

An Integrated Approach to Optimised Airport Environmental Management

Heblij, Sander

DOI

[10.4233/uuid:820d9e6b-4ac4-4283-bbcf-cc090d17fa2c](https://doi.org/10.4233/uuid:820d9e6b-4ac4-4283-bbcf-cc090d17fa2c)

Publication date

2016

Document Version

Final published version

Citation (APA)

Heblij, S. (2016). *An Integrated Approach to Optimised Airport Environmental Management*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:820d9e6b-4ac4-4283-bbcf-cc090d17fa2c>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

AN INTEGRATED APPROACH TO OPTIMISED AIRPORT ENVIRONMENTAL MANAGEMENT



Sander Heblij

AN INTEGRATED APPROACH TO OPTIMISED AIRPORT ENVIRONMENTAL MANAGEMENT

PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus Prof. ir. K.Ch.A.M. Luyben,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen
op maandag 19 december 2016 om 15:00 uur
door

Sander Johannes HEBLIJ

ingenieur Luchtvaart en Ruimtevaart
geboren te Apeldoorn.

Dit proefschrift is goedgekeurd door de:

promotor: Prof.dr. R. Curran en
copromotor: Dr.ir. H.G. Visser

Samenstelling promotiecommissie:

Rector Magnificus	voorzitter
Prof.dr. R. Curran	Technische Universiteit Delft
Dr.ir. H.G. Visser	Technische Universiteit Delft

onafhankelijke leden:

Prof.Dr.-Ing E. Stumpf	RWTH Aachen, Duitsland
Prof. J.-P. Clarke	Georgia Inst. of Technology, VS
Prof.dr. G.P. van Wee	Technische Universiteit Delft
Prof.dr. D.G. Simons	Technische Universiteit Delft
Prof.dr.ir. M. Mulder	Technische Universiteit Delft

overige leden:

Ir. R.W.A. Vercammen	Nederlands Lucht- en Ruimtevaartcentrum
----------------------	-----------------------------------------

This research was funded by NLR and TU Delft.

Copyright © 2016 by S.J. Heblj

All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilised in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without written permission of the author.

Cover design by R. Nonnekens. Printed in the Netherlands by *W.C. den Ouden*.

ISBN-978-90-826306-0-2

Contents

Summary	vii
Samenvatting (Summary in Dutch)	xi
Nomenclature	xvii
1 Introduction	1
1.1 Airport environmental impact	1
1.2 Management of the airport environmental impact	2
1.3 An integrated support system for improving airport environmental management	4
1.4 Research objective and approach	5
1.5 Limitations and scope	6
1.6 Outline of this thesis	7
2 Airport Environmental Impact Management	9
2.1 Aircraft noise	9
2.2 Aircraft gaseous emissions	16
2.3 Third party risk	19
2.4 Other airport environmental issues	21
2.5 Impact reduction based on optimisation	23
2.6 The integrated approach	28
2.7 Discussion	29
3 Advanced noise abatement departure procedures	31
3.1 The concept of custom optimised departure profiles	31
3.2 Aircraft trajectory optimisation model	36
3.3 Numerical examples of departure procedures	43
3.4 Estimating the runway capacity effects	58
3.5 Discussion	66

4	Reducing environmental impact as objective in arrival management	69
4.1	The concept of an environment aware arrival manager	69
4.2	AMAN optimisation model	71
4.3	Numerical examples of an environment aware arrival management	80
4.4	Cumulative exposure and fairness	83
4.5	Discussion	85
5	Environmental impact allocation in airport strategic planning	87
5.1	The concept of strategic environmental impact allocation	88
5.2	Runway allocation optimisation model	89
5.3	Numerical examples of strategic environmental impact allocation .	101
5.4	Discussion	104
6	Multi-level integrated environmental management	107
6.1	The concept of the multi-level integrated system	107
6.2	The multi-level optimisation model	111
6.3	Numerical examples of the multi-level integrated system	117
6.4	Discussion	125
7	Conclusions and recommendations	129
7.1	Conclusions	129
7.2	Limitations and Recommendations	133
A	Aircraft navigation	137
A.1	Traditional aircraft navigation	137
A.2	Area Navigation (RNAV)	138
B	Additional information on the custom optimised departure profiles	141
B.1	Optimal control	141
B.2	Collocation and nonlinear programming	142
B.3	Solving NLP Problems	143
B.4	Aerodynamic modelling	143
B.5	Influence of wind fields	145
B.6	Overview of optimised profiles	146
B.7	Model verification	146
C	Linear Programming	149
C.1	General introduction	149
C.2	Integer Programming	150
C.3	Solving LP Problems	150

D	Additional information on the environment aware arrival manager	153
D.1	Transit times	153
D.2	CPS sequence	153
D.3	Route selection	155
D.4	Earliest landing opportunity	155
D.5	Separation	156
D.6	Model verification	157
E	Additional information on strategic environmental impact allocation	159
E.1	Identification of situations	159
E.2	Feasible runway allocation options	160
E.3	Consequences of the options	162
E.4	Model verification	162
F	Additional information on multi-level integrated environmental management	163
F.1	Operational allocation model	163
F.2	Traffic scenario	164
F.3	Tests with alternative configurations	165
F.4	Model verification	166
	Bibliography	167
	Acknowledgements	177
	List of Publications	179
	About the author	181

