values. However, there is trade-off between between the Annoyance Score and the peak noise at Milton Keynes making it impossible to have reductions in both measures relative to the baselines.

Figure 8.46 shows a subset of the Pareto front solutions. The baseline solutions are shown in black. The solutions in the subset were largely identified from the Low MK L_{Amax} cluster, with an extra restriction requiring the Annoyance Score values to be no greater than that of the D0 procedure. The subset was chosen as it provided a good balance between the objective values. The region contains points that have Annoyance Scores and fuel burn values less than or equal to the current (D0) procedure, L_{Amax} LB values less than the proposed (D1) procedure, and L_{Amax} MK values less than the D0 route but not less than those proposed by the D1 route. Therefore the solutions in the selected region offer an improvement over both the existing and proposed routes in three of the four objectives and offer an improvement in the remaining objective over the current design.

Table 8.2 shows the mean and standard deviation for fuel burn and air quality emissions for the selected Pareto region. The selected region was identified from the Low MK L_{Amax} cluster and has fuel burn, NOx, CO and HC emissions that are equivalent to that cluster. Therefore, it can be seen from Table 8.2 and Figure 8.45 that the selected region has NOx, CO and HC values in line or less than the other noise minimising clusters, with only the low fuel burn cluster having better mean values for all emissions.

Metric	Fuel (kg)	NOx (kg)	CO (kg)	HC (kg)
SD	16.5	0.31	0.022	8.66E-04
Mean	576	13.1	0.354	0.0261

TABLE 8.2: Luton case: Mean and standard deviation for fuel burn and air quality emissions for selected Pareto region

The path, height, speed and thrust profiles for the selected region are shown in Figures 8.47 and 8.48. The paths involve a early turn to the final fix with most trajectories passing directly over the Dunstable community. Prior to passing this community, aircraft thrust is held high for an extended period. This allows the aircraft height over Dunstable to be maximised while maintaining a typical speed schedule. Thrust cutback then occurs over the community to levels that are below average relative to the Pareto solutions in Figure 8.40. The combination of maximised height and low thrust acts to minimise noise on the Dunstable community directly below the route path. Thrust levels as the aircraft subsequently passes Leighton Buzzard, Milton Keynes and intervening smaller communities are maintained low but sufficient for a small continuous acceleration. Once most of Milton Keynes is passed and the target speed for the final fix is reached, higher levels of thrust are reintroduced and the aircraft completes its climb to the final fix.

Figure 8.46 shows that both the D0 and the selected Pareto front procedures provide lower Annoyance Scores than the D1 procedure. This is because the D1 procedure avoids the major population centres only at the expense of using a longer path that ultimately results a larger number of people in smaller communities being exposed to aircraft noise. The D0 procedure has a lower Annoyance Score than the D1 procedure. The D0 procedure path initially takes aircraft away from Luton and Dunstable but does not do so at the expense of passing through a large number of smaller communities and therefore the D0 procedure minimises the > 70 dB(A) SEL levels that dominate the Annoyance Score results. The earlier turn to the final fix of the D0 procedure also reduces the procedure fuel burn relative to the D1 procedure. For the Pareto solutions in the selected region, Figure 8.47 shows that all have procedure paths that turn to the final fix before both baselines, providing further reductions in fuel burn results. Unlike the baselines, the Pareto paths pass directly over the densely populated Dunstable community. However, in this case, the vertical trajectories are used to minimise the noise impact on Dunstable in such a way as to achieve lower Annoyance Scores than that achieved by the D0 procedure. In addition, the earlier turn allows the procedure paths to maximise the distance to the right of Leighton Buzzard and Milton Keynes, allowing for the reduction in peak noise values at those communities relative to the D0 procedure and reductions in the peak noise values at Leighton Buzzard relative to the D1 procedure. The subset of Pareto front solutions implicitly define a range of paths and profiles within which the gains in the objective values remain valid. This is useful when defining the procedure extents, and for defining the aircraft and navigation performance required to fly the procedure.

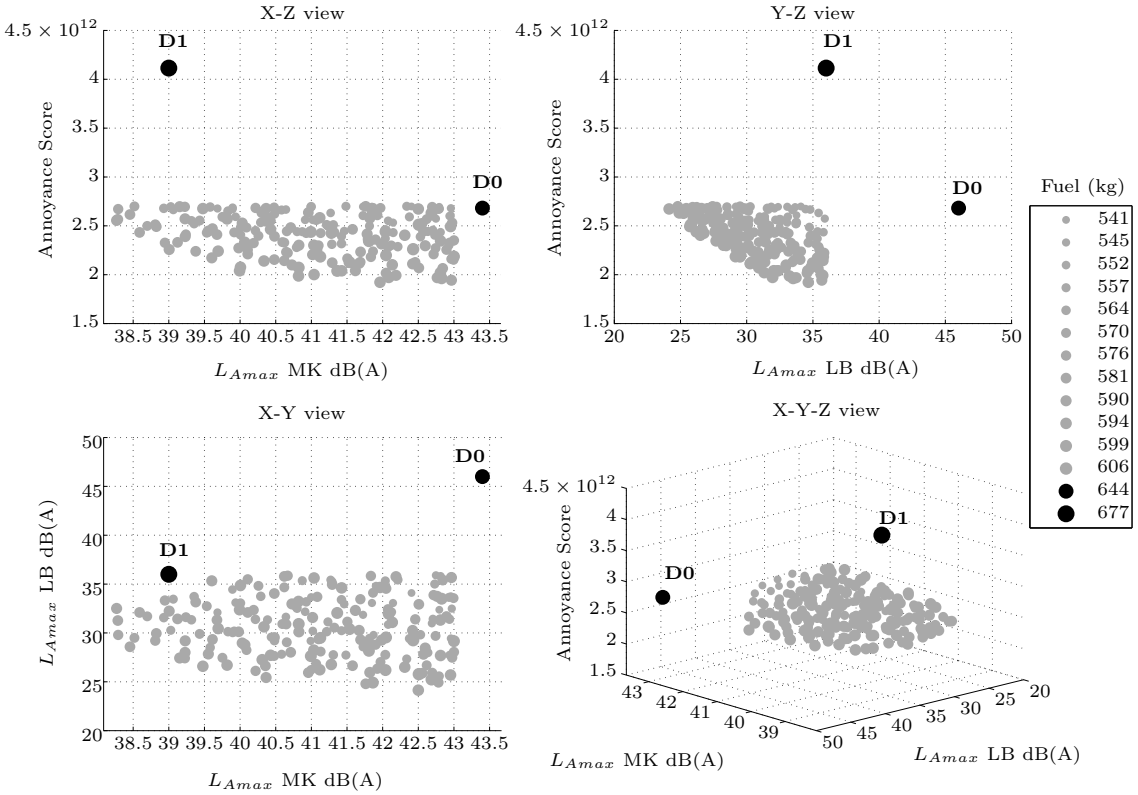

FIGURE 8.46: Luton case: Four view selected 4D Pareto front with baseline solutions

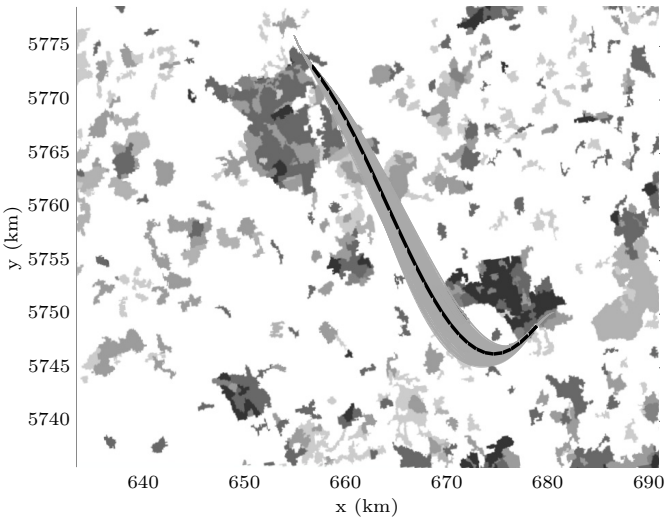

FIGURE 8.47: Luton case: Selected Pareto front trajectory paths

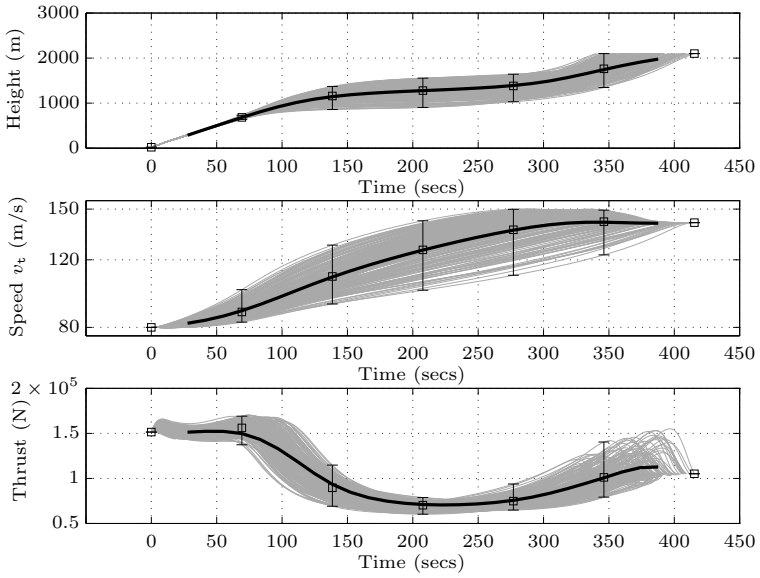

FIGURE 8.48: Luton case: Selected Pareto front trajectory height, speed and thrust profiles from region of interest

8.3.5 Conversion to Procedure

The Pareto optimal results are intended to provide decision support to route designers considering the placement of departure routes and the definition of related height and speed constraints. The results can be used to provide insight into the trade-offs between the competing environmental objectives. The results were produced with consideration of the dominant aircraft on the route. However SID procedures must be designed to serve a number of different aircraft types with differing navigation equipage that impact the aircraft's ability to fly a given procedure. To ensure compatibility, any recommended trajectory set must be converted to a procedure in line with the procedure design guidance in [37] and validated with further simulations and potentially flight trials. In some instances

the multi-use requirement of the procedure may limit the environmental efficiency of the trajectory achievable for any specific aircraft type. In other instances, the procedure constraints may be wide enough to accommodate a wide range of possible trajectories and therefore the emphasis shifts to the aircraft operator to fly the most environmentally balanced trajectory within the procedure constraints.

8.3.6 Summary

This case study proposed the use of the IDVD-DE method to provide a harmonised approach to the planning of environmental abatement climb operating procedures. In the section the proposed methodology was applied to a real world case study. The study highlighted the need for a multi-objective optimisation method that could be applied to problems with more than two objectives. The methodology is designed to support designers in identifying procedures that provide the best trade-offs between these sometimes conflicting environmental objectives.

The method was then used to identify a range of ground paths and related vertical trajectories that, as a group, offered an improvement over both baselines procedures in three of the four objectives and offered an improvement in the final objective over the existing procedure design. Of the metrics used, the Annoyance Score is the least mature and further research is required to determine what constitutes significant changes in Annoyance Score.

8.4 Conclusions: Multi-Objective Procedure Optimisation

This thesis proposes a methodology intended to produce results that support decision makers seeking to plan environmentally efficient climb and descent procedures. The methodology is intended to be used alongside the 3 ICAO environmental goals of reducing community noise impact, reducing local air quality emissions, and reducing climate changing emissions.

In this chapter the IDVD-DE approach was applied to two real world procedure optimisation case studies. The baselines procedures in each study were reached through expert opinion, iterative analysis and in the case of the NEMA work, flight trials. The procedure design case studies presented in this chapter are not intended as definitive results but rather are offered as demonstrations of how control and Pareto based optimisation can be used to define a data driven approach that can be used to seek improvements in aircraft operating procedures. However, it is possible that expert flight and airspace planners would not make the same decisions and therefore not reach the same conclusions based on the same data as those reached in the case studies. Therefore, the application of the IDVD-DE method does not remove the need for expert analysis, but it is shown here

how trajectory optimisation results and in particular Pareto based results can be used to inform and guide such an analysis from the outset.

In the Luton departure and NEMA arrival procedure case studies, the Pareto fronts were used to identify a full range of potential procedure options and to show how each option affected the trade-offs in the environmental objectives. It is intended that this information can be used by airspace and flight procedure planners to guide decision making, allowing them to identify specific procedures or procedure characteristics that are judged to result in the best balance of environmental impacts for each specific study. The IDVD-DE method is presented as a general approach that can tailored to specific local conditions. This, it is intended, will support a very data driven assessment of the environmental trade-offs allowing procedures to be defined in an objective manner that reduces reliance on intuition and iteration.

In the Luton departure procedure definition case study, it was found that avoiding population centres in itself was not an effective strategy for reducing overall levels of community noise impact. Using the IDVD-DE approach, a set of procedure options were proposed that involved a more direct flight path routing to the target fix, reducing fuel consumed relative to the baseline procedures. Overall community noise was also reduced below the baseline procedures by emphasising the management of the vertical profile along the more direct flight path routing.

The trade-offs among LAQ emissions and between LAQ emissions and other forms of environmental impact were explored as part of the Luton departure case study. However they were not used as objectives and were not the subject of a dedicated case study. This was due to the unavailability of real world LAQ procedure optimisation studies that could be used as a reference or benchmark for comparison to the IDVD-DE solutions. It was found, in the literature, that an assumption often made, is that LAQ emissions reduce in line with fuel consumption. Therefore fuel/CO2 is usually chosen to be the design objective. The Luton case study showed that while this was not an unreasonable assumption, that closer examination of the data will show trade-offs existing between LAQ emissions and fuel consumption.

For the NEMA case study, the IDVD-DE method was used to identify approach procedure characteristics that lead to a good balance between noise and fuel reduction goals. These characteristics involved maximising the distance between the route path and the city of Leicester and increasing the radius of the procedure turn onto final approach.

The shallower bank angles facilitated by the increased turn radius allowed aircraft to commence the final approach leg at a higher energy state. This allowed thrust levels to be managed in a manner that mitigated fuel consumption and noise impact on final approach. The work also showed how simple cylindrical constraints could be used effectively to approximate complex airspace structures.

The results of the NEMA optimisation study however were complicated by the inclusion of the noise under the centreline measure as an objective in the IDVD-DE optimisation. The use of the metric had a negative impact on the distribution of the Pareto front solutions affecting the ease with which conclusions could be drawn from the data set.

The metric however was retained in the study to maintain consistency with the baseline study. The results as is were also considered to provide a useful illustration of the how the choice of metrics and the design of experiment can effect the Pareto front solution sets.

The environmental impact analysis of departure and approach procedures forms only a part of the procedure development process. While the work here considered typical safety constraints such as minimum cut-back height and ILS descents, the realisation of an implemented operating procedure requires the consideration of further safety, human factor, political and CNS infrastructure factors. These factors were beyond the scope of this research, but can have a significant affect on the implementability of an environmentally optimised procedure.

8.4.1 Resource Commitment

It would be impossible to estimate the time taken to analyse the case studies in this chapter. The case studies were an integral part of the research, development and debugging of the IDVD-DE methodology and algorithms. They set the goals and the constraints for the optimisation scenarios and provided baseline solutions against which both algorithm implementation and performance were evaluated.

As it stands, the multi-objective optimisation results presented in this chapter and in Chapter 7 took between 8-16 hours, with objective numbers from 2 to 4 and solution numbers from 70-500. The maximum number of generations chosen ranged between 10,000 and 20,000 with trajectory discretisation between 1000 and 6000 nodes. Optimisation scenarios were generally run over-night on Intel Core i7 (3.0GHz, 8 core) 16GB of RAM machines. In all cases, 500 by 500 meter noise evaluation grids were used. However, the IDVD-DE method was implemented as a MATLAB prototype not fully optimised for run-time performance. Therefore, it is expected that substantial improvements in run time performance are possible. This is discussed further in the thesis conclusions in Chapter 10.

The analysis approach demonstrated in this chapter was found to be a necessarily long-form approach. Each trajectory solution is itself composed of state and control profiles, including height, speed, thrust, aerodynamic and flight path routing histories. Therefore there is a significant amount of information to parse to determine which operational factors are driving the significant changes in which environmental objectives. This, for the case studies here was necessarily a time consuming task due to the newness of the methods employed. Evaluating nonintuitive solutions was also similarly very time consuming.

It is anticipated that to be useful, analysts would need to be able to draw a complete set of conclusions from the results in at most a few hours. Therefore there is a need for further tools to reduce the time and effort required for post optimisation analysis. However, it is expected that continued research would quickly mature the optimisation approaches and techniques for analysing the Pareto fronts. With the concept mature, it is expected that such fast turnaround times for optimisation studies becomes very realisable.

Chapter 9

Trajectory Based Flight Planning

9.1 Introduction

In current day operations, the 4D trajectory must be segmented into a series of discrete operational steps in order to be realised by the flight crew and on-board systems (see Chapter 8). However, as discussed in Section 2.6.2, for trajectory based operations, the flight crew and on-board systems are required to track to the 4D trajectory and there ceases to be any difference between the operating procedure and the continuous time 4D trajectory. As the flight tracks to the 4D trajectory, there is no longer a need to utilise a cost index, as the operators cost priorities become implicitly defined by the 4D trajectory.

From the ATM perspective, 4D trajectories can be used to accurately schedule limited airport and airspace resources, allowing precise application of ATM constraints and only when necessary to maintain the safe and expeditious flow of air traffic.

For these reasons, in trajectory based operations, the planning of the 4D trajectory takes on a heightened importance for the aircraft operator in the Business Development Trajectory (BDT) phase of the SES reference trajectory development life-cycle (see Section 2.6.2). Therefore, while previous Chapters focused on the trade-offs in environmental impacts within in the TMA, this chapter will focus from the operators perspective, on planning environmentally efficient flight trajectory within predefined ATM constraints.

Work in this thesis has always looked to adopt real world studies and reference solutions for comparison to IDVD-DE optimisation results. However, as SES trajectory based operations do not exist yet, there was a need to find an alternative reference for comparison. The NATS 3D inefficiency (3Di) score is a flight efficiency method based around the SES user preferred trajectory concept. The 3Di method is used to asses CO2 inefficiencies in flight operations by comparing actual flown flight trajectories to a library of ideal planned trajectories. The ideal planned trajectories of the 3Di method are used in this chapter to define a current best practice trajectory flight planning approach. The ideal 3Di planned

trajectories and the 3Di score itself are used for relative comparison to the CO2 efficient planned flight trajectories generated by the IDVD-DE method.

9.2 Energy Height

Before applying the IDVD-DE method to a more realistic flight planning problem, it was first applied to a simple flight planning problem. The simple flight planning problem is the Minimum Time to Climb (MTTC) problem. For the same problem, the IDVD-DE method was then compared to the Energy Height method, which is method known to provide a good approximate solution to the MTTC problem [168].

The Energy Height method for trajectory optimisation is detailed in the Engineering Sciences Data Unit (ESDU) implementation guidance documents 90012 [168] and 91016 [169]. The method involves parameterising aircraft speed and height as a single energy height parameter H_e, where the aircraft energy is defined as the sum of the kinetic and potential energies such that

$$H_e = h + \frac{v_{\mathtt{t}}^2}{2g} \tag{9.1}$$

and where h is height, $v_{\mathtt{t}}$ is true airspeed and g is gravitational acceleration. It can then be shown that the minimum time to climb from one energy height level to another occurs at the maximum rate of change of energy height at each energy height level. Specifically the problem becomes the minimization of the integral

$$t = \int_{He_2}^{He_1} \frac{dt}{dHe} dHe = \int_{He_2}^{He_1} \frac{1}{\frac{dHe}{dt}} dHe \tag{9.2}$$

where the integral will be a minimum if at all energy heights

$$\frac{\delta}{\delta v_{\mathtt{t}}} \left(\frac{dt}{dH_e} \right) = 0 \tag{9.3}$$

and

$$\frac{dH_e}{dt} = \frac{v_{\mathtt{t}}(T - D)}{mg} \tag{9.4}$$

where T is thrust, D is drag and m is mass.