Summary

Airports around the world continue to face issues related to the environmental impact of aviation. Mitigation measures have therefore been implemented at many of these airports. Attaining an optimal combination of mitigation measures can be a difficult process for several reasons. For example, reduction measures targeted at one particular type of impact may not reduce, or even increase other types of impact. A typical example is a noise reduction measure that would increase pollutant emissions. Furthermore, measures can be active at different levels of aggregation. This can be observed when, for example, comparing land use planning activities with modifying the trajectory of a specific flight. Both measures can, however, only be fully effective if attuned to each other. Finally, mitigation measures are often generic in nature or standardised. By definition, this means that they are not fully optimised with respect to the local population distribution and the specific aircraft capabilities.

It is expected that some of the current inefficiencies of mitigation measures can be eliminated by using a process that is based on three main principles: (i) to use mathematical optimisation in order to select the best mitigation options, (ii) to evaluate multiple performance areas simultaneously, and (iii) to evaluate multiple mitigation options at multiple levels of aggregation simultaneously. While each of these principles has already been applied before, the question remains whether these three principles can also be applied simultaneously, by implementing these principles as functionality in a decision support system. It is also questioned whether such a system would indeed lead to additional benefits with respect to airport environmental impact.

These questions resulted in the objective to develop a prototype of an integrated support system for airport environmental management, thereby exploring both the feasibility as well as the benefits of such a system. The approach that has been followed to develop this system is to break down the design problem along different levels of aggregation. Three different levels are defined, and for each of these three levels, this thesis first presents a seemingly independent support sys-

tem for environmental impact mitigation, employing optimisation techniques in each instance. Ultimately, this is followed by the presentation of a fourth support system, which is geared to address the different levels simultaneously.

The first and lowest level of aggregation is defined as the trajectory level. The system presented at this level is based on an existing aircraft trajectory optimisation tool that has been adapted towards a concept called custom optimised departure profiles. Under this concept, each departing aircraft would execute a departure profile that has been optimised for a specific flight, while still using predefined and published departure routes. This concept is presented as an alternative to full trajectory optimisation, involving optimisation of both the vertical profile and the route. It is expected that restricting the optimisation to the vertical profile will reduce the potential environmental benefits. However, the concept does not suffer from a considerable increase in airspace complexity.

The profile optimisation results show that the use of optimised vertical profiles leads to minor reductions in fuel burn and impressive reductions with respect to the applied indicator for community noise when compared to two different standardised departure procedures. The optimised trajectories feature a speed profile that differs from the current situation, and this observation spurred a study into the effects of the optimised profiles on runway capacity. The results of this capacity study show that using the optimised departure profiles or mixing traditional and optimised departures procedures does not lead to a reduction in runway capacity when compared to the baseline situation.

The intermediate level of aggregation is called the operational level and the support system developed at this level is based on the concept of an *environment aware arrival manager*. An arrival manager is typically a decision support tool for air traffic controllers used to assist in creating an efficient flow of aircraft towards the runway. This noise-aware version has been adapted to also assist in obtaining a particular noise allocation strategy. The desired noise allocation result is achieved through the selection of fixed arrival routing options.

The resulting tool has been used to investigate the trade-off between noise and delay in different traffic demand situations, ranging from busy to quiet. The results showed that in some cases the employed noise performance indicator can be improved without any sacrifice in efficiency. In most cases, however, delay increases if the relative importance of the noise objective is increased. For low demand situations, the resulting delay can even exceed the delay in a non-managed (i.e. first come, first serve) situation. For the busy, high demand situations the trade-off proves to be more favourable.

The highest aggregation level regarded is defined the strategic level. The system developed at this level is able to assign all annual flights to the runways of an airport, while respecting the feasibility of the applied runway configurations in terms of required capacity and wind limits. The objective that is used in this model is

to minimise a particular combination of runway delay, community noise exposure and third party risk. For the ultimate trade-off between the three objectives, the system user is guided by a graphical interactive optimisation procedure. The final result is a balanced runway usage scheme.

The strategic level optimiser shows significantly improved results with respect to the reference situation. The cumulative overall risk is reduced by about 30% and a similar reduction is observed for the performance indicator estimating annoyance due to aircraft noise. However, it should be acknowledged that the comparison between the optimised and reference situation may not be completely fair, because of some assumptions and simplifications in the design of the optimisation model. Still, runway allocation optimisation targeted at minimising environmental impact yields promising results at the strategic level.

A fourth and final multi-level model is developed, by combining selected elements of the previously presented systems, as well as adding new components. The resulting implementation is in fact a decision support system at the intermediate (operational) level that interacts with the adjacent levels. The operational component of the system first of all optimises environmental impact and the corresponding geographical allocation by selecting pre-processed optimal departure profiles for each specific flight. Simultaneously, it also influences aircraft routing and runway usage, by selecting airport operational modes.

Single-objective optimisation runs have been performed first, aimed at the minimisation of either fuel burn or one of the three noise-related performance criteria. The results show that all of the optimised solutions outperform the reference solution on any of the four regarded performance criteria. This observation means that improvements are feasible in all four performance areas simultaneously. Furthermore, the benefit of the multi-level optimisation has been shown by comparing the multi-level results in terms of fuel burn and noise response to results obtained using the trajectory level support system only. Finally, the results revealed that minimization of cumulative, multi-event noise exposure is not automatically achieved by combining different flight trajectories that have been optimised for single-event noise exposure.

Multi-objective optimisation has been performed as well, to explore the trade-off between the reduction of fuel burn and the reduction of the noise response. The results yielded the observation that it was possible to obtain the majority of the potential fuel burn and resulting cost reduction, without simultaneously sacrificing most of the potential noise annoyance reduction.

It was expected that three principles, implemented as capabilities in an integrated support system, could help improve the airport environmental management process. The first capability, related to the use of optimisation in order to maximise the efficiency of mitigation measures, has been demonstrated numerous times in the literature, as well as in this thesis. These applications clearly illustrate the benefits of using optimisation through the impact reductions that can be achieved.

The second capability, evaluating multiple environmental performance areas simultaneously, is occasionally promoted in the literature and also demonstrated in this thesis. It can identify negative side-effects, and when combined with optimisation functionality, the trade-off between conflicting objectives can be explored. Results show that these trade-offs can be favourable, meaning that substantial impact reductions can be achieved in one performance area, at the cost of only a marginal increase in the conflicting performance area.

The third and final capability concerns the use of an integrated approach in the sense of considering multiple mitigation options at different levels of aggregation simultaneously, as was shown using the multi-level support system only. Synergy benefits were demonstrated using a powerful combination of optimisation of departure profiles and runway and routing configurations. This resulted in impact reductions that cannot be achieved by these measures in isolation, or even through combination of these measures when not coordinated properly. It should be realised, however, that all three of the desired capabilities play an important role in the example of the multi-level system. It is especially the combination of these three capabilities integrated into a single support system that can contribute to improving the airport environmental management process.

For all of the developed systems, multiple opportunities for further development have been identified. First of all, not all systems consider both arriving and departing traffic, and most could be extended with more and/or more sophisticated environmental impact models and accompanying dose-response models. Several other aspect, like safety, passenger comfort, and non-standard atmospheric conditions have not been regarded, but are certainly relevant, for some of the developed systems.

It is also recommended to further explore the interaction between the strategic and the operational level. Instead of the approach that has been used, where the strategic level model only generates up-to-date sensitivity coefficients to be used by the operational level model, the recommendation is to first generate a high-level optimisation result at the strategic level. Consequently, the operational support systems should actively guide the realisation of the strategic targets.

Generally speaking, it can be stated that the systems in this thesis have been developed for research purposes, without the technological readiness that would be required for real-world application. However, given the results, it is recommended to further develop these capabilities and to work on the implementation issues, and ultimately to make use of these capabilities. This should not necessarily be at every airport, but especially at airports where the efforts towards environmental impact mitigation are considered significant and complex.

Samenvatting

Over de hele wereld worden luchthavens structureel geconfronteerd met de gevolgen van de milieubelasting van de luchtvaart. Daarom zijn op veel van deze luchthavens mitigerende maatregelen geïmplementeerd. Het verkrijgen van een optimale combinatie van de mitigerende maatregelen kan om verschillende redenen een moeizaam proces zijn. Zo kan een maatregel die gericht is op het reduceren van een bepaald type milieubelasting geen of zelfs een tegengesteld effect hebben indien wordt gekeken naar een ander type milieubelasting. Een typisch voorbeeld hiervan is een geluidsreducerende maatregel die leidt tot meer gasemissies. Verder kunnen maatregelen ingrijpen op verschillende aggregatieniveaus. Vergelijk hierbij bijvoorbeeld activiteiten op het gebied van ruimtelijke ordening met het aanpassen van het vluchtpad van een bepaalde vlucht. Beide maatregelen kunnen echter alleen volledig effectief zijn indien ze op elkaar zijn afgestemd. Ten slotte geldt ook nog dat mitigerende maatregelen vaak generiek of gestandaardiseerd zijn. Dit betekent dat ze per definitie niet volledig geoptimaliseerd zijn naar de lokale verdeling van de bevolking en de mogelijkheden van een bepaald vliegtuig.

De verwachting is een deel van de huidige inefficiënties rondom mitigerende maatregelen kan worden weggenomen indien een proces wordt toegepast dat is gebaseerd op drie hoofdprincipes: (i) het gebruik van wiskundige optimalisatie om de beste maatregelen te selecteren, (ii) het simultaan evalueren van de verschillende type milieubelastingen en (iii) het simultaan evalueren van verschillende maatregelen op verschillende aggregatieniveaus. Hoewel deze principes allemaal eerder al zijn toegepast, is het de vraag of ze ook gelijktijdig kunnen worden toegepast, door deze principes als functionaliteit op te nemen in een beslissingsondersteunend systeem. Vervolgens is het de vraag of een dergelijk systeem ook zou leiden tot additionele voordelen op het gebied de milieubelasting van luchthavens.