The Energy Height method for the solution of the MTTC problem can be visualised in Figure 9.1 for an Airbus A321 by plotting lines of constant energy height and constant $\frac{dH_e}{dt}$ on a $v_{\mathtt{t}}$, h chart. Observing where the contour lines of maximum $\frac{dH_e}{dt}$ intersect each energy height, the locus of points at which the excess power is at a maximum provides the $v_{\mathtt{t}}$, h profiles for the minimum-time to climb trajectory.

For the actual implementation of the Energy Height method, ESDU 91016 proposes a direct search approach for finding the minimum time to climb trajectory. The method requires an algorithm to step iteratively along each line of constant energy height to determine the point at which maximum $\frac{dH_e}{dt}$ occurs. For subsonic aircraft there is only one point for each line. The locus of points of maximum $\frac{dH_e}{dt}$ at each energy height then defines the MTTC height and speed profiles for the aircraft. The MTTC trajectory produced by the ESDU 91016 method for an Airbus A321 from sea level to a chosen cruise level of 29000 ft is shown in Figure 9.2. Figure 9.3 then shows the MTTC trajectory calculated by the ESDU Energy Height (ESDU-EH) method relative to the contours of constant H_e and $\frac{dH_e}{dt}$. It can also be seen in Figure 9.3 that there is good agreement between the theory and the implementation of the method.

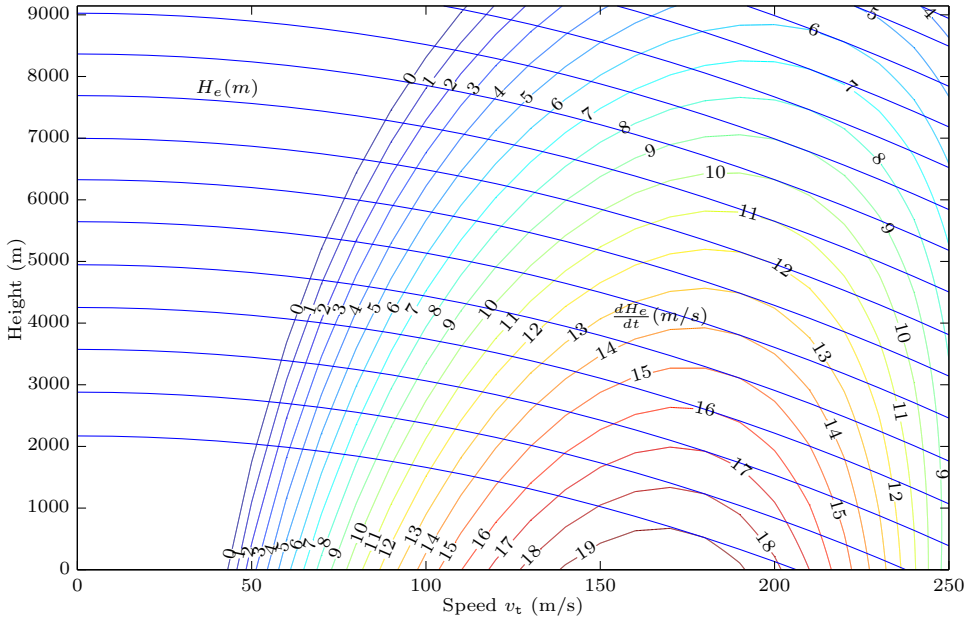

FIGURE 9.1: MTTC case: Contours of constant H_e and $\frac{dH_e}{dt}$

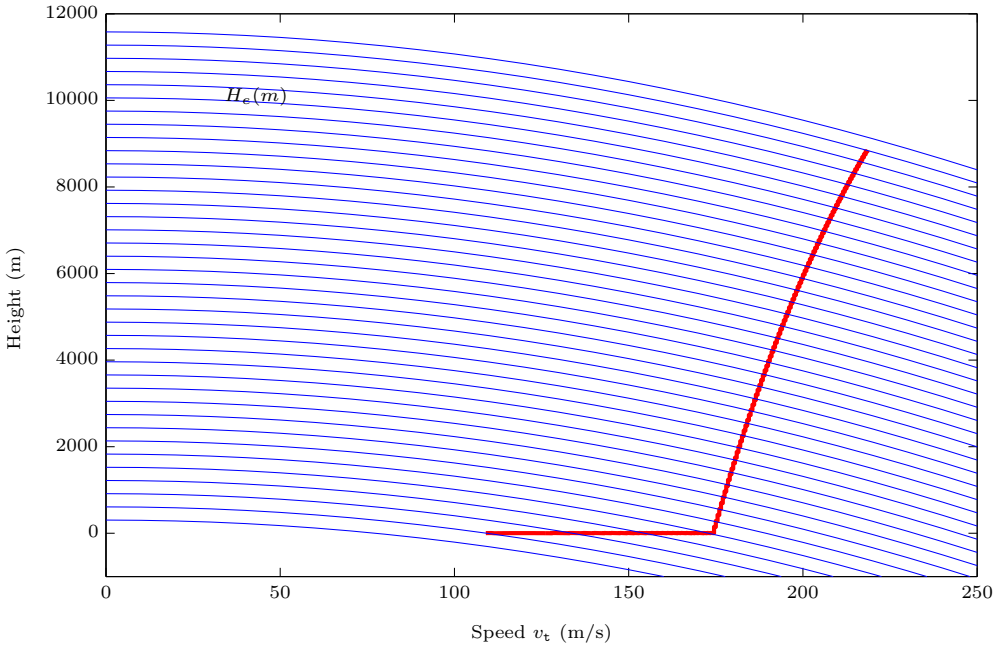

FIGURE 9.2: MTTC case: MTTC trajectory calculated by the ESDU Energy Height method (ESDU-EH)

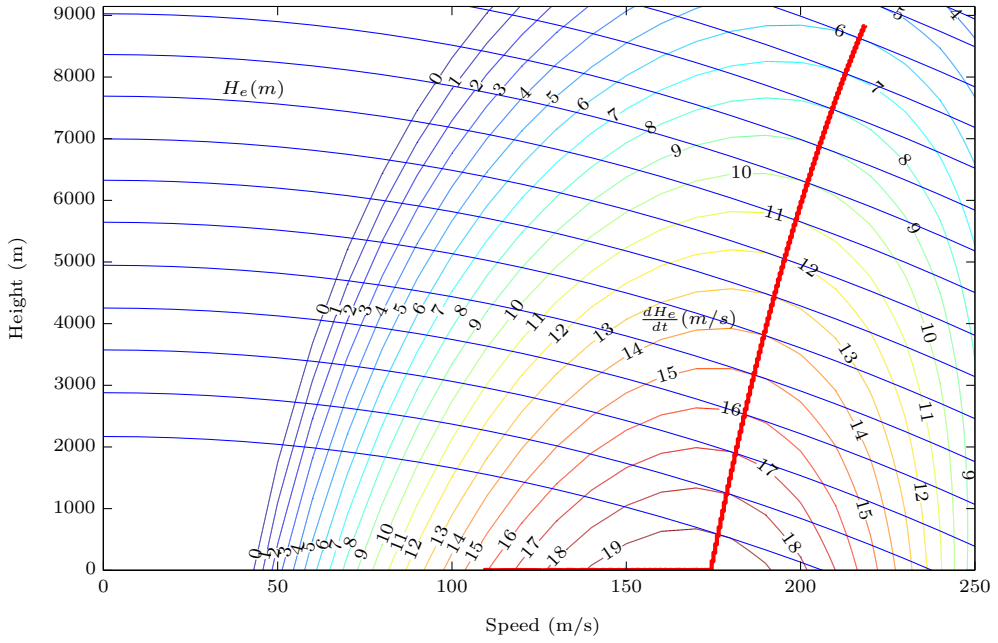

FIGURE 9.3: MTTC case: ESDU-EH MTTC solution relative to contours of constant H_e and $\frac{dH_e}{dt}$

Figure 9.4 shows the time parameterised height, speed, thrust and rate of climb profiles for the ESDU-EH solution to the MTTC problem. The profiles show that the MTTC solution has an initial acceleration at sea level followed by a fast subsonic climb to the target height. This solution occurs as the Thrust to Drag relationship at lower heights favours increases in speed until the speed and therefore the drag are increased to levels where the largest increases in $\frac{dH_e}{dt}$ are achieved by climbing the aircraft. The solution is typical of the optimum MTTC trajectory solution calculated by the ESDU-EH method [168, 169]. However, the Energy Height method merely connects a series of discrete points to define the MTTC trajectory and the method does not attempt to accurately simulate the dynamics of a transition from one discrete point to the next. Therefore the ESDU-EH method can only be considered to be a simple approximation to the true optimal solution to the MTTC problem.

Figures 9.5 to 9.10 compare the ESDU-EH solution to the IDVD-DE trajectory solutions for the same A321 MTTC problem to 29000 ft. The boundary vales used for the IDVD-DE optimisation are shown in Table 9.1.

Time	$x(m)$	$h(m)$	$v_{\mathtt{t}}(m/s)$	$\gamma_{\mathtt{i}}(rad)$	$\chi_{\mathtt{i}}(rad)$
t_o	0	5	109	0.145	0
t_f	176234	8838	218	0.025	0

TABLE 9.1: MTTC case: Scenario boundary values

The trajectory solutions for the 2 methods were very close with the ESDU-EH method calculating a minimum time to climb of 904 seconds and the IDVD-DE method converging on a minimum time to climb of 909 seconds. Figures 9.5 and 9.6 shows the energy height H_e and the rate of change of energy height $\frac{dH_e}{dt}$ histories for both trajectories. The mean difference between the energy height histories was 34 metres and the maximum difference was 46 metres. The mean and maximum differences between the $\frac{dH_e}{dt}$ histories was 0.13 m/s and 1.06 m/s respectively. In Figure 9.6, over the time interval 0-100 seconds, it can be seen that the IDVD-DE solution does not utilise the maximum available excess power. This was due to the dynamics constraints on the accelerations and on the rate of change of flight path angle $\dot{\gamma}$ for the IDVD-DE method. The effect of this difference can be seen in the trajectory height profiles in Figure 9.7, where the IDVD-DE trajectory rotates earlier, gaining height sooner, while accelerating slower than the ESDU-EH trajectory.

It can be seen in Figure 9.8 that the IDVD-DE speed profile is faster than the ESDU-EH profile for times between 200 and 900 seconds. The IDVD-DE method had 10kN more available thrust than the ESDU-EH method to allow for easier satisfaction of the maximum thrust constraint and to support the search for feasible solutions. This allowed the IDVD-DE to maintain a slightly faster speed profile than the ESDU-EH solution at some points along the trajectory path.

Figure 9.9 shows the marginally higher levels of available thrust for the IDVD-DE method relative to the ESDU-EH method. The biggest difference however between the two methods can be seen in rate of climb profiles in Figure 9.10. Here, where both aircraft are

climbing, the ESDU-EH method can be seen to have higher attainable climb rates than the IDVD-DE solution. Once in climb, The ESDU method assumes that all excess power is used to climb the aircraft while accelerations are achieved instantaneously at each energy height level. For the IDVD-DE method, accelerations are modelled more realistically and the excess power must be split between accelerating the aircraft and gaining height. Therefore, to maintain increases in speed, the IDVD-DE cannot avail of the maximum theoretical climb rate at each energy height level.

In summery, both methods for calculating solutions to the MTTC problem, provide very close agreement between their trajectory results. Due to constraints on the dynamics of the aircraft, it is thought that the IDVD-DE provided a more realistic approximation to the optimum flyable MTTC trajectory. The ESDU-EH method could be supplemented with dynamics constraints to provide a closer match to the IDVD-DE solution, but doing so was considered beyond the needs of this work.

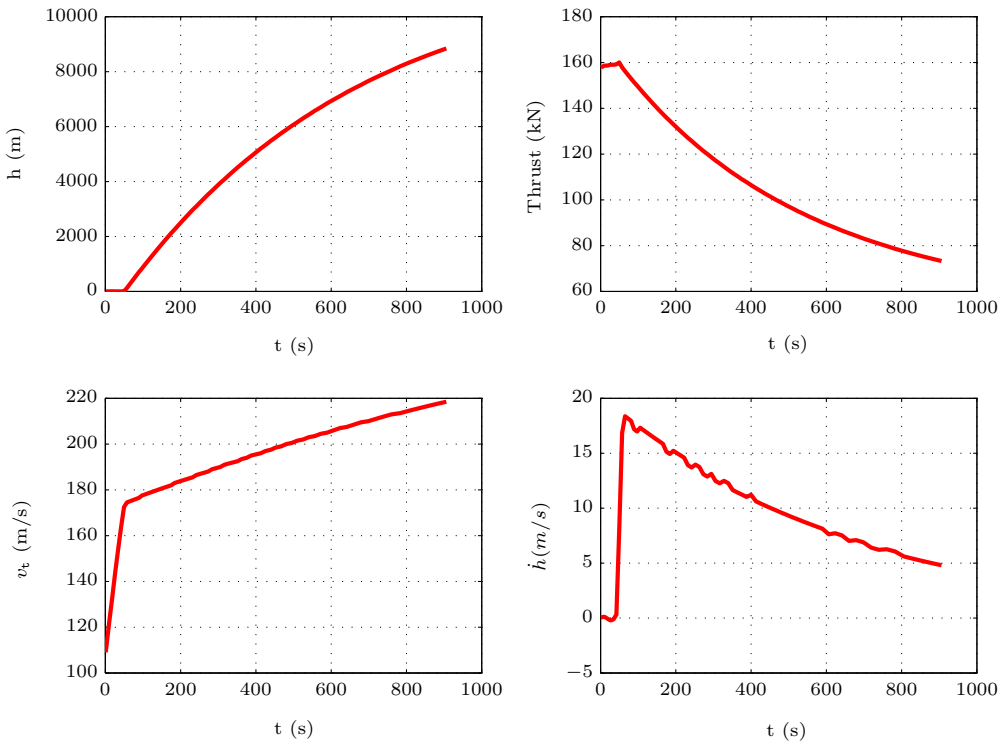

FIGURE 9.4: MTTC case: Time parameterised height, speed, thrust and rate of climb profiles for ESDU-EH solution

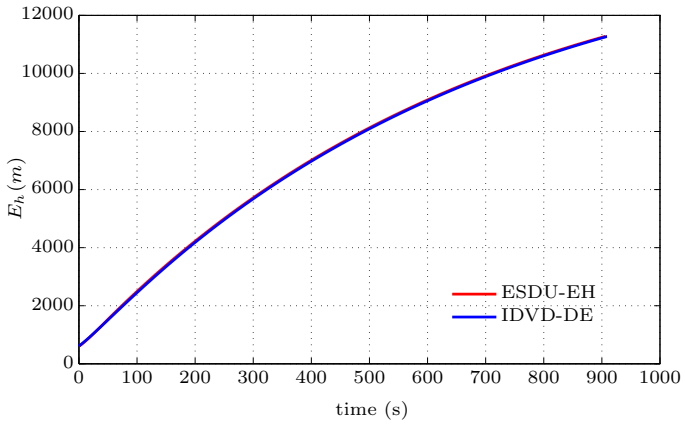

FIGURE 9.5: MTTC case: Comparative H_e profiles for the IDVD-DE and the ESDU-HE methods

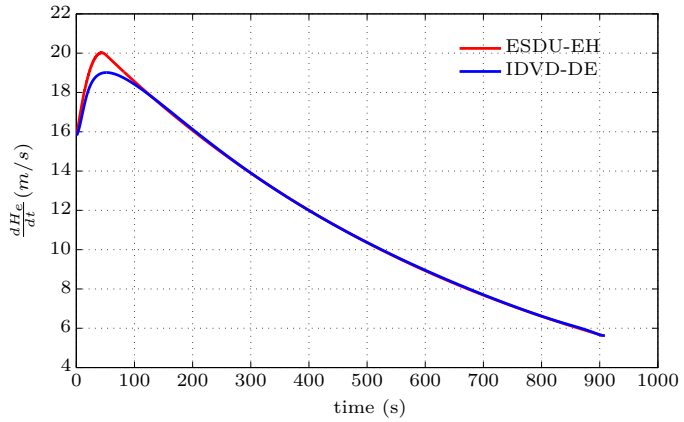

FIGURE 9.6: MTTC case: Comparative $\frac{dH_e}{dt}$ profiles for the IDVD-DE and the ESDU-HE methods

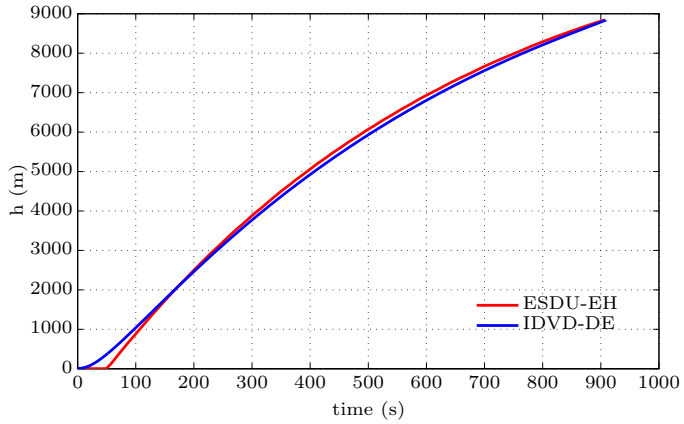

FIGURE 9.7: MTTC case: Comparative trajectory height profiles for the IDVD-DE and the ESDU-HE methods

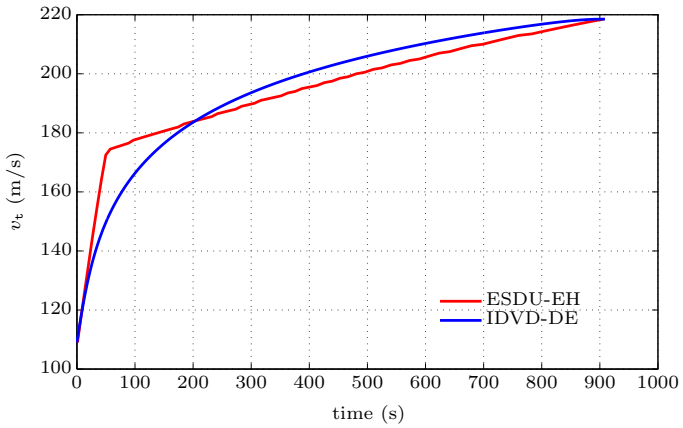

FIGURE 9.8: MTTC case: Comparative trajectory speed profiles for the IDVD-DE and the ESDU-HE methods

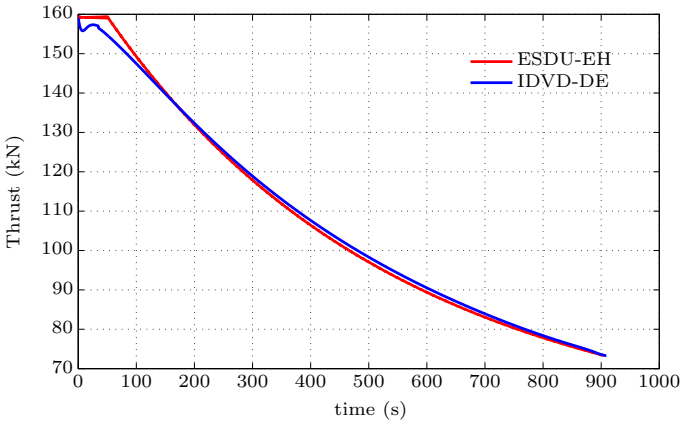

FIGURE 9.9: MTTC case: Comparative trajectory thrust profiles for the IDVD-DE and the ESDU-HE methods

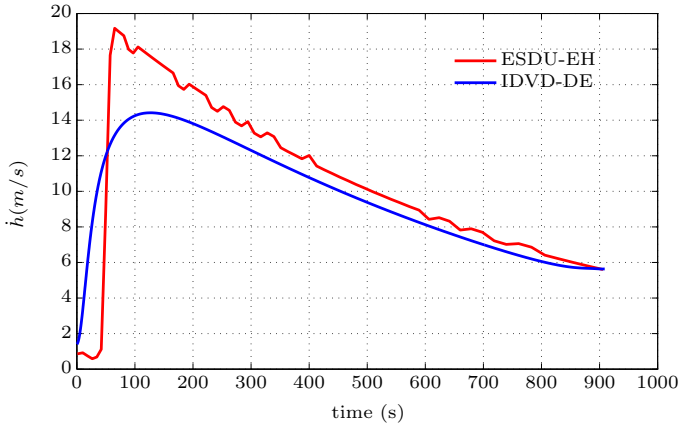

FIGURE 9.10: MTTC case: Comparative trajectory rate of climb profiles for the IDVD-DE and the ESDU-HE methods

9.3 3Di Score Case Study

As a means of measuring progress towards fuel and emissions reduction goals, the 3D inefficiency (3Di) metric was developed by the UK Air Navigation Service Provider (ANSP) NATS. Specifically, the intention of the metric is to provide a measure of

- the flight efficiency of a given flight [170],
- the fuel efficiency of a given flight [39, 82],
- the environmental performance of a given flight [171, 172],
- ANSP performance in delivering a preferred trajectory [82].

In principle the 3Di score is calculated by comparing a flown trajectory to a theoretical fuel/CO2 optimum trajectory. Inefficiencies in the horizontal track and the vertical profile are measured independently and then combined into a weighted expression to determine the 3Di score of a flight. The theoretical optimal is defined as a totally environmentally efficient 4D trajectory that minimises fuel and therefore CO2 [173]. The long term use of the metric is intended to drive fuel burn and CO2 related improvements in trajectories [171].

To determine the coefficients used in the calculation of the 3Di score, the optimal vertical trajectory was defined by NATS as a BADA [128] generated trajectory along the great circle path between departure and arrival airports. The BADA vertical trajectory was modelled using the standard BADA speed schedule as an uninterrupted climb and descent to and from a Requested Flight Level (RFL) and a cruise segment at the RFL. The fuel inefficiency for each trajectory within United Kingdom airspace was then determined by comparing the fuel burn from a BADA generated trajectory ($S_{f_{REF}}$) to the estimated fuel burn from actually flown trajectory ($S_{f_{ACT}}$) for a sample of 174000 flights as

$$I = \frac{S_{f_{ACT}} - S_{f_{REF}}}{S_{f_{REF}}} \tag{9.5}$$

Using regression analysis, the fuel inefficiency I was further simplified to track extension distance σ to represent the horizontal inefficiency while time periods of level flight t_i below the RFL were used to calculate the vertical inefficiency.

The vertical inefficiency ν_i related to periods of level flight away from the BADA trajectory is then calculated as

$$\nu_i = \begin{cases} \frac{t_i(L-l_i)}{T_d L} & l_i \leq L \\ 0 & l_i > L \end{cases} \tag{9.6}$$

where l_i is the flight level during the level flight, T_d is the time duration of the flight and L is the Requested Flight Level (RFL) for the flight. It can be seen that periods above the RFL are regarded as having zero inefficiency. The inefficiency of periods of level flight below the RFL are calculated by multiplying the time duration by the difference between the actual flight level and the requested flight level. This acts to increase the vertical inefficiency value the lower the flown level flight level is from the RFL. To account for differing rates of fuel burn in different phases of flight, the vertical inefficiency is considered by phase of flight where

$$\nu_{CL} = \sum_{CLIMB} \nu_i, \qquad \nu_{CR} = \sum_{CRUISE} \nu_i, \qquad \nu_D = \sum_{DESCENT} \nu_i \tag{9.7}$$

and where the ν_{CL}, ν_{CR}, ν_D are the vertical inefficiency of the climb, cruise and descent phases respectively. The horizontal inefficiency part of the score, σ, is calculated by comparing the actual distance flown D_{ACT} to the minimum Great Circle Distance (GCD) that could have been flown between the same points D_{GCD}

$$\sigma = \frac{D_{ACT} - D_{GCD}}{D_{GCD}} \tag{9.8}$$

The differences between the distances are then considered to be the effect of fuel inefficiency introduced by tactical instructions, procedure and airspace design.

The 3Di inefficiency score ϑ is then determined by combining the horizontal and vertical inefficiencies into an overall inefficiency score

$$\vartheta = a_1\sigma + a_2\nu_{CL} + a_3\nu_{CR} + a_4\nu_D + a_5\nu_{CL}\sigma + a_6 v_{CR}\sigma + a_7\nu_D\sigma \tag{9.9}$$

where $a_1, a_2, a_3, a_4, a_5, a_6, a_7$ are 3Di score regression coefficients.

Metric Evaluation

In 2011 the UK CAA sought stakeholder consultation with regard to the 3Di metric [170]. Eurocontrol, the European organisation for the safety of air navigation responded as follows [174].

"Has NERL (NATS En Route Ltd) endeavoured to develop the best flight efficiency regime?, the answer would have to be in the negative; *there are no attempts to derive what the user considers to be the optimum flight profile beyond what is contained in the flight plan*, which is heavily influenced by vertical restrictions imposed at NERLs request. Moreover, the indicator is measured with reference to a model, which makes the indicator

dependent on the validity of the model. Information on this model is too limited to take a view on its validity.

The horizontal flight efficiency appears straight forward and could certainly be applicable from 2012, being a variation (albeit considerable) of the KPI (Key Performance Indicator) established as part of the SES (Single European Sky) II Performance scheme. The vertical aspects are fundamentally different. *Considering the 2395 standing level agreements in the UK RAD (Route Availability Document), will these be considered as the optimum levels requested by the users or will these simply be removed from the vertical efficiency analysis?* If aircraft are subject to level capping then how will their optimum level be known to compare with the flight profile? The difference in profiles for these flights could be substantial and this will impact arriving traffic."

In summary, Eurocontrol questioned the optimality of the theoretical optimal trajectory and also highlighted the need for any vertical efficiency metric to consider the inefficiency introduced by flight planning restrictions [174].

Aims

In response to the Eurocontrol evaluation of the metric, the remainder of this section will use the Inverse Dynamics trajectory optimisation method to examine a number of aspects of the 3Di score. Principally,

- the suitability of using level segments to define vertical fuel inefficiency,
- the suitability of using a BADA trajectory to define a fuel/CO2/environmentally optimal vertical trajectory,
- the effect of flight planning constraints on a fuel efficient trajectory.

9.3.1 Use of Level Segments to Define Vertical Fuel Inefficiency

To achieve a zero inefficiency 3Di score for the climb phase of flight, NATS recommend that departing aircraft perform a Continuous Climb Departure (CCD), which is then facilitated by the ANSP from an air traffic control perspective [172]. CCDs involve giving the aircraft a direct uninterrupted routing to the top of climb. However the current guidance can encourage continuous gains in height over gains in speed, which may result in aircraft achieving less fuel efficient climbs. To demonstrate, an example simulation was performed. The simulation scenario consisted of an A321 climb to a RFL of 6705 m/22000 ft at a specified distance from take-off. Three trajectories solutions were investigated, a standard BADA speed schedule climb to the RFL followed by a level cruise to the target distance, a constant angle/constant acceleration climb to the target distance and an IDVD-DE generated trajectory to the target distance. The common start and

end speeds for all trajectories were taken from the BADA speed schedule. However, to provide the clearest illustration of the differences between the trajectories, the operational 129 ms^{-1}/250 kts IAS constraint below 3048 m/10000 ft was not enforced.

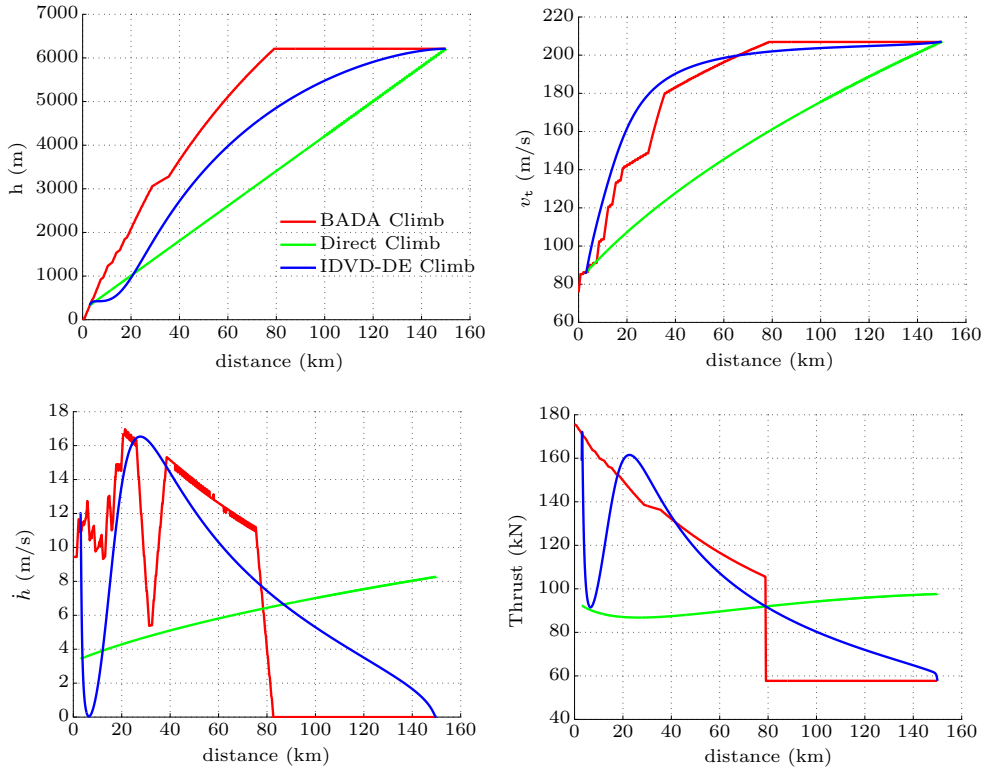

FIGURE 9.11: 3Di case: Climb profile comparisons

From Figure 9.11, examining the BADA generated trajectory solution, it can be seen that the aircraft climbs at approximately 90% maximum available climb thrust using the standard BADA speed schedule directly to the RFL. On reaching the RFL, the aircraft sets thrust equal to drag for the remaining time to the target distance.

For the constant angle climb solution, a constant acceleration, and therefore a linear speed profile climb as proposed in [16] is utilised. Aircraft thrust is determined inversely to deliver the desired dynamics. The speed schedule, although low relative to the other trajectories did not violate the minimum speed constraint.

After an initial climb out common to all solutions, the trajectory generated by the IDVD-DE method consists of a level segment at low flight levels followed by a fast climb to the target conditions. The level segment is used principally to accelerate the aircraft using excess thrust. Although fuel burn is expensive at lower flight levels, the low level acceleration is utilised to achieve a faster, more direct climb to the target conditions. It can be seen from the thrust profile that the thrust utilised is limited by the BADA acceleration constraint and that this requires the thrust to be reduced during acceleration before higher levels are reintroduced when the flight recommences the climb at higher speeds.

Of the 3 trajectories, the IDVD-DE method had both the shortest flight time and the lowest fuel burn 1265 kg/807 sec. The next lowest fuel burn was provided by the BADA trajectory, which had the next shortest flight time, 1376 kg/843 sec. The constant acceleration climb trajectory had the longest flight time and was also the trajectory with the largest fuel burn, 1458 kg/1046 sec. While the trajectories presented above are considered to be operationally achievable they are unlikely to be operationally desirable. They do however illustrate the importance of speed management in delivering a fuel efficient climb trajectory. Therefore the results suggest that fuel efficient climbs are achieved by prioritising energy management and not by exclusively focusing on continuously climbing the aircraft.

Examining the trajectory solutions relative to the 3Di score, the BADA and the constant angle climb trajectories would have 3Di scores of zero inefficiency due to neither trajectory having a level flight segment below the RFL. This is despite there being considerable differences in the fuel burn for both those solutions. The IDVD-DE method generated the most fuel efficient trajectory, but would have been graded by the 3Di score as being the least fuel efficient trajectory, as the trajectory has a level segment of flight at low flight levels.

9.3.2 Optimum Vertical Profile

In the 3Di score calculation, BADA vertical trajectories are assumed to define the theoretical optimum trajectory. To examine the fuel efficiency of a BADA generated trajectory, a simulation scenario was created comparing an IDVD-DE generated trajectory to a BADA generated trajectory. For the scenario, the 3Di demonstration flight was used for reference [175, 39]. The demonstration flight involved an uniterupted A321 flight from London to Edinburgh, unconstrained by typical ATM constraints and with a cruising height of 10363 m/34000 ft [175, 39]. Therefore a scenario was setup involving an A321 flight for a great circle distance equivalent to the Heathrow to Edinburgh distance (544 km). The BADA trajectory used for comparison with the IDVD-DE method was generated by the BADA performance calculation tool, with a specified cruise height set to 10363 m/34000 ft. The input boundary values for the IDVD-DE method were also defined by the BADA solution.

For the IDVD-DE solution, to improve the accuracy of the optimisation, the trajectory was treated as a series of maneuvers, where three piecewise polynomial trajectories ($\mathcal{P}1, \mathcal{P}2, \mathcal{P}3$) were optimised and the end states of one polynomial determined the initial conditions for the next such that

$$\mathcal{T}_{[\tau_o,\tau_f]} = \mathbf{r}(\tau) = \begin{cases} \mathbf{r}_{\mathcal{P}1}(\tau), & \tau \in [\tau_o, \tau_1] \\ \mathbf{r}_{\mathcal{P}2}(\tau), & \tau \in [\tau_1, \tau_2] \\ \mathbf{r}_{\mathcal{P}3}(\tau), & \tau \in [\tau_2, \tau_f] \end{cases}, \quad \forall \tau \in [\tau_o, \tau_f] \tag{9.10}$$

and where the optimisation variables became

$$\Xi = \begin{bmatrix} x_0''', y_0''', h_0''', v_0'', \\ x_1, y_1, h_1, x_1', y_1', h_1', x_1'', y_1'', h_1'', x_1''', y_1''', h_1''', v_1'', \\ x_2, y_2, h_2, x_2', y_2', h_2', x_2'', y_2'', h_2'', x_2''', y_2''', h_2''', v_2'', \\ x_f''', y_f''', h_f''', v_f'', \\ \tau_1, \tau_2, \tau_f \end{bmatrix} \tag{9.11}$$

Comparing the IDVD-DE trajectory to the BADA generated trajectory in Figure 9.12, it can be seen that during the departure climb phase that the IDVD-DE solution has a shallow near level acceleration segment at a height of approximately 1000 m. The shallow segment was used to accelerate the aircraft, with much of the gain in kinetic energy then used to increase the climb rate of the aircraft. Higher levels of thrust are then re-introduced to maintain the high climb rate at the higher speeds. This results in the IDVD-DE trajectory having a faster initial climb out relative to the BADA trajectory, but requires the IDVD-DE trajectory to use maximum climb thrust where the BADA trajectory utilises 90% thrust. Higher in the climb phase, the thrust is reduced and used preferentially for climbing while the speed is allowed to level out. However, the IDVD-DE solution is then required to perform a second acceleration segment so that it does not violate the BADA defined minimum speed constraint, shown in orange in Figure 9.12. The IDVD-DE trajectory reaches a maximum cruise height of 10973 m/36000 ft before commencing a shallower, slower descent than the BADA trajectory until below 3048 m/10000 ft, after which both trajectories assume a 3 degree descent slope to final approach.

At 2176 kg of fuel, the IDVD-DE generated trajectory used almost 10% less fuel than the 2412 kg of fuel used by the BADA trajectory. Therefore, the results suggest that the BADA trajectories used in the calculation of the 3Di score are not fuel or CO2 optimal trajectories. It can be seen that IDVD-DE trajectory cruised at a slower more economical speed than the BADA trajectory, having a longer flight time of 50 minutes relative to the BADA trajectory flight time of 43 minutes. However, airlines normally fly the aircraft using a cost index that is a balance between the fuel burn cost and the operating time cost of the aircraft. It can also be seen that the IDVD-DE trajectory involved an extended use of maximum climb thrust on climb out. Airlines usually prefer to minimise the use of maximum thrust levels due to the wear and tear it causes on the engine. Therefore the results highlight that the most fuel/CO2 efficient trajectory may not be the user preferred trajectory. The results also highlighted that it is likely that there is a trade-off between fuel burn and maintenance costs. It is suggested that a multiobjective trajectory optimisation study, as in [176], could be used to investigate the trade-offs between trajectories optimised for fuel, operating and maintenance costs.

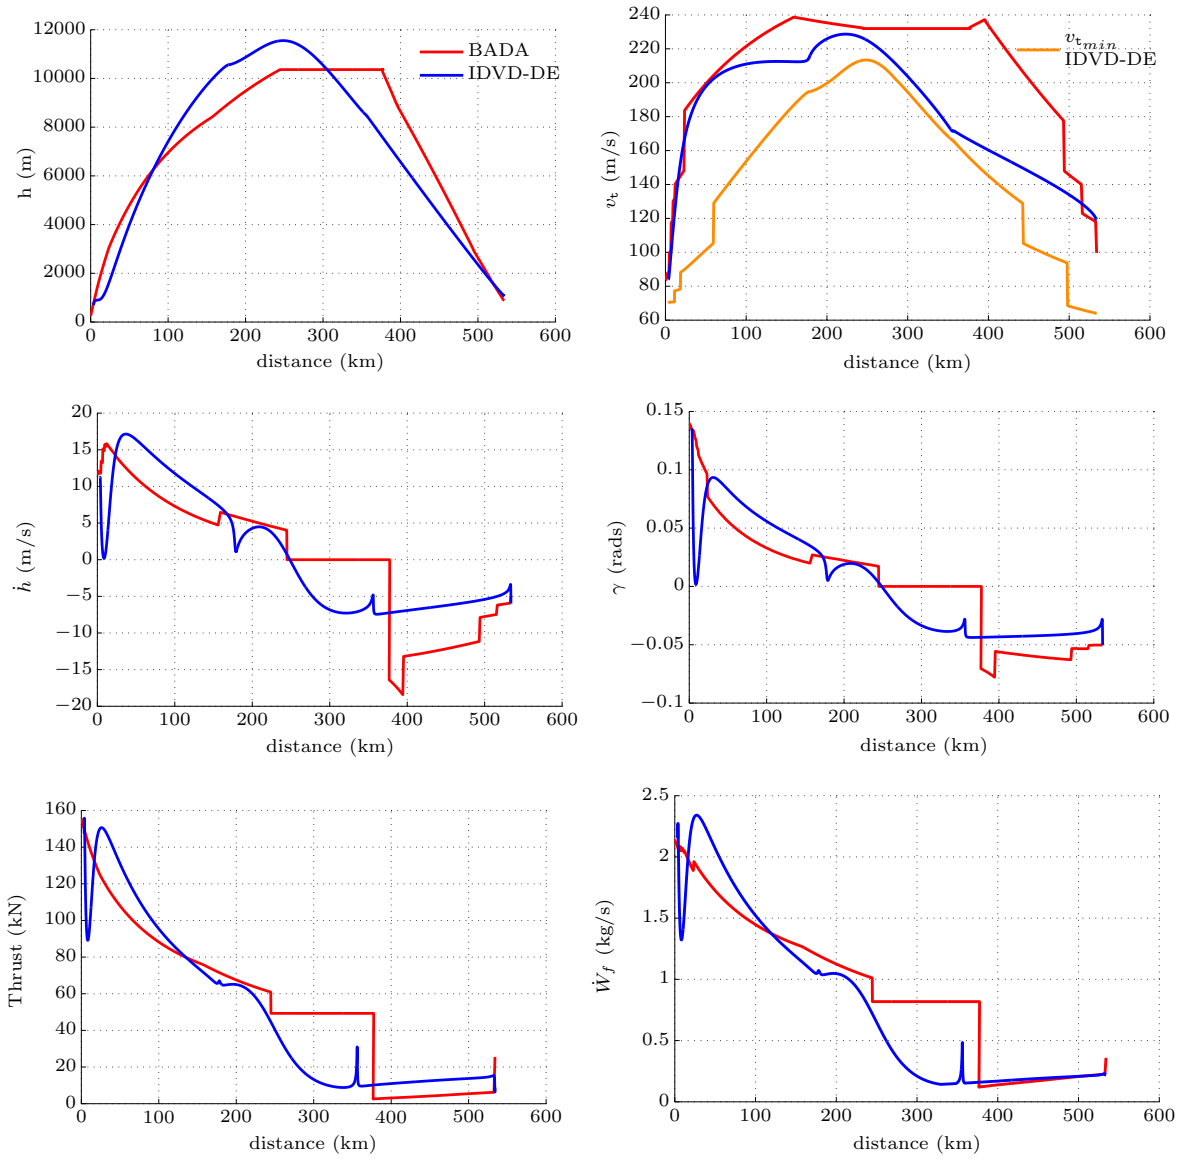

FIGURE 9.12: 3Di case: Fuel efficient trajectory comparisons

9.3.3 Flight Planning Constraints

To demonstrate the impact of flight planning constraints on a preferred trajectory, the IDVD-DE method was used to generate procedurally constrained and unconstrained fuel efficient trajectories for a representative flight from London Gatwick (EGKK) to Paris Charles de Gaule (LFPG). The fuel inefficiency introduced by the constraints was then analysed with consideration of the 3Di score. The flight plan generated from [177], required the constrained trajectory to follow the Hardy SID from Gatwick, The UM605 airway from Hardy to the DPE and to descend on the DPE 4W STAR towards Paris. The trajectory waypoint constraints for the constrained trajectory are shown in Table 9.2.