Deze vraagstelling heeft geleid tot de doelstelling om een prototype van een integraal beslissingsondersteunend systeem voor het beheer van de milieubelasting van luchthavens te ontwikkelen, om daarmee zowel de haalbaarheid als de voordelen van een dergelijk systeem te verkennen. Wat betreft de aanpak om dit

systeem te ontwikkelen, is ervoor gekozen om het ontwerp uit te splitsen langs de verschillende aggregatieniveaus. Er zijn hierbij eerst drie niveaus gedefinieerd en voor elk van deze drie niveaus wordt eerst een ogenschijnlijk onafhankelijk systeem gepresenteerd, wat met behulp van optimalisatietechnieken ondersteuning geeft aan het mitigeren van de milieubelasting. Uiteindelijk volgt ook nog een vierde systeem, welke in staat om de verschillende niveaus simultaan te adresseren.

Het eerste en laagste aggregatieniveau is gedefinieerd als het vliegbaanniveau. Het gepresenteerde systeem is gebaseerd op een bestaand softwarepakket voor vliegbaanoptimalisatie, welke is aangepast om te voldoen aan een concept genaamd individueel geoptimaliseerde startprofielen. Bij dit concept volgt elk vliegtuig tijdens de start een verticaal vliegprofiel, welke geoptimaliseerd is voor die specifieke vlucht, maar waarbij nog wel langs de vaste en gepubliceerde startroutes wordt gevlogen. Dit concept wordt gepresenteerd als een alternatief voor volledige baanoptimalisatie, waarbij zowel het verticale profiel als ook het grondpad worden geoptimaliseerd. De verwachting is dat het enkel optimaliseren van het verticale profiel zal leiden tot minder milieuwinst, maar tegelijkertijd leidt het gepresenteerde concept ook niet tot een toename van de complexiteit van het luchtruim.

Uit de resultaten van de profieloptimalisatie blijkt dat het gebruik van de optimaliseerde profielen leidt tot een kleine afname in brandstofverbruik en een indrukwekkende afname van de gebruikte indicator voor de geluidsbelasting, indien deze profielen worden vergeleken met twee standaard startprocedures. De geoptimaliseerde profielen wijken wat betreft snelheidsverloop wel af van de huidige praktijk en deze observatie heeft geleid tot een extra studie naar het effect van deze profielen op de startbaan capaciteit. Uit de resultaten blijkt dat het gebruik van de geoptimaliseerde profielen, al dan niet in combinatie met de traditionele profielen, in vergelijking tot de referentiesituatie niet tot een afname van de baan capaciteit leidt.

Het middelste aggregatieniveau wordt het operationele niveau genoemd en het ontwikkelde systeem is gebaseerd op een *geluidsbewuste arrivalmanager*. Een arrivalmanager is een ondersteunend systeem dat normaal gesproken verkeersleiders helpt om naderende vliegtuig zo efficiënt mogelijk achter elkaar op te lijnen voor de landing. De geluidsbewuste versie is zo aangepast dat deze ook kan helpen om een bepaalde verdeling van het geluid over de omgeving te verkrijgen. Deze verdeling wordt hierbij bereikt door een keuze te maken uit verschillende vaste naderingsroutes.

Het systeem is vervolgens ingezet om de afweging tussen geluid en vertraging te onderzoeken, bij verschillende verkeersaanbod, variërend van hoog tot laag. De resultaten laten zien dat in bepaalde gevallen de prestatie-indicator voor geluid kan worden verbeterd, zonder op efficiëntie in te leveren. Echter, in de meeste gevallen zal de vertraging toenemen naarmate het relatieve belang van de doelstel-

ling voor geluid wordt verhoogd. Voor situaties met weinig verkeersaanbod kan de optredende vertraging zelfs hoger worden dan in de situatie waarbij vliegtuigen worden afgehandeld op basis van volgorde van binnenkomst. Voor situaties met veel verkeersaanbod is de afweging duidelijk gunstiger.

Het hoogst beschouwde aggregatieniveau is gedefinieerd als het strategisch niveau. Het ontwikkelde systeem kan alle vluchten van een jaar toewijzen aan de start- en landingsbanen, waarbij de geschiktheid van de ingezette baancombinaties wat betreft windlimieten en benodigde capaciteit wordt gerespecteerd. De doelstelling van dit model is om een bepaalde combinatie van de optredende vertraging, de geluidsbelasting en het externe veiligheidsrisico te minimaliseren. Voor de onderlinge afweging tussen deze drie doelen wordt de gebruiker ondersteund met een interactieve, grafische optimalisatieprocedure. Dit resulteert uiteindelijk in een weloverwogen baangebruiksplan.

De resultaten van de optimalisatie laten een aanmerkelijke verbetering zien wanneer deze worden vergeleken met de referentiesituatie. Het cumulatieve veiligheidsrisico is met ongeveer 30% gereduceerd en een vergelijkbare reductie is zichtbaar bij de prestatie-indicator die dient om de hinder als gevolg van de geluidsbelasting in te schatten. Echter, hierbij moet worden opgemerkt dat de vergelijking tussen de geoptimaliseerde en referentiesituatie niet volledig zuiver is, als gevolg van aannames en vereenvoudigingen in het ontwerp van het optimalisatiemodel. Desondanks blijven de resultaten van een op strategisch niveau geoptimaliseerde baantoewijzing, gericht op het minimaliseren van de milieubelasting, veelbelovend te noemen.