Waypoint	Lat	Long	Diameter (m)	H (m)	H tolerance	V (m/s)	V tolerance
MIDD10	51.128	-0.358	5000	1220	460	-	-
OCKD13	51.088	-0.428	5000	1297	383	-	-
OCKD18	51.005	-0.421	5000	1525	155	-	-
OCKD23	50.923	-0.415	5000	1846	166	-	-
OCKD28	50.84	-0.409	5000	-	-	-	-
BOGNA	50.702	-0.252	10000	-	-	-	-
HARDY	50.473	0.485	10000	-	-	-	-
DPE	49.923	1.171	5000	6401	153	-	-
SOKMU	49.334	1.419	5000	3962	152	-	-
MEURE	49.301	1.851	10000	2743	153	129	5

Table 9.2: 3Di case: Flight planning constraints

The waypoint constraints were taken from the relevant AIP and SID/STAR charts. From the DPE STAR chart, there was an additional below 144 ms^{-1}/280 kts IAS constraint at DPE, however it was found that this constraint conflicted with the minimum BADA defined flight speed constraint at the DPE constraint target height of 6401 m/22000 ft. Therefore, the constraint was removed for this study. A corridor constraint [132] was considered to represent the UM605 airway, however it was found that the waypoint constrained trajectory pre-satisfied the airway constraint and therefore the corridor constraint is not shown here. In addition to the waypoint constraints the constrained trajectory was also subject to the 129 ms^{-1}/250 kts IAS below 3048 m/10000 ft rule. Both trajectories were required to be established in line with the runway on final approach by 6 DME and to fly three degree descending approaches on the ILS below 1036 m/3400 ft.

Figure 9.13 shows the constrained and unconstrained trajectory solutions to the London-Paris scenario where the height and speed profiles are shown relative to the circular waypoint crossing and the cylindrical height and speed constraints. Comparing the solutions in Figure 9.13, it can be seen clearly that, close to Gatwick, there are large differences between the speed profiles of the two trajectories. The unconstrained trajectory is not constrained by the less than 129 ms^{-1}/250 kts below 3048 m/10000 ft restriction, nor the HARDY SID constraints, therefore, thrust is initially applied to preferentially increase

aircraft speed at lower flight levels leading to a faster climb to higher flight levels. When the aircraft reaches a height of approximately 6000 m, it begins to trade-off some of its speed for continued gains in height. The constrained trajectory solution climbs from EGKK following the HARDY SID restrictions. At 3048 m/10000 ft the aircraft clears the 129 ms^{-1}/250 kts IAS speed restriction and begins an acceleration to 200 ms^{-1}/390 kts while also continuing to climb. Once the aircraft in the constrained solution nears its cruise height, it also begins to trade-off some of its speed for continued gains in height.

Examining the descent from cruise, it can be seen that the unconstrained trajectory, unconstrained by the DPE STAR, begins its descent to Paris earlier than the constrained trajectory. It assumes a shorter path to LFPG, reducing speed early, helping it to minimise fuel burn. For the constrained solution, the aircraft must stay higher for longer due to the STAR constraints. However, once the aircraft does commence its descent, it quickly assumes similar descent rates, speed and fuel consumption profiles as those of the unconstrained trajectory.

In terms of fuel burn, the ATM constraints considered introduced a 321kg/17% fuel inefficiency relative to the unconstrained trajectory solution. Analysing the results relative to the 3Di score, it is reasonably expected that the path extension factor σ would account for horizontal path inefficiencies introduced by the constraints. However, for the vertical efficiency, it is clear that the ATM constraints alter the most efficient cruising height and therefore the cruising height requested by an airline in the flight plan. Therefore, there is an RFL related fuel inefficiency included in the flight plan submitted to the ANSP. As inefficiencies in the RFL are not considered in the calculation of the 3Di score, the constraints have introduced a vertical inefficiency into the trajectory that is not quantified by the metric. The results also showed that including the waypoint constraints in the flight planning caused them to be navigated with no periods of level flight. So, similar to the RFL inefficiency, the vertical fuel inefficiency related to those waypoints is again unquantified by the 3Di score.

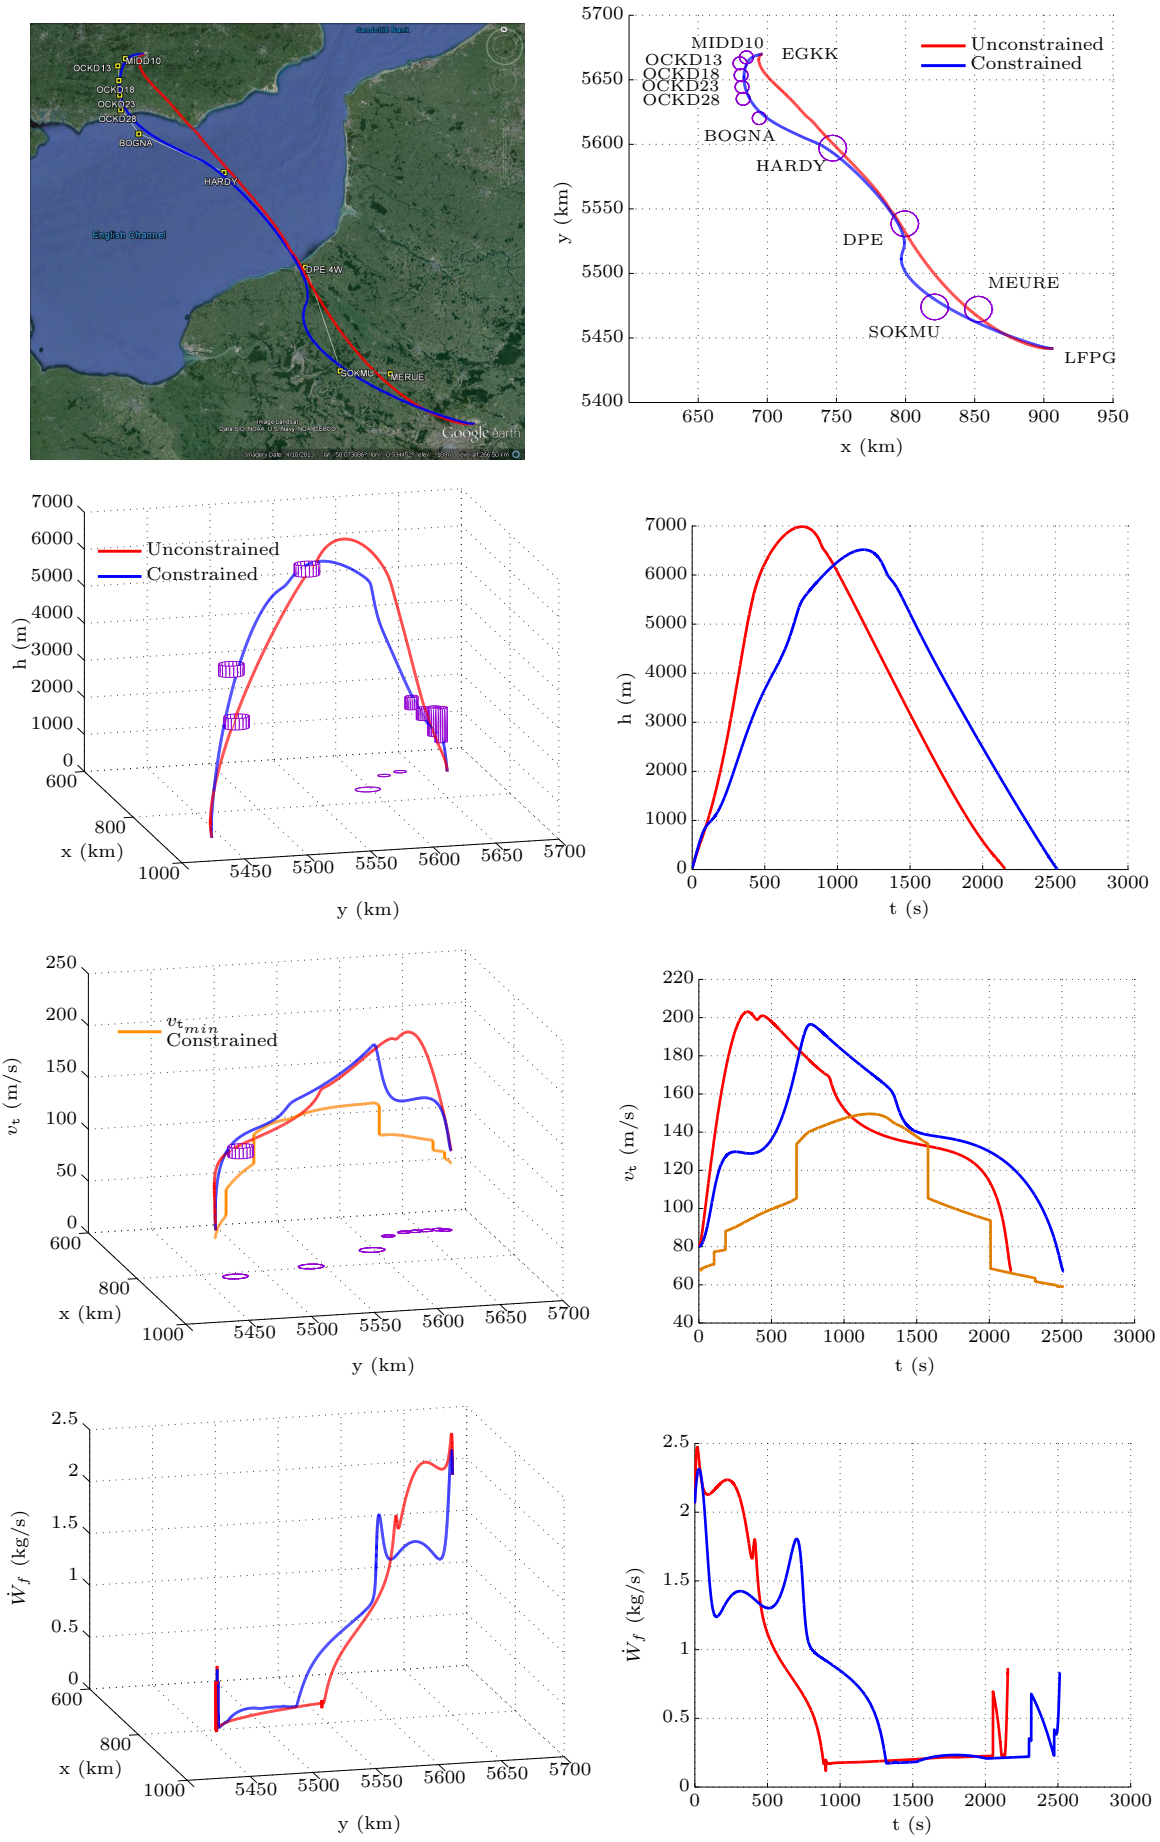

FIGURE 9.13: 3Di case: London-Paris constrained and unconstrained trajectory solutions

9.3.4 3Di Case Study Conclusions

In response to a Eurocontrol review of the metric, the 3Di score has been analysed using a trajectory optimisation method based on optimal control.

The results suggest that further development is required for the metric to be considered as a flight or fuel efficiency metric. The results show that BADA trajectories used in the 3Di score to define the optimum fuel efficient operation of the aircraft are not optimal fuel efficient trajectories. The results also highlight the importance of considering flight planning restrictions when calculating the fuel inefficiency of a trajectory.

The fuel efficiency of a flight trajectory is a collaboration between the flight crew and ATC. Separating ATM introduced fuel inefficiencies from operator introduced inefficiencies is a subject that still requires further research. However, the ANSP has significant influence over trajectory height, speed and climb rates through tactical instructions and the design of the ATM system, including climb and descent procedures and standing agreements, and also through the guidance it issues on efficient climb and descent operations. The results have shown that currently the 3Di score is insensitive to a wide number of operational changes that can be applied to the vertical trajectory to reduce the fuel burn of a flight. Therefore, if as proposed in [171], the 3Di score is to drive long term operational improvements in fuel related trajectory operations, the metric should be sensitive to operational changes in the vertical trajectory that significantly impact fuel consumption, and not solely be sensitive to periods of level flight away from a BADA trajectory.

ICAO defines 3 prominent types of aviation related environmental impacts as climate changing emissions, air quality emissions and noise. It is well established that there are a wide number of trade-offs between the objectives. Particularly for noise objectives, which are very sensitive to local conditions. In CO2, the 3Di score only attempts to look at the inefficiencies related to one form of environmental impact. Also results here have shown that the BADA trajectory used to benchmark the inefficiencies of flown trajectories is not a CO2 optimal trajectory. Defining a environmentally efficient trajectory is likely to involve a trade-off analysis between the competing environmental objectives [176].

The results in this analysis offer a preliminary analysis of the 3Di score using a trajectory optimisation method. The results must still be confirmed by further simulation results, potentially using more accurate trajectory optimisation methods. Further consideration must also be given to the impact of wind, aircraft weight, flight planning restrictions, environmental trade-offs and airline operating costs when defining reference trajectories intended to benchmark inefficiencies in flown trajectories

9.4 Summary: Trajectory Based Flight Planning

This chapter discusses a shift in the ATM concept of operations towards a system based on trajectory flight planning and operation. It is shown here how a trajectory based analysis can be used to study the fuel efficiency of a flight. It is then shown how, by limiting the energy management of the flight, ATM constraints can introduce fuel inefficiencies to planned trajectories. In the cases considered, it was found that ATM related fuel inefficiencies were introduced by

- track-extension,
- constrained flight speed schedule management,
- constrained RFL

Surprisingly, in the work, ATM related fuel inefficiencies were not typified by level flight segments. However, this is not to say that fuel inefficiencies are not introduced by level flight segments. More so that by analysing the effects of ATM constraints on flight energy management, a more detailed understanding of the origin of a fuel inefficiency may be achievable. More long term, it is expected that managing these inefficiencies on a trajectory by trajectory basis will be key to achieving the SES fuel efficiency goals.

Although fuel consumption is one of the most significant costs to an operator, it quickly became clear when analysing the results in this chapter that there comes a point where the most CO2 efficient trajectory diverges from the operator preferred trajectory. However, determining where that point is was beyond the scope of the work in this thesis, which focused on trade-offs in environmental objectives.

Determining the trade-offs between fuel/CO2 efficient and operator preferred trajectories can be approached in the same way as the multi-environmental-objective trade-off case studies in Chapter 8. However, in this instance, methods and metrics would not be solely environmental, but would include models and objectives specific to airline operating costs. These are likely to include, among others, time costs of crew and aircraft, CO2 permit cost, airport and ATM service costs and engine maintenance costs.

Chapter 10

Conclusions and Recommendations for Future Work

10.1 Procedure Optimisation

Aircraft climbing and descending to and from airports are highly constrained by procedural ATM and operating constraints. As defined by the thesis goals set out in Section 2.6.3, a major aim of the work in this thesis, is the optimisation of these procedure constraints in a manner that supports the environmental mitigation of flight operations. Procedure optimisation, it is proposed, is best enabled through informed decision making on the trade-offs in environmental impacts that occur from the different ways of operating the aircraft within the TMA.

It is understood that within the TMA, neither the adoption of the most community noise nor fuel optimal operation of a flight defines an operational procedure that best balances competing environmental impacts. Therefore, to fully explore the trade-offs in environmental impacts, the procedure optimisation problem is treated as a multi-objective optimal control problem that is solved for many objective Pareto fronts using a direct optimal control method combined with a global heuristic NLP solver.

The environmental measures and calculation methods used in this work are defined in Chapters 3 and 4. It is discussed in these chapters how policy choices affect the prioritisation of different measures and also how all the measures can be calculated as functions of the states and controls of the aircraft trajectory. The importance of assessing noise impact relative to populations close to airports is emphasised in Chapter 4, and drives the adoption of the noise Annoyance Score metric for a number of case studies in the thesis.

In Chapter 5, methods for converting the optimal control problem into an NLP problem are reviewed. The Inverse Dynamics in the Virtual Domain method is detailed in this

chapter and it is discussed that the method is adopted due to its low parameter space and balance between computational efficiency and trajectory accuracy.

In Chapter 6, the NLP methods considered for use with the IDVD trajectory method are discussed. The desire for global solutions and Pareto fronts led to the choice of a global heuristic method. The specific choice of the Differential Evolution heuristic was due to encouraging work by Drury in [79] who had previously combined the IDVD and DE methods. Chapter 6 also defines the algorithmic structure of the PDE and DEMO algorithms, which are multi-objective variants of the DE method, and combine DE with the nondominated sorting and crowding distance methods defined by Deb [153] for the NSGAII algorithm. It is shown in this chapter and in Appendix A how the basic PDE and DEMO algorithms perform well for two objective problems but perform poorly when applied to optimisation problems with higher numbers of objectives (sometimes referred to as many objective problems).

Recognising this as a significant limitation, Kukkonen and Deb proposed the kNN crowding distance measure and pruning method intended to improve the performance of PDE and DEMO like algorithms for many objective optimisation problems. As a number of the environmental procedure optimisation studies in this thesis were for more than two objectives, the crowding measure and pruning method were adopted for this work.

Sections 6.4.3 and 6.4.4 define the integration of the kNN pruning method into the PDE and DEMO algorithms. Although the algorithms used have largely been defined in the previous literature, they are scattered over a wide number of documents. In Section 6.4 the algorithms are presented in a consolidated and harmonised form that makes clear the function, reliance's and interactions of each algorithm.

In Chapter 7, a multi-objective Pareto analysis technique is defined that colour clusters Pareto front solutions and related trajectory profiles according to the balance of objectives on the Pareto front. The aim of the technique is to use the Pareto front objective values to identify distinct groupings of trajectory behaviour. In Chapter 7, the convergence and diversity properties of the multi-objective IDVD-DE approach are also investigated for a realistic but not real-world environmental objective test problem. The algorithm performed well, identifying significant trade-offs in trajectory objectives and consistently converging on Pareto fronts close to the defined pseudo global optimal.

In Chapter 8, the IDVD-DE approach was applied to two real world procedure optimisation case studies. The Luton departure procedure (Section 8.3) and NEMA arrival procedure (Section 8.2) case studies were both studies with existing environmentally optimised solutions that provided baseline and reference information for comparison to the IDVD-DE generated trajectory solutions. The baseline procedures in each study were reached through more traditional approaches, using expert opinion and iterative analysis. In the Luton departure procedure definition case study, The IDVD-DE method was used to identify a range of ground paths and related vertical trajectories that, as a group,

offered an improvement over both baselines procedures in three of the four objectives and offered an improvement in the final objective over the existing procedure design.

For the NEMA case study (Section 8.2), environmental performance values for the the baseline procedure were not available for direct comparison to the performance values of IDVD-DE generated solutions. However the IDVD-DE method did identify a number of environmentally advantageous approach procedure characteristics. While the baseline NEMA study proposed the use of the Low Power Low Drag (LPLD) descent technique to reduce fuel and noise impacts of approach operations, the technique was not found to be a defining feature of the IDVD-DE generated Pareto solutions. Instead the IDVD-DE generated solutions suggested an alternative descent strategy where increased thrust levels were introduced higher in the descent in order to minimise thrust usage lower in the descent.

The goal of the multi-objective IDVD-DE approach used, is not to have the method dictate a single optimal solution, but to have it generate a range of potential Pareto solutions. The Pareto information is then intended to inform flight and airspace planners, allowing them to draw conclusions on how changes in operational procedures affect key trade-offs in environmental objectives. This, it is intended, will allow planners to environmentally optimise procedure designs in a data driven manner not currently achievable with existing approaches to the problem. The combination of control theory, a global heuristic method and Pareto based environmental analysis of flight climb and descent operations was not observed in the literature prior to work in this thesis.

10.1.1 Implementation and Time Investment

As computational performance was seen numerous times in the literature to be a significant limitation to achieving multi-objective environmental procedure optimisation (see Section 2.6.3), there was an emphasis on computational efficiency in the choice of the IDVD-DE method. In particular, the low parameter space offered by the IDVD-DE method allowed the quick evolution of potential trajectory solutions. However, the version implemented in this thesis was a MATLAB prototype not fully optimised for run-time performance.

For studies in this thesis with noise objectives, an external noise model was used (see Section 4.3). The run time performance of INM then largely determined the run time performance of the optimisations. In turn, the choice of grid size, density of grid, population density and number of trajectory discretisation nodes all had an impact on the run-time performance of the noise model. As INM was an external model, inefficient reading and writing of input and output files was required.

INM, however is based on an open standard Noise Power Distance (NPD) method. Although grid based, the NPD approach is effectively a corrected look-up table approach. If

implemented natively with the IDVD-DE method, it should be well suited to faster execution times through code optimisation techniques such as vectorisation. If implemented natively, then IDVD-DE optimisations could also be parallelised on a single computer or cluster of computers, significantly improving run time performance.

The *k*NN pruning algorithm implemented (see Section 6.4.3) was simplified from the initial implementation to take advantage of Matlab logical indexing, which was very fast for populations of less than 1000, to which it was applied to. However, linked lists were used in the initial implementation of the algorithm to manage solution indexing and it is expected that their use leads to a faster implementation of the algorithm. Further, it is expected that the run time performance of the methodology could be further improved through general code optimisation in a programming language such as C/C++.

As it stands, the multi-objective optimisation runs considered in this thesis (see Chapters 7 and 8) took between 8-16 hours, with objective numbers from 2 to 4 and solution numbers from 70-500. The maximum number of generations chosen ranged between 10,000 and 20,000 with trajectory discretisation between 1000 and 6000 nodes. Optimisation scenarios were generally run over-night on Intel Core i7 (3.0GHz, 8 core) 16GB of RAM machines.

Surprisingly, the most time consuming element of each study was not the simulation run time, but the post-optimisation analysis. Each trajectory solution is itself composed of state and control profiles, including height, speed, thrust, aerodynamic and flight path routing histories. Therefore, there is a significant amount of information to parse when determining which trajectory factors are driving the significant changes in in which environmental objectives. This process, even with the clustering algorithm acting to group trajectories with similar behaviours, is typically very time consuming. This is especially true when analysing non-intuitive trajectory solutions.

Design of experiment and metrics chosen also have a significant impact on the ease with which conclusions can be drawn from the optimisation results. This was particularly true for the NEMA case study (Section 8.2), where the choice of the noise under the centreline metric interacted with the other metrics chosen in a manner that added sparsity to the Pareto front, making the results difficult to analyse. It is expected though, that the repeated use of the techniques described here will lead to even the non-intuitive solutions becoming more conventional and widely accepted. Over time this should allow for the parsing of Pareto front and related trajectory information to become quicker and more intuitive.

As with any large dataset however, it is beneficial to have techniques that can highlight potentially interesting or useful information within the set. Therefore, for future work, it would be very beneficial if further techniques could be developed and used alongside the Pareto front information to recommend potentially interesting or well-balanced solutions from the front.

10.1.2 Procedure Optimisation: Further Investigation

A number of interesting topics related to procedure optimisation, and the application of the IDVD-DE method to procedure optimisation, were not fully explored in this work. These topics would be suitable for further investigation as part of future work.

In the NEMA and Luton procedure optimisation case studies (see Chapter 8), the end conditions for climb and the start conditions for descent were specified by the baseline solutions. However, instead of matching the optimisation boundary values to those of the baselines, potentially better environmental trade-off solutions than those identified in the case studies could be available by allowing greater flexibility of either the initial or terminating boundary conditions.

Disappointingly, there was no real world studies on which to base a local air quality procedure optimisation scenario. While LAQ emissions are often monitored in procedure design studies, fuel consumption rather than LAQ emissions is typically used as the procedure design goal. This is done with the assumption that minimising fuel minimises all emissions. Results from the Luton departure procedure optimisation case study in Section 8.3.3 show that this is a very reasonable assumption. However the results do show that there are trade-offs between LAQ emissions and fuel burn and also between individual LAQ emissions. These trade-offs, from a climb and descent operational perspective, have really not been explored in the literature and would be suitable for further investigation.

Some previous work was conducted looking at the differences between NOx emissions calculated using the lower fidelity time in mode calculation relative to the higher fidelity advanced calculation method [98, 99] (see also Section 3.3). The latter advanced method is used in this work (see Section 3.4). The work was not specifically aimed at assessing aircraft operational procedures, but showed that the 'time in mode' calculation significantly underestimated NOx emissions relative to the advanced method [98, 99]. This was related, at least in part, to the greater sensitivity of the advanced method to how aircraft were being operated. Therefore, it would be interesting, for future work, to use multi-objective optimisation to examine how the differences in two metrics affect general guidance on the most NOx efficient climb and descent operations.

10.2 Trajectory Based Operations

In addition to the thesis goals set out in Section 2.6.3 for defining a multi-objective procedure optimisation approach, thesis goals set out in Section 2.6.4 include the aim to further apply the IDVD-DE method to the planning of user-preferred flight trajectories.

It is discussed in Section 2.5 how Trajectory Based Operations is a defining element of the SES target concept. For TBO, the trajectory is the fundamental unit of the air traffic system. It defines the intent and the cost in terms of energy and the environmental impact of a flight movement. The trajectory is shared by operators with ANSPs and is collaboratively updated and revised to maintain flight safety and a balance between flight and network system efficiency.

For these reasons, in trajectory based operations, the planning of the 4D flight trajectory takes on a heightened importance for the aircraft operator. In Chapter 9 the IDVD-DE method is applied to the trajectory based planning of CO2 efficient flight trajectories. To improve the accuracy of the IDVD-DE approach for the 3Di case study considered in Section 9.3, each of the climb, cruise and descent phases of flight were represented by piecewise polynomial trajectories. The ideal planned trajectories of the 3Di method are used in the case study to define a current best practice trajectory flight planning approach.

In all cases, the IDVD-DE method defined flight trajectories between 5-10% more fuel/CO2 efficient than the 3Di defined optimal solutions. The IDVD-DE method was further applied to a real world trajectory flight planning problem and was used to quantify the flight efficiency impact of SID/STAR and airway constraints.

The most interesting part of the work in the 3Di case study however, is not the performance of the IDVD-DE method in itself, but how the IDVD-DE method, and by extension, trajectory optimisation methods in general, can be applied to analyse commonly made assumptions regarding the most fuel/CO2 optimal operation of commercial flights.

In the Continuous Climb Departure simulation scenario in Section 9.3.1, it is shown how conventional guidance on continuously climbing the aircraft can result in the inefficient management of the ascent speed schedule when planning fuel optimal climbs. Unexpectedly, level segments were often a feature of fuel efficient climbs, typically used to accelerate the aircraft at lower flight levels. Through the application of the IDVD-DE method, it could be seen that while low level acceleration was expensive in terms of rate of fuel consumption, that the increased energy at lower flight levels could be used to better expedite the climb, in general reducing climb time and minimising overall fuel to climb.

In Section 9.3.2, by comparing the more fuel efficient IDVD-DE generated trajectory to the 3Di reference trajectory, it was discussed how the minimum CO2 cruising speed could be considerably lower than the user-preferred cruising speed.

Often in flight operations planning, it is assumed that climb, cruise and descent phases of flight can be planned independently. The IDVD-DE results for the London-Edinburgh (Section 9.3.2) and London-Paris (Section 9.3.3) simulation scenarios however show that, when planning fuel efficient short duration flights, there is a coupling between the climb,

cruise and descent segments. Similarly, it is often assumed that a flight will seek to spend as much as possible at the aircraft's most efficient cruising level. The same scenarios show that, again for short durations flights, IDVD-DE generated trajectories generally sought to descend from cruise early, performing long, shallow flight path angle descents that took advantage of the best lift to drag glide ratio. While this observation is in keeping with well understood engineering principles, it introduces a trade-off with initiatives aimed at increasing the steepness of approach procedures for noise mitigation purposes.

10.2.1 Flight Planning: Further Investigation

A recurring observation made throughout the results in this thesis, is the importance of the flight speed schedule, and how changes to the speed schedule can drive significant changes in noise and emissions performance measures. While it is common for operators to manage the efficiency of the the flight through the use of the flight speed schedule, it is rarely considered how ATM constraints and recommended procedures affect a flights planned speed schedule. This, it is thought, is a subject worthy of further investigation.

Through these results then, the IDVD-DE method is shown to provide a number of interesting insights regarding CO2 efficient flight operations. However, the results are not definitive and must be further confirmed and potentially improved upon using higher fidelity trajectory optimisation methods.

The CO2 efficient trajectory planning problem was always intended to be an operator focused, single objective case study. However, the results for the 3Di case study in Section 9.3 showed that the fuel planning of a commercial aircraft flight trajectory is in itself a multi-objective problem, where trade-offs in fuel consumption need to be weighed against fuel, time and maintenance costs.

It is therefore also proposed that, for future work, trajectory optimisation methods and Pareto front analysis be applied to explore the cost trade-offs for aircraft operators along with further trade-offs in environmental impacts relative to operator costs.

References

[1] Airbus, "Flying on demand: Global market forecast 2014-2033," 2014.

[2] SESAR Joint Undertaking, "European ATM master plan, edition 2," 2012.

[3] International Civil Aviation Organization, "Environmental report 2013," 2013.

[4] S. K. Ribeiro, S. Kobayashi, M. Beuthe, J. Gasca, D. Greene, D. S. Lee, Y. Muromachi, P. J. Newton, S. Plotkin, D. Sperling *et al.*, "Transportation and its infrastructure," *Institute of Transportation Studies*, 2007.

[5] European Federation for Transport and Environment, "Grounded: How ICAO failed to tackle aviation and climate change and what should happen now," 2010.

[6] M. Grote, I. Williams, and J. Preston, "Direct carbon dioxide emissions from civil aircraft," *Atmospheric Environment*, vol. 95, pp. 214–224, 2014.

[7] International Civil Aviation Organization, "Environmental report 2010," 2010.

[8] Council of the European Union, "Directive 2008/101/EC amending Directive 2003/87/EC so as to include aviation activities in the scheme for greenhouse gas emission allowance trading within the Community [2009] OJ L8/3."

[9] ——, "Regulation (EU) No 421/2014 amending Directive 2003/87/EC establishing a scheme for greenhouse gas emission allowance trading within the Community, in view of the implementation by 2020 of an international agreement applying a single global market-based measure to international aviation emissions [2014] OJ L129/1."

[10] European Commission, "The EU Emissions Trading System EU ETS factsheet," http://ec.europa.eu/clima/publications/docs/factsheet_ets_en.pdf/, 10 2014.

[11] International Civil Aviation Organization, "ICAO Assembly Resolution A38-18: Consolidated statement of continuing ICAO policies and practices related to environmental protection Climate change," 4 October 2013.

[12] ——, "Doc. 10018, Report of the assessment of market-based measures," 2013.

[13] Advisory Council Aeronautics Research in Europe, "Strategic Research Agenda. Volume 1," September 2012.

[14] ——, "Strategic Research Agenda. Volume 2," September 2012.

[15] G. Horton, "Future aircraft fuel efficiencies - final report," UK Department for Transport, 2010.

[16] SESAR Joint Undertaking, "SESAR and the environment," 2010.

[17] Eurocontrol Performance Review Body, "Advice to the Commission in the setting of Union-wide performance targets for RP2," September 2013.

[18] Eurocontrol Performance Review Commission , "Performance Review Report An assessment of air traffc management in Europe during the calendar year 2013," May 2014.

[19] Eurocontrol, "European Route Network Improvement Plan (ERNIP) Part 1 European airspace design methodology guidelines," November 2014.

[20] Council of the European Union, "Regulation (EU) No 598/2014 on the establishment of rules and procedures with regard to the introduction of noise-related operating restrictions at Union airports within a Balanced Approach and repealing Directive 2002/30/EC Community [2014] OJ L173/65."

[21] International Civil Aviation Organization, "Doc 9829 AN/451, Guidance on the balanced approach to aircraft noise management," 2008.

[22] UK Civil Aviation Authority, "CAP725, CAA Guidance on the application of the airspace change process," 2007.

[23] Committee on Aviation Environmental Protection, "Review of Continuous Descent Approach (CDA) implementation and associated benefits," *Working Paper CAEP/7-WP/26*, 2006.

[24] MIME Consortium, "EEC MIME WP6.2, Final report of the MIME project," 2011.

[25] International Civil Aviation Organisation, "Doc 4444, Procedures for Air Navigation Services: Rules of the air and air traffic services," 1996.

[26] ——, "Annex 11 to the Convention on International Civil Aviation, Volume I: Air traffic services," 2001.

[27] A. Cook, *European Air Traffic Management: Principles, Practice, and Research*. Ashgate Publishing, Ltd., 2007.

[28] NATS Enroute Limited, "NERL Strategic plan for safety," April 2011.

[29] A. Majumdar, W. Y. Ochieng, J. Bentham, and M. Richards, "En-route sector capacity estimation methodologies: An international survey," *Journal of Air Transport Management*, vol. 11, no. 6, pp. 375–387, 2005.

[30] Eurocontrol Performance Review Commission, "Vertical flight efficiency," March 2008.

[31] Eurocontrol, "Sixty thousand nautical miles a day wasted," http://www.eurocontrol.int/sites/default/files/content/documents/official-documents/skyway/articles/2014-summer-skyway-focus-sixty-thousand-nautical-miles-a-day.pdf, Summer 2014.

[32] ——, "As the crow flies free route airspace Maastricht," HI, March 2011.

[33] International Civil Aviation Organization, "ICAO Circular 317 AT/136, Effects of PANS-OPS noise abatement departure procedures on noise and gaseous emissions," 2008.

[34] ——, "Doc. 9750-AN/963, Global air navigation plan," 2013.

[35] E. Lassooij, "Performance-based navigation seen as key to global harmonization," *ICAO Journal*, vol. 61, no. 3, 2006.

[36] J. Mcgraw, J. Williams, and H. Shahidi, "Implementation of performance-based navigation making notable progress," *ICAO Journal*, vol. 61, no. 3, 2006.

[37] International Civil Aviation Organization, "Doc. 8168, Aircraft Operations. Volume I: Flight Procedures," 2007.

[38] ——, "Review of noise abatement procedure research & development and implementation results discussion of survey results," 2007.

[39] C. Nutt, "OA 1161 Version 1.3, NATS Fuel efficiency metric," 2012.

[40] International Civil Aviation Organization, "Doc 9993 AN/495, Continuous Climb Operations (CCO) manual," 2015.

[41] UK Civil Aviation Authority, "Noise from arriving aircraft: An industry code of practice," 2001.

[42] T. Symmans, "Point merge a more efficient way of sequencing arrivals," https://www.eurocontrol.int/sites/default/files/field_tabs/content/documents/events/2011-cda-point-merge.pdf, 2011.

[43] Eurocontrol, "Point merge integration of arrival flows enabling extensive RNAV application and continuous descent," July 2010.

[44] SESAR Joint Undertaking, "Doc id 05.07.04-D05 & D06, Operational Service and Environment Definition (OSED) for point merge in complex TMA," 2004.

[45] Eurocontrol, "Flexible use of airspace," http://www.eurocontrol.int/services/point-merge-concept, retreived October 2015.

[46] International Civil Aviation Organization, "Doc 9976, Flight planning and fuel management manual," 2015.

[47] www.rocketroute.com, "Eurocontrol route availability document," http://www.rocketroute.com/blog/eurcontrol-route-availability-document, retreived March 2015.

[48] UK Civil Aviation Authority, "CAP694, The UK flight planning guide," January 2007.

[49] Eurocontrol, "CFMU/URB/ATFCM-STRAT-01-00, Air traffic flow & capacity management strategy," April 2004.

[50] www.skybrary.aero, "Sector over-deliveries and overloads," http://www.skybrary.aero/index.php/Sector_Over-deliveries_and_Overloads, retreived March 2015.

[51] B. Roberson, "Fuel conservation strategies: cost index explained," *Boeing Aero Magazine*, 2010.

[52] International Civil Aviation Organization, "ICAO Circular 303 AN/176, Operational opportunities to minimize fuel use and reduce emissions," 2003.

[53] www.skybrary.aero, "Functional Airspace Block (FAB)," howpublished = http://www.skybrary.aero/index.php/Functional_Airspace_Block_(FAB),, retreived March 2015.

[54] SESAR Joint Undertaking, "SESAR concept of operations step 1," May 2012.

[55] M. Castelletti, "Europe targets airspace reform," *journal of the civil air navigation services organisation*, no. 22, 2023.

[56] Eurocontrol, "Functional airspace block (FABs) - defragmenting european airspace," https://www.eurocontrol.int/functional-airspace-block-fabs-defragmenting-european-airspace, retreived March 2015.

[57] www.skybrary.aero, "Flexible use of airspace," http://www.skybrary.aero/index.php/Flexible_Use_of_Airspace, retreived March 2015.

[58] Eurocontrol, "FASTI making change in en-route air traffic control," June 2007.

[59] SESAR Joint Undertaking, "Milestone deliverable D3: The ATM target concept," 2017.

[60] Environmentally Responsible Air Transport, "D1-2 Version 1.0, Inventory of potential operational measures, effects and required enabling technologies, September 2008," February 2011.

[61] ——, "Final report, 13 April 2011," February 2011.

[62] UK Civil Aviation Authority, "Airspace for tomorrow 2," 2011.

[63] SESAR Joint Undertaking, "D65-011, SESAR concept of operations," 2012.

[64] Thales ATM Inc., "Analysis of SESAR 4D trajectory activities," November 2011.

[65] L. Prez Sanz, "Precision Trajectory Clearance (PTC): A new separation mode,Towards higher levels of automation in Air Traffic Management lecture slides," 2011.

[66] Sustainable Aviation, "Sustainable aviation noise road-map."

[67] Sourdine Consortium, "Report D5, Study of Optimisation procedURes for Decreasing the Impact of NoisE (SOURDINE)," 2001.

[68] H. G. Visser and R. A. Wiljnen, "Optimization of noise abatement departure trajectories," *Journal of Aircraft*, vol. 38(4), pp. 620–627, 2001.

[69] ——, "Optimization of noise abatement arrival trajectories," *The Aeronautical Journal*, vol. 107(1076), pp. 607–615, 2003.

[70] C. Hargraves and S. Paris, "Direct trajectory optimization using nonlinear programming and collocation," in *Astrodynamics 1985*, vol. 1, 1986, pp. 3–12.

[71] S. J. Hebly and H. G. Visser, "Advanced noise abatement departure procedures: custom optimized departure profiles," in *AIAA Guidance, Navigation and Control Conference, AIAA 2011-6598*, Portland OR, August 2011.

[72] M. Houacine and S. Khardi, "Gauss pseudospectral method for less noise and fuel consumption of aircraft operations," *Journal of Aircraft*, vol. 47, no. 6, pp. 2152–2158, 2010.

[73] X. Prats, V. Puig, J. Quevedo, and F. Nejjari, "Lexicographic optimisation for optimal departure aircraft trajectories," *Aerospace Science and Technology*, vol. 14, no. 1, pp. 26–37, 2010.

[74] X. Prats, J. Quevedo, V. Puig, and F. Nejjari, "Hierarchical and sensitivity analysis for noise abatement departure procedures," in *Proceedings of the 8th AIAA Aviation Technology, Integration, and Operations(ATIO) Conference*, 2008.

[75] H. Pervier, D. Nalianda, R. Espi, V. Sethi, P. Pilidis, D. Zammit-Mangion, J.-M. Rogero, and R. Entz, "Application of genetic algorithm for preliminary trajectory optimization," *SAE International Journal of Aerospace*, vol. 4, no. 2, pp. 973–987, 2011.