Het vierde en laatste model adresseert meerdere aggregatieniveaus. Dit model is ontwikkeld door het combineren van geselecteerde elementen uit de eerder gepresenteerde systemen, alsmede het toevoegen van nieuwe elementen. Het resultaat hiervan is een beslissingsondersteunend systeem dat werkt op het middelste (operationele) aggregatieniveau en daarbij beide naastgelegen niveaus betreft. De operationele component optimaliseert de milieubelasting en de bijbehorende ruimtelijke allocatie daarvan door voor elke vlucht een vooraf geprepareerd optimaal startprofielen te selecteren. Tegelijkertijd wordt ook de routing van de vliegtuigen en het baangebruik beïnvloed door te kiezen uit verschillende operationele modi voor de luchthaven.

Allereerst is geoptimaliseerd met één doelstelling, waarbij geminimaliseerd is naar brandstofverbruik of één van de drie geluidgerelateerde indicatoren. Uit de resultaten blijkt dat alle geoptimaliseerde oplossingen op alle vier de aspecten beter presteren dan de referentieoplossing. Dit betekent dat een reductie op alle vier prestatie-indicatoren tegelijkertijd mogelijk is. Het voordeel van optimaliseren op meerdere aggregatieniveaus is gedemonstreerd door de resultaten in termen van brandstofverbruik en geluidsbelasting te vergelijken met de resultaten die eerder zijn behaald met enkel het systeem op vliegbaanniveau. Ten slotte blijkt uit de resultaten nog dat het minimaliseren naar de cumulatieve geluidsbelasting van

verschillende vluchten samen, niet automatisch wordt bereikt door het combineren van vluchten die zelf zijn geoptimaliseerd voor minimale geluidsniveaus van elke vlucht afzonderlijk.

Ook is er geoptimaliseerd met meerdere doelstellingen om zo het spanningsveld tussen brandstofverbruik en geluid in kaart te brengen. Uit deze resultaten blijkt het mogelijk is om het merendeel van de potentiële reductie in brandstofverbruik en bijbehorende kosten te behalen, zonder dat tegelijkertijd daarvoor ook het grootste deel van de potentiële reductie in geluidshinder moet worden ingeleverd.

De verwachting was dat drie hoofdprincipes, als functionaliteit opgenomen in een beslissingsondersteunend systeem, zouden kunnen helpen om het beheer van de milieubelasting van luchthavens te verbeteren. Het eerste principe, namelijk het gebruik van optimalisatie om de efficiëntie van mitigerende maatregelen te maximaliseren, is al vaak in de literatuur, maar ook meerdere keren in deze dissertatie gedemonstreerd. Deze toepassingen illustreren duidelijk de voordelen van optimalisatie op basis van de reducties in milieubelasting die kunnen worden bereikt.

Het tweede principe, het gelijktijdig evalueren van de verschillende type milieubelastingen, wordt ook soms aangeraden in de literatuur en is ook toegepast in deze dissertatie. Hiermee kunnen ongewenste neveneffecten worden geïdentificeerd en indien gecombineerd met optimalisatiemogelijkheden, kan ook het spanningsveld tussen tegenstrijdige doelstellingen worden verkend. Uit de resultaten blijkt dat de afwegingen gunstig kunnen zijn, wat leidt tot een relatief grote reductie van het ene type milieubelasting, op kosten van slechts een kleine toename van een ander type milieubelasting.

Het derde en laatste principe betreft het gebruik van een integrale aanpak in de zin van het simultaan evalueren van verschillende maatregelen op verschillende aggregatieniveaus, wat alleen het vierde en laatste ontwikkelde systeem kan. Met dit systeem zijn synergievoordelen gedemonstreerd op basis van een sterke combinatie van optimalisatie van startprofielen en de toegepaste baan- en route configuraties. Dit heeft geresulteerd in reducties die op basis van de losse maatregelen niet gehaald kunnen worden, ook niet indien deze maatregelen gelijktijdig maar ongecoördineerd worden toegepast. Hierbij moet echter worden bedacht dat alle drie de principes een belangrijke rol spelen bij de resultaten van dit systeem. Het is vooral de combinatie van deze principes, geïmplementeerd in een beslissingsondersteunend systeem, die kan bijdragen aan het verbeteren van het beheer van de milieubelasting van luchthavens.

Voor alle ontwikkelde systemen geldt dat er diverse mogelijkheden voor verdere ontwikkeling zijn geïdentificeerd. Allereerst geldt dat niet alle systemen zowel aankomend als vertrekkend verkeer beschouwen en ook geldt meestal dat ze zouden kunnen worden uitgebreid met extra of meer geavanceerde milieubelastingsmodellen en bijbehorende dosis-respons curves. Diverse andere aspecten, zoals veiligheid, passagierscomfort en niet-standaard atmosferische condities zijn niet beschouwd, maar zeker wel relevant voor een deel van de ontwikkelde systemen.

Verder wordt aanbevolen de interactie tussen het strategisch en operationele niveau verder te onderzoeken. In plaats van de gebruikte aanpak, waarbij het strategisch niveau enkel de actuele gevoeligheidscoëfficiënten bepaalt ten behoeve van het operationele niveau, is de aanbeveling om eerst op het strategisch niveau een optimum te genereren. Vervolgens zou dan het operationele systeem zo moeten sturen dat de strategische doelen bereikt worden.