[76] R. Torres, J. Chaptal, C. Bes, and J. Hiriart-Urruty, "Optimal, environmentally friendly departure procedures for civil aircraft," *Journal of Aircraft*, vol. 48, no. 1, pp. 11–22, 2011.

[77] J. Betts, "Survey of numerical methods for trajectory optimization," *Journal of Guidance, Control, and Dynamics*, vol. 21, no. 2, 1998.

[78] C. Lai and J. Whidborne, "Real-time trajectory generation for collision avoidance with obstacle uncertainty," in *AIAA Guidance, Navigation and Control Conference, AIAA 2011-6598*, Portland OR, August 2011.

[79] R. Drury, *Trajectory Generation for Autonomous Unmanned Aircraft Using Inverse Dynamics*. PhD Thesis Cranfield University, 2010.

[80] N. Akthar, J. F. Whidborne, and A. K. Cooke, "Real-time trajectory generation technique for dynamic soaring uavs," *Department of Aerospace Sciences, Cranfield University*, (2008).

[81] H. Alturbeh, *Collision Avoidance Systems for UAS Operating in Civil Airspace*. PhD Thesis Cranfield University, 2014.

[82] S. Hammond, J. Civil, M. Ross, and K. Slater, "Air traffic control, business regulation and CO2 emissions," *OR insight*, vol. 25, no. 3, pp. 127–149, 2012.

[83] C. N. Jardine, "Calculating the environmental impact of aviation emissions," *Environmental Change Institute-Oxford University*, 2005.

[84] U. Burkhardt and B. Kärcher, "Global radiative forcing from contrail cirrus," *Nature climate change*, vol. 1, no. 1, pp. 54–58, 2011.

[85] Cleansky ITD, Systems for Green Operations, "D_3.2.1_1, problem definition report," 12 February 2010.

[86] MIME Consortium, "Report on airport/ATM and market simulation, Volume 3 Appendices," February 2011.

[87] T. Thompson and P. Hullah, "Achieving climate-optimal trajectories," in *Proceedings of the 2014 AIAA Aircraft Noise and Emissions Reduction Symposium*, 16-20 June 2014, Atlanta, GA 2014.

[88] J. Faber, D. Greenwood, D. Lee, M. Mann, P. M. de Leon, D. Nelissen, B. Owen, M. Ralph, J. Tilston, A. van Velzen *et al.*, "Lower NOx at higher altitudes policies to reduce the climate impact of aviation NOx emission," Tech. Rep., 2008.

[89] R. Wit, B. Boon, A. Van Velzen, M. Cames, O. Deuber, and D. Lee, "Giving wings to emission trading," *Inclusion of aviation under the European emission trading system (ETS): design and impacts. Report for the European Commission, DG Environment, Delft*, 2005.

[90] L. R. Jenkinson, P. Simpkin, D. Rhodes, and L. Jenkison, *Civil Jet Aircraft Design*. Arnold London, 1999, vol. 7.

[91] J. Green, "Air travel-greener by design mitigating the environmental impact of aviation: Opportunities and priorities-report of the air travel-greener by design science and technology sub group london," *Aeronautical Journal*, vol. 109, no. 1099, pp. 361–416, 2005.

[92] International Civil Aviation Organization, "Doc 9889, Airport Air Quality Manual," 2011.

[93] K. Kugele, F. Jelinek, and R. Gaffal, "Aircraft particulate matter emission estimation through all phases of flight," *EEC/SEE/2005/0014, EUROCONTROL Experimental Centre, France*, 2005.

[94] M. Schaefer, *Methodologies for Aviation Emission Calculation: a comparison of alternative approaches towards 4D global inventories*. PhD Thesis Univ.-Bibliothek der Techn, 2006.

[95] UK Civil Aviation Authority, "Environmental charging, Review of impact of noise and NOx landing charges," 2007.

[96] European Civil Aviation Conference, "Recommendation ECAC/27-4 NOx emission classification scheme," 3 September 2011.

[97] S. Baughcum *et al.*, *Scheduled Civil Aircraft Emission Inventories for 1992: Database Development and Analysis*. National Aeronautics and Space Administration, Langley Research Center, 1996.

[98] E. Fleuti and P. Hofmann, "ADAECAM case study Zurich airport 2006," 25 September 2007.

[99] Federal Office of Civil Aviation, Swiss Confederation, "Validation of ADAECAM (Advanced Aircraft Emissions Calculation Method) report on fuel calculation," 10 August 2007.

[100] DfT (Department for Transport), "Project for the sustainable development of heathrow: Air quality technical report," Report of the PSDH Technical Panels to the DfT, 2006.

[101] B. Kim, G. Fleming, S. Balasubramanian, A. Malwitz, J. Lee, J. Ruggiero, I. Waitz, K. Klima, V. Stouffer, L. Long *et al.*, "System for assessing aviation's global emissions (sage). version 1.5, technical manual, federal aviation administration, office of environment and energy," FAA-EE-2005-01, Tech. Rep., 2005.

[102] B. Figlar, S. Ottl, G.and Schwanke, R. M. Rodrguez, T. Gjestland, A. Guhnemann, H. Harwatt, B. G. Tight, M., and P. Hullah, "State of the art on tradable permits, noise legislation, noise restriction methods and noise modelling," in *EEC MIME WP6.1 SotA*, January 2010.

[103] M. J. Smith, *Aircraft Noise*. Cambridge University Press, 2004, vol. 3.

[104] International Civil Aviation Organization, *Annex 16 to the Convention on International Civil Aviation, Volume I*, 2005.

[105] I. L. N. Granøien and T. Gjestland, "Mime - market-based impact mitigation for the environment," in *Baltic-Nordic Acoustics Meeting (BNAM)*, Bergen, Norway, May 2010.

[106] R. A. Wijnen and H. G. Visser, "Optimal departure trajectories with respect to sleep disturbance," *Aerospace science and technology*, vol. 7, no. 1, pp. 81–91, 2003.

[107] S. Hartjes, H. G. Visser, and S. J. Hebly, "Optimization of RNAV noise and emission abatement departure procedures," in *9th AIAA Aviation Technology, Integration, and Operations Conference (ATIO) and Aircraft Noise and Emissions Reduction Symposium (ANERS)*, 2009.

[108] M. Richter, F. Fisch, and F. Holzapfel, "Noise-minimal landing and take-off trajectories under procedural and safety regulations," in *Proceedings of the Third International Air Transport and Operations Symposium 2012*. IOS Press, 2012, p. 77.

[109] H. Yu and J. Mulder, "Noise abatement trajectory optimization using genetic algorithms," in *AIAA Guidance, Navigation, and Control Conference*, 2012.

[110] ECAC, "Standard method of computing noise contours around civil airports ECAC. CEAC Doc. 29," European Civil Aviation Conference, Tech. Rep., 2005.

[111] J. Olmstead, G. Fleming, J. Gulding, C. Roof, P. Gerbi, and A. Rapoza, "Integrated Noise Model (INM) version 7.0 technical manual," Report FAA-AEE-02-01, Office of Environment and Energy, Federal Aviation Administration, Tech. Rep., 2002.

[112] J. Ollerhead, *The CAA Aircraft Noise Contour Model: ANCON Version 1*, 1992.

[113] D. E. Kirk, *Optimal Control Theory: an Introduction*. DoverPublications, 2012.

[114] D. G. Hull, "Fundamentals of airplane flight mechanics," Springer, 2007.

[115] A. K. Nandakumaran, "Calculus of variations to optimal control: A short introduction," https://www.icts.res.in/media/uploads/Program/Files/Nands_part%201_1.pdf, 12 2012.

[116] J. T. Betts, *Practical Methods for Optimal Control and Estimation using Nonlinear Programming*. Siam, 2010, vol. 19.

[117] A. V. Rao, "A survey of numerical methods for optimal control," *Advances in the Astronautical Sciences*, vol. 135, no. 1, pp. 497–528, 2009.

[118] D. G. Hull, "Conversion of optimal control problems into parameter optimization problems," *Journal of Guidance, Control, and Dynamics*, vol. 20, no. 1, pp. 57–60, 1997.

[119] M. Bittner, F. Fisch, and F. Holzapfel, "A multi-model gauss pseudospectral optimization method for aircraft trajectories," in *AIAA Atmospheric Flight Mechanics Conference*, 2012.

[120] F. Fahroo and I. Ross, "Direct trajectory optimization by a chebyshev pseudospectral method," in *Proceedings of the 2000 American Control Conference*, vol. 6. IEEE, 2000, pp. 3860–3864.

[121] O. Yakimenko, "Direct method for rapid prototyping of near-optimal aircraft trajectories," *AIAA Journal of Guidance, Control, and Dynamics*, vol. 23, no. 5, 2000.

[122] V. T. Taranenko, "Experience on application of Ritzs, Poincares, and Lyapunovs methods in solving of flight dynamics problems," *Air Force Engineering Academy Press Moscow*, 1968.

[123] M. Etchemendy, *Flight Control and Optimal Path Planning for UAVs.* MSc Thesis Cranfield University, 2007.

[124] I. D. Cowling, O. A. Yakimenko, J. F. Whidborne, and A. K. Cooke, "Direct method based control system for an autonomous quadrotor," *Journal of Intelligent & Robotic Systems*, vol. 60, no. 2, pp. 285–316, 2010.

[125] R. Drury, A. Tsourdos, and A. Cooke, "Real-time trajectory generation: Improving the optimality and speed of an inverse dynamics method," in *2010 IEEE Aerospace Conference.* IEEE, 2010, pp. 1–12.

[126] W. Glover and J. Lygeros, "A stochastic hybrid model for air traffic control simulation," in *Hybrid Systems: Computation and Control.* Springer, 2004, pp. 372–386.

[127] R. Alligier, D. Gianazza, and N. Durand, "Learning the aircraft mass and thrust to improve the ground-based trajectory prediction of climbing flights," *Transportation Research Part C: Emerging Technologies*, vol. 36, pp. 45–60, 2013.

[128] A. Nuic, "User manual for the Base of Aircraft Data (BADA) revision 3.10," 2012.

[129] K. Culligan, M. Valenti, Y. Kuwata, and J. P. How, "Three-dimensional flight experiments using on-line mixed-integer linear programming trajectory optimization," in *American Control Conference, 2007. ACC'07.* IEEE, 2007, pp. 5322–5327.

[130] A. G. Richards, *Trajectory optimization using mixed-integer linear programming.* PhD Thesis Massachusetts Institute of Technology, 2002.

[131] P. Schneider and D. H. Eberly, *Geometric Tools for Computer Graphics.* Morgan Kaufmann, 2002.

[132] D. Mellinger and V. Kumar, "Minimum snap trajectory generation and control for quadrotors," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on.* IEEE, 2011, pp. 2520–2525.

[133] J. García-Heras, M. Soler, and F. J. Sáez, "A comparison of optimal control methods for minimum fuel cruise at constant altitude and course with fixed arrival time," *Procedia Engineering*, vol. 80, pp. 231–244, 2014.

[134] D. Delahaye, S. Puechmorel, P. Tsiotras, and E. Feron, "Mathematical models for aircraft trajectory design: a survey," in *Air Traffic Management and Systems*. Springer, 2014, pp. 205–247.

[135] G. Fasano and J. Pintér, *Modeling and Optimization in Space Engineering*. Springer, 2013.

[136] F. Fahroo and I. Ross, "Advances in pseudospectral methods for optimal control," in *AIAA Guidance, Navigation and Control Conference and Exhibit*, 2008.

[137] F. Fahroo and I. M. Ross, "On discrete-time optimality conditions for pseudospectral methods," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*. AIAA, 2006.

[138] G. T. Huntington, *Advancement and Analysis of a Gauss Pseudospectral Transcription for Optimal Control Problems*. PhD Thesis Massachusetts Institute of Technology, 2007, vol. A.

[139] J.-F. Bonnans, J. C. Gilbert, C. Lemaréchal, and C. A. Sagastizábal, *Numerical optimization: theoretical and practical aspects*. Springer, 2006.

[140] P. T. Boggs and J. W. Tolle, "Sequential quadratic programming," *Acta numerica*, vol. 4, pp. 1–51, 1995.

[141] P. E. Gill, W. Murray, and M. A. Saunders, "Snopt: An sqp algorithm for large-scale constrained optimization," *SIAM journal on optimization*, vol. 12, no. 4, pp. 979–1006, 2002.

[142] K. Schittkowski, "NLPQL: A fortran subroutine solving constrained nonlinear programming problems," *Annals of operations research*, vol. 5, no. 2, pp. 485–500, 1986.

[143] R. T. Marler and J. S. Arora, "Survey of multi-objective optimization methods for engineering," *Structural and multidisciplinary optimization*, vol. 26, no. 6, pp. 369–395, 2004.

[144] S. Hartjes, J. Dons, and H. Visser, "Optimization of area navigation arrival routes for cumulative noise exposure," *Journal of Aircraft*, vol. 51, no. 5, pp. 1432–1438, 2014.

[145] K. Price, R. Storn, and J. Lampinen, *Differential Evolution: a Practical Approach to Global Optimization*. Springer-Verlag New York Inc, 2005.

[146] N. Madavan, "Multiobjective optimization using a pareto differential evolution approach," in *Proceedings of the 2002 Congress on Evolutionary Computation, 2002. CEC'02*, vol. 2. IEEE, 2002, pp. 1145–1150.

[147] M. J. Riley, T. Peachey, D. Abramson, and K. W. Jenkins, "Multi-objective engineering shape optimization using differential evolution interfaced to the nimrod/o tool," in *IOP Conference Series: Materials Science and Engineering*, vol. 10, no. 1. IOP Publishing, 2010, p. 012189.

[148] 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.

[149] R. G. Drury, "Performance of NLP algorithms with inverse dynamics for near-real time trajectory generation," in *AIAA Guidance, Navigation, and Control Conference*, August 2011.

[150] C. M. Fonseca, P. J. Fleming *et al.*, "Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization." in *ICGA*, vol. 93, 1993, pp. 416–423.

[151] N. Srinivas and K. Deb, "Muiltiobjective optimization using nondominated sorting in genetic algorithms," *Evolutionary computation*, vol. 2, no. 3, pp. 221–248, 1994.

[152] J. Lampinen *et al.*, "DEs selection rule for multiobjective optimization," *Lappeenranta University of Technology, Department of Information Technology, Tech. Rep*, pp. 03–04, 2001.

[153] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: Nsga-ii," *Evolutionary Computation, IEEE Transactions on*, vol. 6, no. 2, pp. 182–197, 2002.

[154] T. Robic and B. Filipic, "DEMO: Differential evolution for multiobjective optimization," in *Evolutionary Multi-Criterion Optimization*. Springer, 2005, pp. 520–533.

[155] S. Kukkonen and J. Lampinen, "GDE3: The third evolution step of generalized differential evolution," in *Evolutionary Computation, 2005. The 2005 IEEE Congress on*, vol. 1. IEEE, 2005, pp. 443–450.

[156] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler, *Scalable Test Problems for Evolutionary Multiobjective Optimization*. Springer, 2005.

[157] S. Kukkonen and K. Deb, "A fast and effective method for pruning of non-dominated solutions in many-objective problems," *Lecture Notes in Computer Science*, vol. 4193, pp. 553–562, 2006.

[158] K. Deb, *Multi-Objective Pptimization using Evolutionary Algorithms*. John Wiley & Sons, 2001, vol. 16.

[159] S.-W. Ra and J. Kim, "A fast mean-distance-ordered partial codebook search algorithm for image vector quantization," *Circuits and Systems II: Analog and Digital Signal Processing, IEEE Transactions on*, vol. 40, no. 9, pp. 576–579, 1993.

[160] S. Baek and K.-M. Sung, "Fast k-nearest-neighbour search algorithm for nonparametric classification," *Electronics Letters*, vol. 36, no. 21, pp. 1821–1822, 2000.

[161] Sourdine Consortium, "Report D3 Study of Optimisation procedURes for Decreasing the Impact of NoisE (SOURDINE)," 2001.

[162] J. DErrico. (2006) Surface fitting using gridfit. MATLAB Central File Exchange. [Online]. Available: http://www.mathworks.com/matlabcentral/fileexchange/8998

[163] N. Lopez, O. Aguirre, J. F. Espiritu, and H. A. Taboada, "Using game theory as a post-pareto analysis for renewable energy integration problems considering multiple objectives," in *In Proceedings of the 41st International Conference on Computers & Industrial Engineering*, 2011, pp. 678–683.

[164] J. J. Durillo, A. J. Nebro, and E. Alba, "The jmetal framework for multi-objective optimization: Design and architecture," in *2010 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, 2010, pp. 1–8.

[165] T. G. Reynolds, "Advanced noise abatement approach activities at nottingham east midlands airport, uk," in *Seventh USA/Europe Air Traffic Management R&D Seminar (ATM 2007), Barcelona*, 2007.

[166] NATS, "Consultation document for the Terminal Control North (TCN) airspace change proposal," *NATS*, 2008.

[167] F. J. Gallego, "A population density grid of the european union," *Population and Environment*, vol. 31, no. 6, pp. 460–473, 2010.

[168] ESDU, *ESDU 90012, Energy height method for flight path optimisation*, ESDU Std., 07 1990.

[169] ——, *ESDU 91016, Energy height method for flight path optimisation. Addendum A: numerical method suitable for rapid use on IBM PC compatible computers*, ESDU Std., 09 1991.

[170] UK Civil Aviation Authority, "Consultation on NATS (En Route) plc (NERL) flight efficiency performance regime," http://www.caa.co.uk/docs/5/Flight%20efficiency%20consultation%20letter%2002082011.pdf, 02 2011.

[171] NATS, "Environmental performance," http://www.nats.aero/environment/3di/, 11 2014.

[172] ——, "3Di environmental performance measure," http://www.nats.aero/wp-content/uploads/2012/07/3di_Infocard.pdf, 07 2012.

[173] ——, "NATS records airspace efficiency improvement in 2013," http://www.nats.aero/news/nats-records-airspace-efficiency-improvement-2013/, 03 2014.

[174] Eurocontrol, "Eurocontrol response, consultation on NATS (En Route) plc (NERL) flight efficiency performance regime," https://www.caa.co.uk/docs/5/20111013%20Eurocontrol.pdf, 10 2011.

[175] NATS, "NATS, British Airways and BAA in UK-first with perfect flight," http://www.nats.aero/news/nats-british-airways-and-baa-in-uk-first-with-perfect-flight/, 07 2010.

[176] Q. McEnteggart and J. Whidborne, "A multiobjective trajectory optimisation method for planning environmentally efficient trajectories," in *2012 UKACC International Conference on Control*. IEEE, 2012, pp. 128–135.

[177] RouteFinder, "Route generator for PC flight simulation use," http://rfinder.asalink.net/free/.

[178] J. J. Durillo and A. J. Nebro, "jmetal: A java framework for multi-objective optimization," *Advances in Engineering Software*, vol. 42, no. 10, pp. 760–771, 2011.

[179] E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonseca, and V. G. Da Fonseca, "Performance assessment of multiobjective optimizers: An analysis and review," *Evolutionary Computation, IEEE Transactions on*, vol. 7, no. 2, pp. 117–132, 2003.

[180] D. Poles, "Base of Aircraft Data (BADA) aircraft performance modelling report," EEC Technical/Scientific Report No. 2009-009, March 2009.

[181] M. R. Spiegel, "Mathematical handbook of formulas and tables," 1968.

Appendix A

Metaheuristic Performance

A.1 Metrics

The metrics used to asses the performance of the multiobjective heuristics in this thesis are summerised in this section. The metrics are intended to measure the convergence performance of the algorithms to the Pareto optimal front and the diversity of solutions in the converged front. The metrics used to measure convergence were Generational Distance and Epsilon. Metrics used principally to measure diversity were Spacing, Spread and Maximum Spread. The Hypervolume metric was also used as this provided a combined measure of both criteria.

Generational Distance
Generational Distance measures the closeness between an approximated Pareto front $\mathcal{Q}$ and a true Pareto front $\mathcal{Q}^*$ [158]. To calculate the Generational Distance, the minimum Euclidean distance d_i from each point on the approximate front to the nearest point on the Pareto optimal front is first determined. The Generational Distance (GD) is then calculated as

$$GD = \frac{\sqrt{\sum_{i=1}^{n_s} d_i^2}}{n_s} \tag{A.1}$$

where n_s is the number of solutions in the approximate Pareto front. A value of $GD = 0$ indicates that all the solutions in the approximated Pareto front are also also members of the true Pareto front.