Algemeen gezien geldt dat de systemen voor onderzoeksdoeleinden zijn ontwikkeld en niet de technologische volwassenheid hebben die vereist zou worden om ze in de dagelijkse praktijk te gebruiken. Gezien de resultaten wordt het echter aanbevolen verder te werken aan de implementatie van de drie hoofdprincipes. Daarnaast geldt de aanbeveling om deze principes uiteindelijk ook in te zetten, niet noodzakelijkerwijs op alle luchthaven, maar vooral voor die luchthavens waarbij de inspanningen gericht op het reduceren van de milieubelasting omvangrijk en complex zijn.

Nomenclature

List of Acronyms

ADF	Automatic Direction Finder
ANSP	Air Navigation Service Provider
ARTCC	Air Route Traffic Control Center
ATC	Air Traffic Control
ATS	Air Traffic Services
BIP	Binary Integer Programming
BT	Business Trajectory
CDA	Continuous Descent Approach
CPS	Constrained Position Shifting
CTAS	Center-TRACON Automation System
DME	Distance Measuring Equipment
DSS	Decision Support System
ECAC	European Civil Aviation Conference
EDP	Expedite Departure Planning
EFR	Environmental Fiscal Reforms
EPNL	Effective Perceived Noise Level
FAA	Federal Aviation Administration

FAST Final Approach Spacing Tool
FCFS First Come First Serve
FICAN Federal Interagency Committee on Aircraft Noise
FMS Flight Management System
GUI Graphical User Interface
ICAO International Civil Aviation Organization
INM Integrated Noise Model
IP Integer Programming
LP Linear Programming
LTO Landing and Take-off
LVNL Air Traffic Control the Netherlands
MIP Mixed Integer Programming
NADP Noise Abatement Departure Procedure
NAP Noise Avoidance Planner
NARSIM NLR ATC Research Simulator
NASA National Aeronautics and Space Administration
NATS National Air Traffic Services
NDB Non-directional Beacon
NEI Noise Exposure Index
NLP Nonlinear Programming
NLR Netherlands Aerospace Centre
NPD Noise Power Distance
PSZ Public Safety Zone
QC Quota Count
QP Quadratic Programming
RNAV Area Navigation
RPK Revenue Passenger Kilometres

SEL Sound Exposure Level

SQP Sequential Quadratic Programming

TMA Terminal Area

TPR Third Party Risk

TRACON Terminal Radar Approach Control

UHC Unburned Hydrocarbons

VOR VHF Omnidirectional Range

WHO World Health Organization

Chapter 1

Introduction

The air transport market has always been characterised by very high growth rates. Measured in tonne kilometres carried, air transport grew 14-15% per year on average during the 1950s and 1960s^[48]. Although the growth rate decreased as the market matured, passenger kilometres still increased by a factor 4.6 between 1970 and 1995 which is still more than 6% on average^[68]. And even though the European and North American markets are now quite mature, and the growth has been impacted by the financial crisis, still a factor 2.3 in Revenue Passenger Kilometres (RPK) is expected for Europe for the next twenty years^[24].

The high growth rates in the air traffic market are typically associated with the development of airport environmental problems. A second important factor in the development of these problems and subsequent complaints was the rapid growth in population combined with increasing levels of urbanisation. This made cities expand beyond their original boundaries, generally also encroaching upon the land around the airports. Both effects resulted in a situation where an increasing number of people were exposed to increasing levels of environmental impact. Airport environmental impact management arose from the recognition that something needed to be done.

This chapter first discusses the importance of airport environmental management and its current inefficiencies. This is followed by an analysis on how an integrated support system can improve this process and how such a system could look like. The final part of this chapter presents the outline of this thesis, which focuses on the development of such a support system for airport environmental management.

1.1 Airport environmental impact

When regarding transportation noise in a general sense, community exposure is widespread. The World Health Organization (WHO) estimates that about half of

the 500 million European Union citizens live in areas that do not ensure acoustical comfort to residents and more than 30% are exposed to night-time noise levels that are disturbing to sleep^[32;95].

The WHO links exposure to noise to several adverse health effects. These effects include noise-induced hearing impairment and the development of hypertension. Other effects related to human well-being are interference with speech and loss of performance in cognitive task. Finally, social and behavioural effects are mentioned, as well as the development of annoyance^[32].

Noise exposure does not only influence human well-being, but also results in an economic impact. One economic component is the need for additional investments, such as sound proofing of buildings. The total cost for such a program for Amsterdam Airport is currently estimated at 600 million euro^[22] and is paid for by airlines landing at the airport^[104]. A second example is loss of property value due to noise impact, often estimated using hedonic pricing^[103]. The results differ from study to study, but a depreciation of 0.81% per dB is found as an average value^[96]. A final economic impact can be identified when noise exposure limits further economic development. Noise generated by aircraft operations is considered one of the major constraints in further air traffic growth, especially for large airports in the developed world^[12;25].

Air pollutant emissions by transportation vehicles are not of less importance. Apart from the contribution to global emissions and the associated climate change, vehicles equipped with combustion engines also have a negative effect on local air quality. Air pollution has an important negative impact on the health of individuals. According to the WHO, there is increasing evidence for adverse effects of air pollution not only on the respiratory system, but also on the cardiovascular system^[79]. Similarly as for noise, purely economical effects can be identified as well, such as damage to crop growth, particularly via the formation of low-level ozone^[51].

Of course, aviation is not the sole factor in both environmental pollution problems, but in areas near airports it is an important factor. This is mainly caused by the combination of a concentration of air traffic in that area and the relatively low altitude of flight movements. On top of that, busy airports also attract a lot of road traffic, to get the passengers and freight to and from the airport, further contributing to the noise and emissions that can be associated with airports.

1.2 Management of the airport environmental impact

To reduce the negative side-effects of aviation a range of mitigation measures has been implemented at airports located close to sensitive communities^[50;56]. Some of these measures are in the form of operating restrictions. Others try to influence demand from the airport's point of view, as they favour certain segments of the market or the use of more environmentally friendly aircraft, but do not directly limit demand. A third group of measures considers air traffic demand as a given,

and aims for dealing with all flights such that environmental impact is minimal.

Achieving an optimal management of the mitigation measures is a complex process, first of all because of the many actors^[46]. A wide range of airlines together with the airport operator and the Air Navigation Service Provider Air Navigation Service Provider (ANSP) are the three main parties that are involved around the actual operation of aircraft at a specific airport. However, regulators at national, regional and municipal level are also involved when it comes to environmental problems, just like different community groups. The fact that these actors usually do not have the same objectives makes this problem even more complicated and may require compromises that are less than ideal.

Adding to the the complexity of the impact mitigation process is the fact that the measures are active at different levels of aggregation. Modifying the trajectory of a specific flight to change the noise impact at that particular time and location is a measure that takes place at a far lower level than decisions in the strategic (or political) domain, such as restricting residential development near the airport in land use planning assignments. But for both measures to be fully effective, the lower level mitigation efforts should be in line with the high-level measures and policy, which is not always evident.

A further mismatch in mitigation measures may arise if the different types of environmental impact are not considered simultaneously. For example, the use of preferential runways for reasons of noise, may increase aircraft taxi times, potentially increasing the emission of pollutant gasses. Naturally, a trade-off between noise and emissions is required in the case of this example, but if that trade-off is unclear, making the right decision is practically impossible.

Another potential inefficiency in the mitigation process can be found in the fact that noise abatement operating procedures for aircraft are standardised. Although the reasoning behind standardising procedures is not disputed in today's operational environment, it still means that these procedures are by definition not fully optimised with respect to the local demographic situation and specific aircraft capabilities. However, with the technological advances in the guidance, navigation and automation capabilities of modern aircraft, opportunities appear to arise to shift towards a more site-specific approach^[105], preferably also based on aircraft capabilities.

To summarise, an optimal management of the environmental impact is not always evident. Even if conflicting interests of different stakeholders do not result in sub-optimal compromises, a mismatch in mitigation measures may still arise if measures at the different levels of aggregation, or regarding the different types of environmental impact, are not attuned to each other. Removing some of these inefficiencies is expected to improve the overall result. Especially in combination with a customisation and optimisation approach, enabled by technological developments, there appear to be opportunities to improve the process in order to reduce airport environmental impact.

1.3 An integrated support system for improving airport environmental management

In the ideal situation, all environmental impact mitigation efforts taking place at the different levels of aggregation should be in concert and managed concurrently. When using such a form of integrated environmental management, it can be ensured that all actions taken to minimise the nuisance caused by aircraft environmental impact will be consistent, complement each other, and make use of synergy benefits. At the same time, it helps avoiding that a certain decision (partly) reduces the effectiveness of another measure at a different level of aggregation, or made by a different stakeholder.

There have been efforts toward a more integrated approach in the impact analysis part of airport environmental management by combining multiple environmental impact models into a single system. A recent example is the development of the Aviation Environmental Design Tool^[35;85]. This new tool will replace and at the same time combine five popular environmental models, three for noise and two for emissions, and is likely to obtain regulatory status. The main advantage of using such an integrated analysis model is that the models can use a single set of input data, saving the user time while guaranteeing consistency in the input. However, such tools do not offer optimisation functionality and decision support is limited to presenting information in a convenient way.

There have also been attempts to improve a specific mitigation effort through the application of mathematical optimisation. A typical example is the design of an optimisation framework for calculating the optimal runway use preference list with respect to the noise load around the airport^[82]. While the study demonstrates the power of applying optimisation, it regards a single mitigation measure active at a single level of aggregation and is aimed at noise impact only.

Flight procedure optimisation with respect to airport environmental impact is applied as well. A popular field of study is the design and further optimisation of the Continuous Descent Approach (CDA)^[19;44;91]. The general principle of this procedure is to keep aircraft higher and at reduced thrust while performing the approach. This procedure does result in benefits with respect to both the performance area of noise and emissions, and on top of that, the procedure reduces fuel burn as well. But again, it is based on a single mitigation measure. Besides, this measure is also generic in nature, not necessarily optimal with respect to the local land use situation.