Hypervolume
The Hypervolume calculates the volume in the objective space, covered by members of a non-dominated set of solutions $\mathcal{Q}$ [178]. The objective space enclosed by the non-dominated front is constructed with the aid of a reference point O_{ref}, which can be defined as a vector of the worst objective function values [178]. The enclosed space is divided into hypercubes. The Hypervolume (HV) is then calculated as the union of all the hypercubes

$$HV = \text{volume}\left(\bigcup_{i=1}^{|\mathcal{Q}|} v_i\right) \tag{A.2}$$

where for each nondominated solution $s_i \in \mathcal{Q}$, v_i is the individual hypercube volume.

Figure A.1 shows a 2D set of nondominated solutions $\{s_1, s_2, s_3\} \in \mathcal{Q}$, where the volumes of the rectangular hypercubes are calculated by multiplying the length by the breadth of each rectangle and the Hypervolume is then the sum of all the individual volumes. Generally, a larger value of Hypervolume is better.

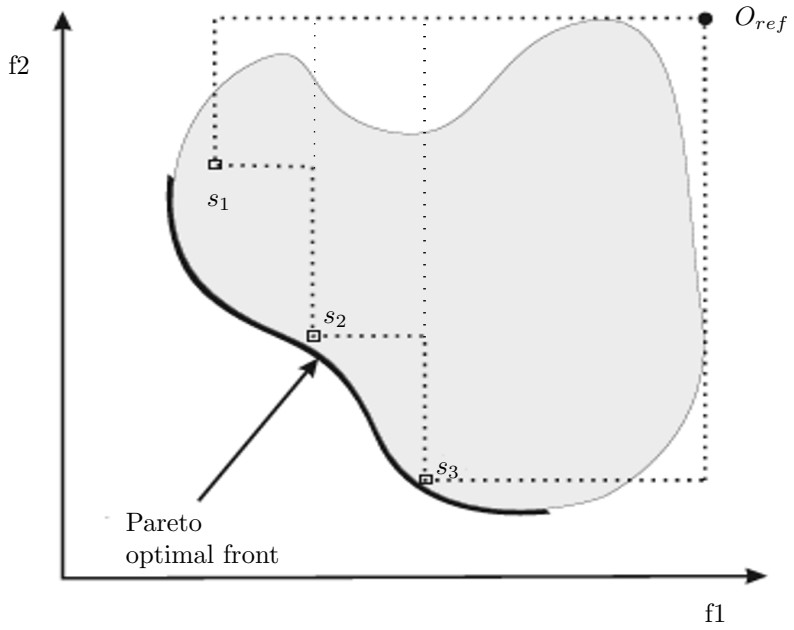

FIGURE A.1: Hypervolume for a 2D Pareto front. Adapted from [178]

Epsilon

The additive ϵ-indicator is a difference value by which a nondominated solution set is worse than another nondominated solution set with respect to all objectives [179]. For two sets of solutions $\mathcal{Q}_1$ and $\mathcal{Q}_2$, the additive ϵ is the minimum difference such that for any solution in $\mathcal{Q}_2$ there is at least one solution in $\mathcal{Q}_1$ that is not worse by a difference of ϵ in all objectives. More formally, where $\boldsymbol{z}^1 = \{z_1^1, z_2^1, \ldots, z_n^1\}$ and $\boldsymbol{z}^2 = \{z_1^2, z_2^2, \ldots, z_n^2\}$ are objective vectors of length n, the additive ϵ-indicator $I_{\epsilon+}$ is deined as:

$$I_{\epsilon+}(\mathcal{Q}_1, \mathcal{Q}_2) = \inf_{\epsilon \in \mathbb{R}} \{\forall\, \boldsymbol{z}^2 \in \mathcal{Q}_2\, \exists \boldsymbol{z}^1 \in \mathcal{Q}_1 : \boldsymbol{z}^1 \succeq_{\epsilon+} \boldsymbol{z}^2\} \tag{A.3}$$

where $\boldsymbol{z}^1 \succeq_{\epsilon+} \boldsymbol{z}^2$ if and only if $\forall\, 1 \leq i \leq n : z_i^1 \leq \epsilon + z_i^2$

Spacing

Spacing is a metric that provides a measure of the relative distance between adjacent solutions in a nondominated set [158]. Spacing is calculated as follows

$$SPC = \sqrt{\frac{1}{|\mathcal{Q}|}\sum_{i=1}^{|\mathcal{Q}|}(d_i - \bar{d})^2} \tag{A.4}$$

where $\mathcal{Q}$ is a set of solutions that form a nondominated front. The values for d_i are the Euclidean distance between consecutive solutions $s_i \in \mathcal{Q}$, and $\bar{d}$ is the mean of the individual distances. Smaller values of SPC indicate a more uniform distribution and $SP = 0$ is an ideal distribution.

Spread

The Spread metric measures the extent or spread of the solutions for a nondominated front $\mathcal{Q}$. The Spread is calculated as

$$SPD = \frac{\sum_{m=1}^{M} d_m^e + \sum_{i=1}^{|\mathcal{Q}|} |d_i - \bar{d}|}{\sum_{m=1}^{M} d_m^e + |\mathcal{Q}|\bar{d}} \tag{A.5}$$

where d_i is again the Euclidean distance between consecutive solutions, and $\bar{d}$ is the mean of these distances. The parameter d_m^e is the distance between the extrema solutions of the nondominated front $\mathcal{Q}$ and the true Pareto optimal front $\mathcal{Q}^*$ for the m^{th} objective function [158]. For an ideal distribution of solutions, $SPD = 0$ [158].

Maximum Spread

Figure A.2 shows the Maximum Spread SPD_M for a 2D Pareto front, which, in this case, is the Euclidean distance between the two extrema solutions. More generally, the Maximum Spread is defined as the length of the diagonal of a hyperbox [158], and can be calculated as

$$SPD_M = \sqrt{\sum_{m}^{|M|} = \left(\max_{i=1}^{|\mathcal{Q}|} f_m^{s_i} - \min_{i=1}^{|\mathcal{Q}|} f_m^{s_i}\right)} \tag{A.6}$$

where s_i is the i^{th} solution $s_i \in \mathcal{Q}$ and m is the m^{th} objective value. Larger values of SPD_M show that there is a larger spread between the extreme solutions.

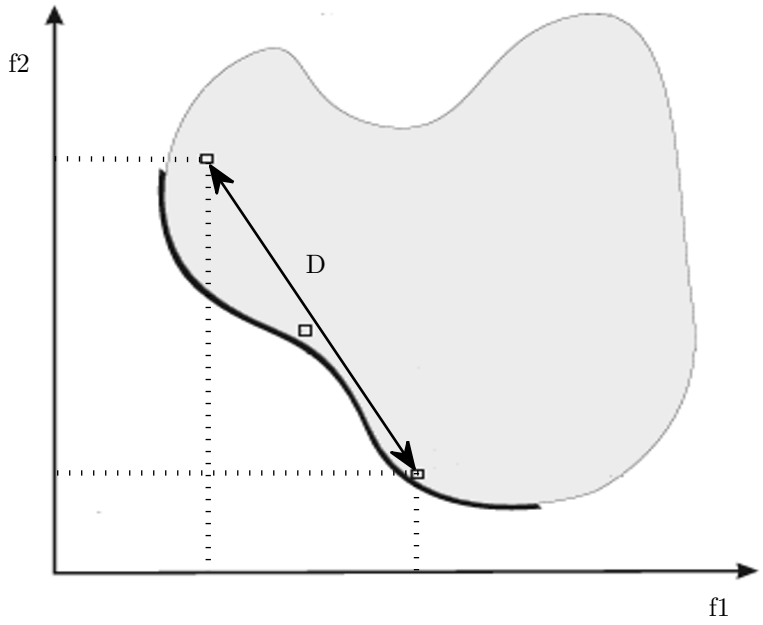

FIGURE A.2: Maximum Spread for a 2D Pareto front

A.2 Test Problems

The DEMO and the DEMO,KNN algorithms used in this Thesis were highly similar to the GDE3,CD and the GDE,MNN algorithms from [155] and [157]. In [157] Kukkonen assesses the performance of the GDE3,CD and the GDE,MNN algorithms for the ZDT1, ZDT2, ZDT3, ZDT4, and ZDT6 test problems using the Hypervolume, Spacing and Maximum Spread metrics. The results are reproduced in Table A.1 where they are also compared against the results of the DEMO and DEMO,KNN algorithms for the same metrics and test problems.

For all of the algorithms used with the ZDT1, ZDT2, ZDT3 and ZDT6 test problems, the differential mutation factor and the crossover coefficient were both set to 0.2. For the ZDT4 problem, the differential mutation factor was set to 0.5 and the crossover coefficient was also set to 0.0. For the calculation of the Hypervolume metric, the reference point W was set to $[2.0, 2.0]$

Each test was repeated 100 times, and the mean and standard deviations for each metric and each algorithm was recorded in Table A.1. Table A.1 then shows that the performances of the DEMO and DEMO,KNN algorithms implemented in this Thesis were equivalent to the performance results for the GDE3,CD and the GDE,MNN algorithms when compared for the same test problems and metrics.

It can also be seen from Table A.1 and from from the Pareto front plots in Figure A.3, that the DEMO and the DEMO,KNN algorithms exhibit similar convergence performance

for the two objective ZDT problems. However, the distribution of the solutions along the Pareto fronts were slightly more uniform for DEMO,KNN algorithm than for the DEMO algorithm. Unlike the DEMO algorithm, the DEMO,KNN algorithm updates the crowding distance measures of the solutions in the nondominated solution set every time the most crowded solution is removed. This acts to improve the improve the diversity of the solutions in the converged front. However, from Table A.2 it can be seen that over the 5 ZDT test problems, the DEMO,KNN algorithm was an average of 12 seconds (2.75%) slower than the DEMO algorithm.

		HV_mean	HV_std	SPC_mean	SPC_std	SPD_M_mean	SPD_M_std
ZDT1	DEMO	3.6594	0.0039436	0.00623	0.00045236	1.4104	0.0025608
	GDE3, CD	3.661	0.0012739	0.006424	0.00055946	1.4113	0.0014634
	DEMO,KNN	3.6609	0.0023577	0.0022007	0.0002167	1.4104	0.0018651
	GDE3,M-NN	3.6616	0.0011364	0.0026305	0.00028921	1.4115	0.0015376
ZDT2	DEMO	3.3242	0.010472	0.0062694	0.00055112	1.4132	0.0021792
	GDE3, CD	3.3274	0.0025905	0.0063904	0.00054882	1.414	0.0005474
	DEMO,KNN	3.3235	0.013965	0.0021797	0.00019249	1.413	0.0029428
	GDE3,M-NN	3.3281	0.003556	0.002764	0.00022573	1.414	0.00074998
ZDT3	DEMO	4.8149	0.00087774	0.026685	0.00015046	1.9628	0.0022321
	GDE3, CD	4.8151	0.0001063	0.0043573	0.00041398	1.9639	0.0017943
	DEMO,KNN	4.8152	0.00079262	0.02579	3.18E-05	1.9628	0.0024211
	GDE3,M-NN	4.8154	2.5151E-05	0.0016415	0.00016563	1.9639	0.0015796
ZDT4	DEMO	3.6513	0.042809	0.01409	0.00058074	1.4098	0.012892
	GDE3, CD	3.6613	0.00016689	0.0060789	0.00055549	1.4112	0.0018113
	DEMO,KNN	3.6535	0.039554	0.0029066	0.00015299	1.41	0.011093
	GDE3,M-NN	3.6589	0.030332	0.0022912	0.00029246	1.4123	0.0083372
ZDT6	DEMO	2.8417	0.47954	0.0049917	0.00055303	1.1335	0.070173
	GDE3, CD	3.0209	0.13448	0.0066047	0.00065536	1.1648	0.02526
	DEMO,KNN	3.03685	0.15814	0.0019079	0.00019492	1.1677	0.02845
	GDE3,M-NN	3.0114	0.1818	0.0026386	0.00028531	1.1632	0.031709

TABLE A.1: Algorithm performance for the ZDT test problems

Problem	Method	time_mean	time_std
ZDT1	DEMO	335.22	3.3547
	DEMO,KNN	345.23	1.4886
ZDT2	DEMO	335.97	2.6674
	DEMO,KNN	344.16	1.1388
ZDT3	DEMO	321.22	2.4867
	DEMO,KNN	323.54	3.4789
ZDT4	DEMO	596	14.04
	DEMO,KNN	626.27	20.144
ZDT6	DEMO	361.43	1.748
	DEMO,KNN	370.66	2.435

TABLE A.2: Algorithm run times for the ZDT test problems

To test the performance of the DEMO and the DEMO,KNN algorithms for problems with more than 2 objectives, the DTLZ1, DTLZ2, DTLZ4, DTLZ6 and DTLZ7 test problems were used. The test problems were developed by Zitzler in [156] who also noted that the performance metrics used for 2 objective problems are not suitable for assessing problems having three or more objectives. Therefore the Pareto front plots for the DTLZ problems for both the DEMO and the DEMO,KNN algorithms were compared against each other in Figures A.4 and A.5. The approximated fronts were also checked against the Pareto optimal fronts published in [157]. In Figures A.4 and A.5 it can be clearly seen that for the same problems, the solutions for DEMO,KNN algorithm have much better diversity than the solution for the DEMO algorithm and that this allowed the

DEMO,KNN algorithm to provide a much better approximation of the Pareto optimal front than the DEMO algorithm.

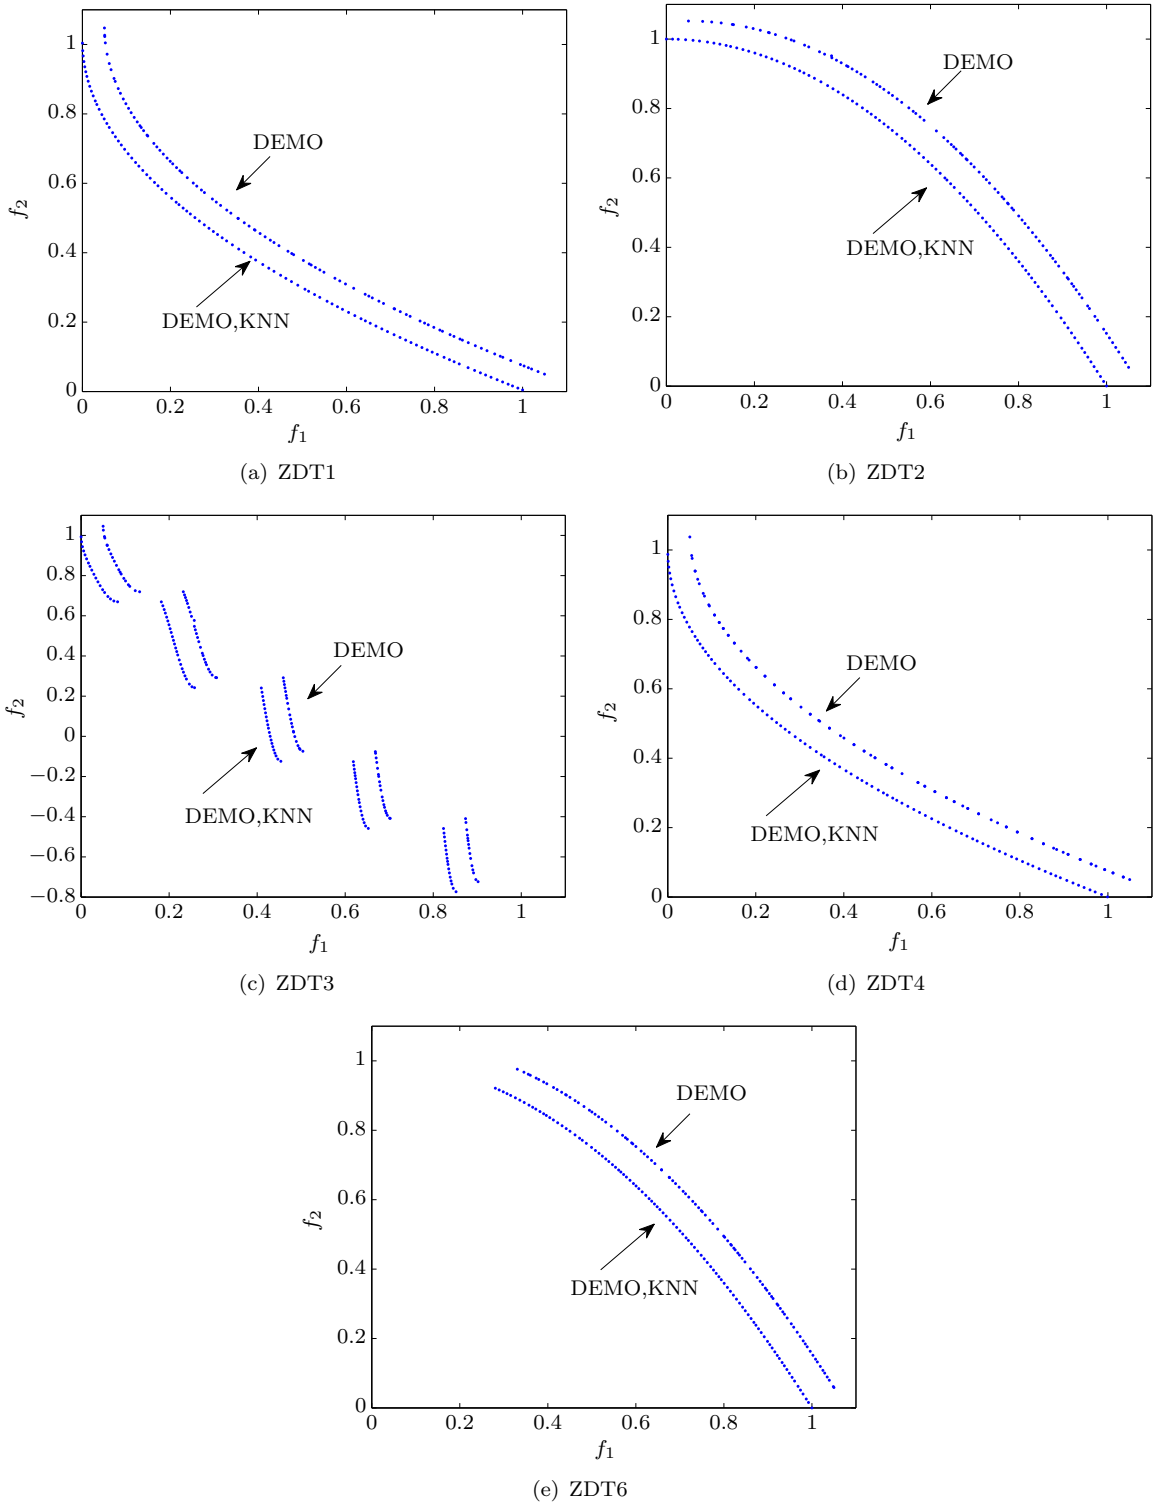

(a) ZDT1

(b) ZDT2

(c) ZDT3

(d) ZDT4

(e) ZDT6

FIGURE A.3: ZDT Pareto front solutions for the DEMO and DEMO,KNN algorithms (sets of solutions have been offset by 0.05 units along both objectives for comparison purposes)

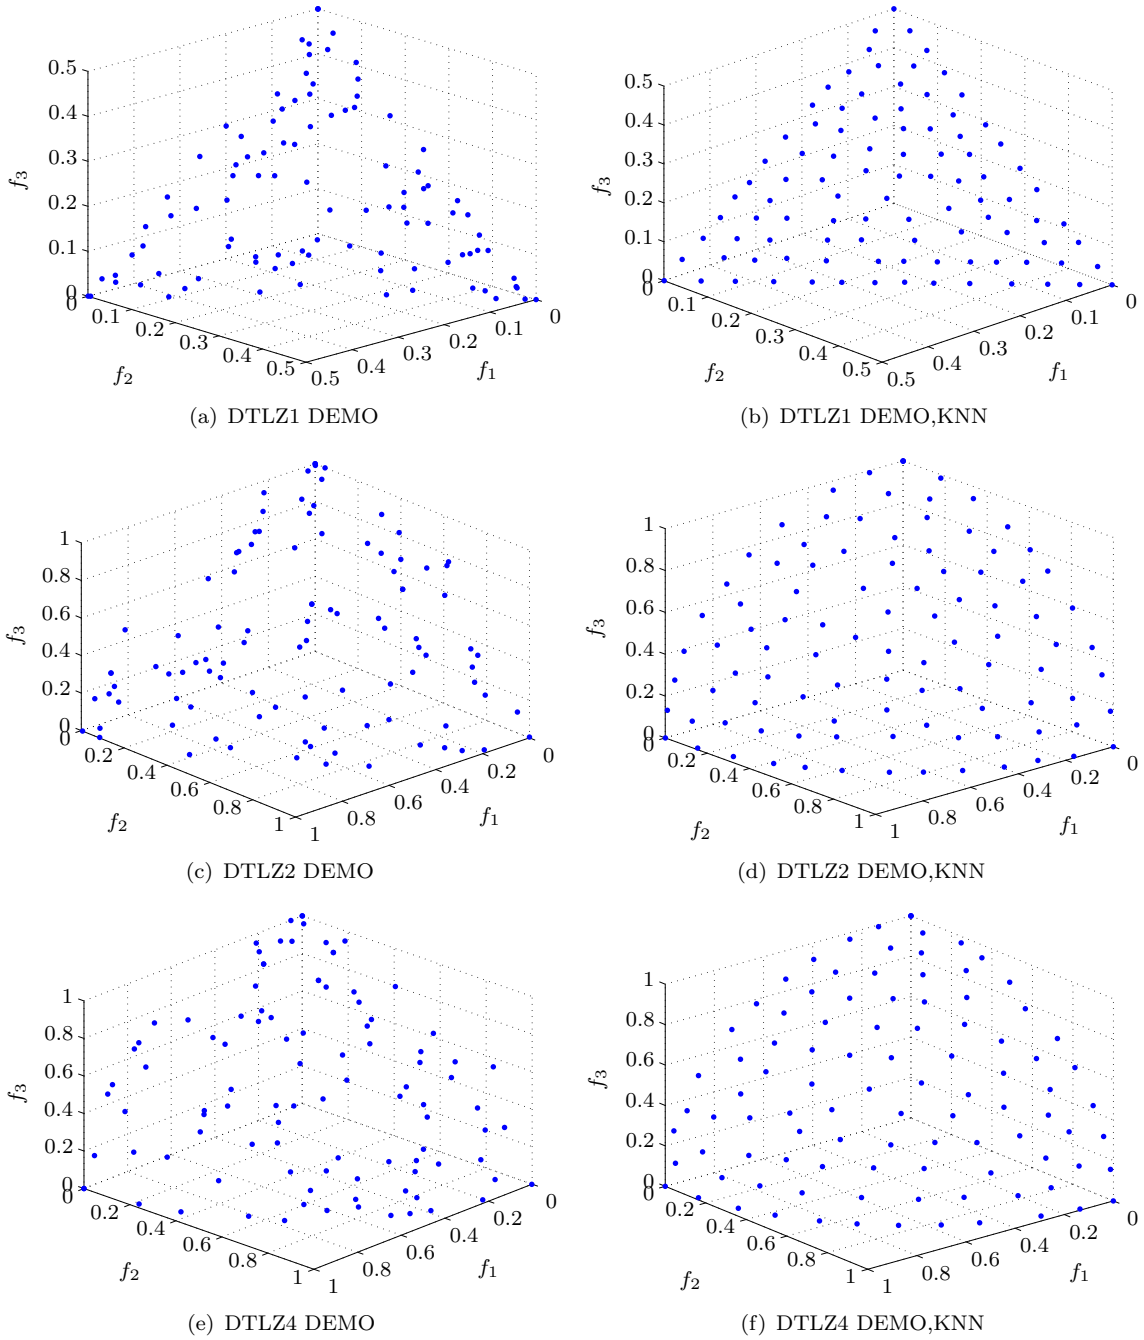

(a) DTLZ1 DEMO

(b) DTLZ1 DEMO,KNN

(c) DTLZ2 DEMO

(d) DTLZ2 DEMO,KNN

(e) DTLZ4 DEMO

(f) DTLZ4 DEMO,KNN

FIGURE A.4: DTL Pareto front solutions for the DEMO and DEMO,KNN algorithms

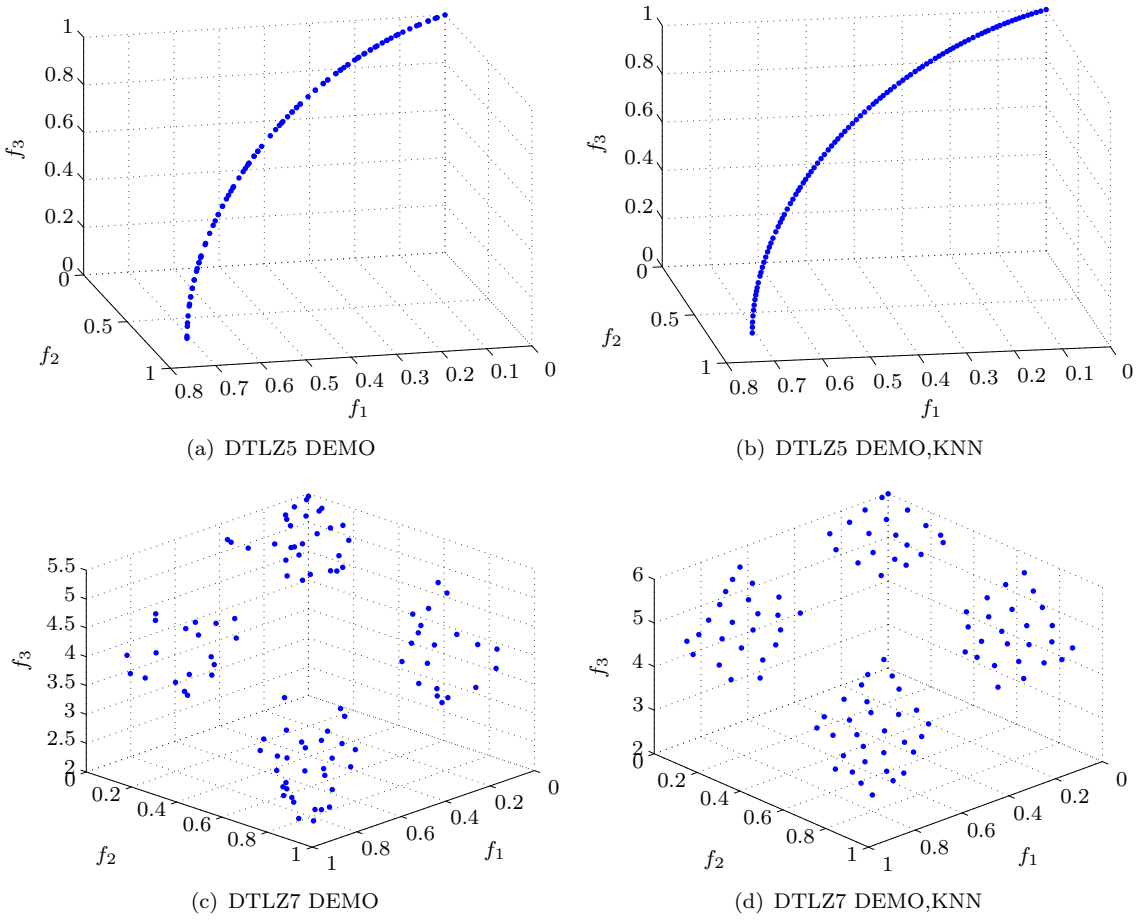

(a) DTLZ5 DEMO

(b) DTLZ5 DEMO,KNN

(c) DTLZ7 DEMO

(d) DTLZ7 DEMO,KNN

FIGURE A.5: DTL Pareto front solutions for the DEMO and DEMO,KNN algorithms

Appendix B

BADA

BADA is an aircraft performance model developed by Eurocontrol for use in trajectory simulation and prediction algorithms within the domain of Air Traffic Management [128]. The BADA database contains parameters for the generation of aircraft trajectories with typical aircraft operational height, speed and configuration schedules as well as coefficients for the modelling of related trajectory thrust, drag and fuel flows for 399 different aircraft types [128].

To determine the BADA aircraft performance parameters, data is taken from a variety of sources, including aircraft manufacturers documentation, radar data and expert analysis [180]. The data is used to generate reference trajectories consisting of profiles for height, speed and rate of climb/descent for a number of different aircraft masses, under a number operational conditions and at a number of atmospheric conditions. The BADA coefficients for the calculation of drag, thrust and fuel flow are derived using a non-linear multivariate parameter estimation process that then matches as closely as possible the reference trajectories with trajectories generated using the BADA equations of motion [180].

For trajectory generation, the BADA model uses a Total Energy Model (TEM), resolving the forces calculated from parameterised engine thrust and airframe drag models into vertical and horizontal accelerations [128]

$$\frac{dh}{dt} = \left[\frac{(T-D)v_t}{mg}\right] ES \tag{B.1}$$

where h is the aircraft height, T is thrust, D is drag, v_t is true airspeed, m is the aircraft mass and $ES \in [0, 1]$ is the energy share index that determines how much of the available power is used to accelerate or decelerate the aircraft. Generally the ES is set to one of a set of constant values that depends on the speed schedule of the aircraft.

The speed schedules of each aircraft are defined in the BADA procedures model. The speed schedules are designed to reflect the typical airline speed schedules for climb, cruise,

and descent operations. The default speed schedules in the BADA procedures model may be replaced by user defined speed schedules tailored to better reflect local operating constraints and procedures. In Section 7.2 the BADA model was used to generate baseline trajectories with user defined speed schedules and flexible ESI to better capture specific operating procedures. Details of modifying the BADA default speed schedule can be found in [128] and [180].

In addition to the generation of baseline trajectories, the BADA equations and coefficients were used with the IDVD-DE method to model aircraft drag and to set the aircraft flight envelope, including setting stall speed constraints and maximum thrust level limits.

The stall speed v_{stall} in knots calibrated airspeed (CAS) for climbing and cruising aircraft is defined as

$$v_{stall} = \begin{cases} v_{stall,TO} & : \text{if } h_{ft}(t) < 400 \text{ ft} \\ v_{stall,IC} & : \text{if } 400 \text{ ft} < h_{ft}(t) < 2000 \text{ ft} \\ v_{stall,CR} & : \text{if } h_{ft}(t) > 2000 \text{ ft} \end{cases} \tag{B.2}$$

where $v_{stall,TO}$, $v_{stall,IC}$, $v_{stall,CR}$ are the BADA take-off, initial climb, and cruise stall speeds defined for each individual aircraft type. All the BADA speeds referred to in this section are defined in knots. h_{ft} is aircraft height in feet.

For descending aircraft v_{stall} is defined as

$$v_{stall} = \begin{cases} v_{stall,LD} & : \text{if } h_{ft}(t) < 3000 \text{ ft } \wedge\ v_{CAS}(t) < v_{min,APP} + 10 \text{ kts} \\ v_{stall,AP} & : \text{if } h_{ft}(t) < 3000 \text{ ft } \wedge\ v_{CAS}(t) > v_{min,APP} + 10 \text{ kts} \\ v_{stall,AP} & : \text{if } 3000 \text{ ft } < h_{ft}(t) < 8000 \text{ ft } \wedge\ v_{CAS}(t) < v_{min,CR} + 10 \text{ kts} \\ v_{stall,CR} & : \text{if } 3000 \text{ ft } < h_{ft}(t) < 8000 \text{ ft } \wedge\ v_{CAS}(t) > v_{min,CR} + 10 \text{ kts} \\ v_{stall,CR} & : \text{if } h_{ft}(t) > 8000 \text{ ft} \end{cases} \tag{B.3}$$

where $v_{CAS}(t)$ is the calibrated airspeed in knots, $v_{stall,LD}$, $v_{stall,AP}$, $v_{stall,CR}$ are the BADA landing, approach and cruise stall speeds defined for each individual aircraft type. The minimum flight speed v_{min} is then calculated in CAS as

$$v_{min} = \begin{cases} v_{min,\text{Buffet}} & \text{if: } \mathsf{H_p} > 15000 \text{ ft} \wedge v_{min,\text{Buffet}} > v_{stall} \\ 1.2 \times v_{stall} \times C_{mass} & \text{if: } \mathsf{H_p} < 15000 \text{ ft} \wedge \text{if in take-off} \\ 1.3 \times v_{stall} \times C_{mass} & \text{if: } \mathsf{H_p} < 15000 \text{ ft} \wedge \text{if in any other phase of flight} \end{cases} \tag{B.4}$$

where

$$C_{mass} = \sqrt{\frac{m}{m_{ref}}} \tag{B.5}$$

and where m_{ref} is the BADA reference mass for the aircraft type. $v_{min,\text{Buffet}}$ (kts) is calculated as

$$v_{min,\text{Buffet}} = 1.9438 \times M_B \times \sqrt{\kappa \mathsf{R_g} \mathsf{T}} \tag{B.6}$$

where κ is the adiabatic index of air, $\mathsf{R_g}$ is the real gas constant for air [$\text{m}^2/(\text{K}{\cdot}\text{s}^2)$] and T is the temperature (K). M_B is buffeting limit Mach number, which is found from the roots of the cubic expression

$$M_B^3 - \frac{C_{Lbo(M_B=0)}}{k_{cl}} M_B^2 + \frac{\frac{W_{ac}}{S_w}}{0.583\,\mathsf{P}\,k_{cl}} = 0 \tag{B.7}$$

where k_{cl} is the lift coefficient gradient. $C_{Lbo(M=0)}$ is the initial buffet onset lift coefficient for $M_B = 0$, S_w is wing area (m^2) and W_{ac} is the aircraft weight (N). All of these coefficient values are available from the BADA database. P is then pressure (Pa). In practice, from [181], the roots of the equation are found as

$$\begin{aligned} X_1 &= 2\sqrt{-Q}\cos\left(\frac{\theta}{3}\right) - \frac{a_1}{3} \\ X_2 &= 2\sqrt{-Q}\cos\left(\frac{\theta}{3} + 120^o\right) - \frac{a_1}{3} \\ X_3 &= 2\sqrt{-Q}\cos\left(\frac{\theta}{3} + 240^o\right) - \frac{a_1}{3} \end{aligned} \tag{B.8}$$

where

$$\theta = \cos^{-1}\left(\frac{R}{\sqrt{-Q^3}}\right) \tag{B.9}$$

$$\begin{aligned} a_1 &= -\frac{C_{Lbo(M=0)}}{k_{cl}} \\ a_2 &= 0, \qquad\qquad a_3 = \frac{\frac{W_{ac}}{S_w}}{0.583\,\mathsf{P}\,k_{cl}} \\ Q &= \frac{(3a_2 - a_1^2)}{9}, \quad R = \frac{(9a_1a_2 - 27a_3 - 2a_1^3)}{54} \end{aligned} \tag{B.10}$$

The solutions of X_1, X_2 and X_3 give the possible values of M_B.

The solution chosen as the buffeting limit is the one with the lowest positive value. In a similar manner as the stall speeds, the coefficient of drag for climbing aircraft is defined as

$$v_{stall} = \begin{cases} C_{D0,TO} + C_{D2,TO}(C_L)^2 & : \text{if } h_{ft}(t) < 400 \text{ ft} \\ C_{D0,IC} + C_{D2,IC}(C_L)^2 & : \text{if } 400 \text{ ft} < h_{ft}(t) < 2000 \text{ ft} \\ C_{D0,CR} + C_{D2,CR}(C_L)^2 & : \text{if } h_{ft}(t) > 2000 \text{ ft} \end{cases}. \tag{B.11}$$

where $C_{D0,TO}$, $C_{D0,IC}$ and $C_{D0,CR}$ are the BADA take-off, initial climb, and cruise parasitic drag coefficients. $C_{D2,TO}$, $C_{D2,IC}$, $C_{D2,CR}$ are the induced drag coefficients for the same three flight phases. For descending aircraft the coefficient of drag is defined as

$$C_D = \begin{cases} C_{D0,LDG} + C_{D0,\Delta LDG} + C_{D2,LDG}(C_L)^2 & : \text{if } \begin{cases} h_{ft}(t) < 3000 \text{ ft} \\ \wedge \\ v_{CAS}(t) < v_{min,APP} + 10 \text{ kts} \end{cases} \\ C_{D0,APP} + C_{D2,APP}(C_L)^2 & : \text{if } \begin{cases} h_{ft}(t) < 3000 \text{ ft} \\ \wedge \\ v_{CAS}(t) > v_{min,APP} + 10 \text{ kts} \end{cases} \\ C_{D0,APP} + C_{D2,APP}(C_L)^2 & : \text{if } \begin{cases} 3000 \text{ ft } < h_{ft}(t) < 8000 \text{ ft} \\ \wedge \\ v_{CAS}(t) < v_{min,CR} + 10 \text{ kts} \end{cases} \\ C_{D0,CR} + C_{D2,CR}(C_L)^2 & : \text{if } \begin{cases} 3000 \text{ ft } < h_{ft}(t) < 8000\text{ft} \\ \wedge \\ v_{CAS}(t) > v_{min,CR} + 10 \text{ kts} \end{cases} \\ C_{D0,CR} + C_{D2,CR}(C_L)^2 & : \text{if } \quad h_{ft}(t) > 8000 \text{ ft} \end{cases} \tag{B.12}$$

where $C_{D0,LDG}$, $C_{D0,APP}$ and $C_{D0,CR}$ are the BADA landing, approach and cruise parasitic drag coefficients. $C_{D2,LDG}$, $C_{D2,APP}$, $C_{D2,CR}$ are the induced drag coefficients for the same three flight phases. $C_{D0,\Delta LDG}$ is the landing gear parasite drag coefficient. Drag (N) is then calculated as

$$D = \frac{C_D \cdot \rho \cdot (1.9438 \cdot v_{t,kts})^2 \cdot S_w}{2} \tag{B.13}$$

where ρ is the air density (kg/m^3) and S_w (m^2) is the wing area.

For the calculation of maximum thrust T_{Max} (N)

$$T_{Max} = \begin{cases} C_{Tc,1}\left(1 - \frac{\mathsf{H_p}}{C_{Tc,2}} + C_{Tc,3}{\mathsf{H_p}}^2\right) & \text{: if Engine type = Jet} \\ \frac{C_{Tc,1}}{v_{t,kts}}\left(1 - \frac{\mathsf{H_p}}{C_{Tc,2}}\right) + C_{Tc,3} & \text{: if Engine type = Turboprop} \\ C_{Tc,1}\left(1 - \frac{\mathsf{H_p}}{C_{Tc,2}}\right) + \frac{C_{Tc,3}}{v_{t,kts}} & \text{: if Engine type = Piston} \end{cases} \tag{B.14}$$

where $\mathsf{H_p}$ is the geopotential pressure altitude in feet, $C_{Tc,1}$, $C_{Tc,2}$ and $C_{Tc,3}$ are the first, second and third BADA thrust coefficients that define the maximum thrust polynomial and $v_{t,kts}$ is the true airspeed in knots. In cruise, the maximum cruise thrust $T_{Cruise,Max}$ was calculated to be 95% of T_{Max}. The nominal fuel flow $\dot{W}_{f_{nom}}$ (kg/s) is then calculated by multiplying the thrust specific fuel consumption η (kg/(min-kN)) by the thrust T (N)

$$\dot{W}_{f_{nom}} = \frac{\eta T}{60000} \tag{B.15}$$

where η (kg/(min-kN)) is defined as

$$\eta = \begin{cases} C_{f1}\left(1 + \frac{v_{t,kts}}{C_{f2}}\right) & \text{: if Engine type = Jet} \\ C_{cf1}\left(1 - \frac{v_{t,kts}}{C_{f2}}\right)\left(\frac{v_{t,kts}}{1000}\right) & \text{: if Engine type = Turboprop} \end{cases} \tag{B.16}$$

and where C_{f1} and C_{f2} are the 1st and 2nd thrust specific fuel coefficients. The minimum fuel flow (kg/s) can be calculated as

$$\dot{W}_{f_{min}} = \frac{C_{f3}}{60}\left(1 - \frac{\mathsf{H_p}}{C_{f4}}\right) \tag{B.17}$$

where C_{f3} and C_{f4} are the 3rd and 4th thrust specific fuel coefficients.