Chapter 2 will introduce more examples of mitigation efforts, structured along three main principles: (i) to use mathematical optimisation in order to select the best mitigation options, (ii) to evaluate multiple performance areas simultaneously, and (iii) to evaluate multiple mitigation options at multiple levels of aggregation simultaneously. However, none of these examples apply all three principles simultaneously and this thesis addresses this gap.

1.4 Research objective and approach

The question remains whether these three principles can indeed be combined into an integrated solution in the form of a support system, and if such a system would contribute to additional benefits with respect to airport environmental impact. Formulated alternatively:

Can an integrated support system improve the airport environmental management process?

Addressing this research question results in the objective to develop a prototype of an integrated support system for airport environmental management, thereby exploring both the feasibility as well as the benefits of such a system.

The approach that has been followed to develop the system is to break down the design problem along different levels of aggregation. As a result, this thesis first presents three different and seemingly independent support systems for different actors, applying one or multiple mitigation measures at a single level of aggregation, while using optimisation techniques. However, these levels will ultimately be addressed simultaneously with the development of a multi-level system.

The three levels of aggregation that have been identified for the three single-level systems are depicted in figure 1.1 and are defined as:

1. The **trajectory level** comprises mitigation measures based on the design or adaptation of the trajectory of a single flight, where each flight is considered as an isolated event. Only the part of the trajectory that is relevant for the environmental impact of a particular airport is considered, resulting in a focus on arrival and departure trajectories.
2. The **operational level** comprises mitigation measures related to the actual day-to-day operation of the airport, especially the provision of air traffic control in the context of environmental performance. As this level can be seen as the level where the actual environmental impact is caused and where the mitigation measures are to be brought into practice, it is readily clear that this level is of paramount importance. Unfortunately, it is also the level where environmental objectives often conflict with other interests, such as safety and efficiency.
3. The **strategic level** comprises mitigation measures with the focus on long-term planning. This includes regulatory issues, such as land use planning and noise compatibility. Other actions at the strategic level involve policy making for airport use and development, as far as this is relevant for the environmental impact.

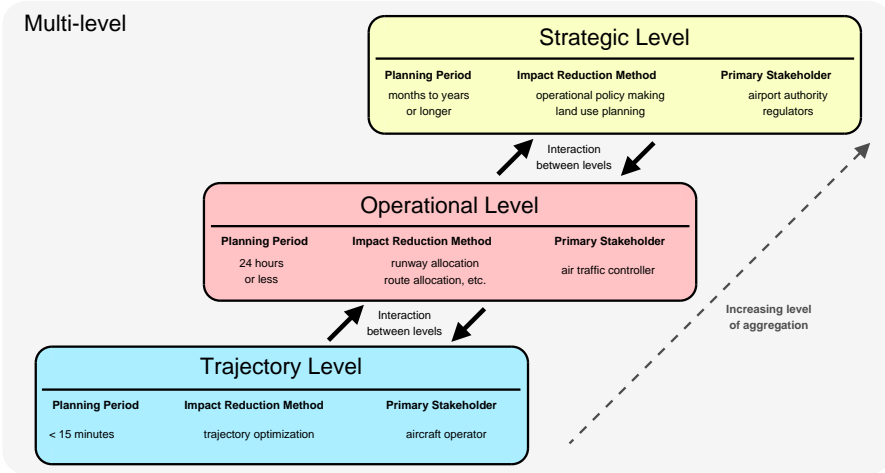


Figure 1.1: *The levels of aggregation for the support systems*

The intended functionality of these systems and related requirements differ for the three different levels. At the same time, the primary actor -the stakeholder that is the intended user of support system- also varies, which is possible because these are independent systems. The trajectory and operational level systems are targeted at airlines and air traffic controllers, respectively. The strategic system is designed to provide decision support for regulators or the airport authority.

Towards the end of the thesis, the three separate support systems -or tools- are combined into a multi-level, hierarchical system. For the sake of this integration, it is analysed how the three levels interact and what the desired hierarchical structure is for this multi-level system. Additionally, the multi-actor problem is analysed to provide a concept that involves all actors, but ultimately provides decision support for a single stakeholder.

1.5 Limitations and scope

The assumptions and limitations that apply to the four developed support systems vary among these different systems, and so does the scope. Therefore the chapters presenting these systems will each provide more detail on the related context, assumptions and limitations. On a higher level however, several areas can be identified that have been regarded to be outside the scope of this project.

Community annoyance that arises with exposure to aircraft noise or transportation noise in general is dependent on both acoustical as well as non-acoustical factors. Examples of non-acoustical variables that may influence the level of annoyance are the predictability and transparency of the noise situation, and the

use of consultation when setting the noise strategy^[53;97]. This thesis, however, focuses on the physical, acoustical side of noise. This means that when reference is made to noise levels, computed sound levels are intended and when referring to annoyance, this is actually an indicator for annoyance, based on an empirically determined dose-response relationship (to be introduced in Chapter 2).

Regarding the definition of the levels of aggregation, a fourth level was identified during this study and classified as the *source level*, positioned one level below the trajectory level. Mitigation measures at this level relate to improving the environmental performance at the design stage of an aircraft type, or making improvements to a particular aircraft during the service life, such as the installation of hush-kits to meet noise regulations, or more present-day, the installation of winglets to reduce fuel burn and global emissions. However, aircraft design and modification were deemed to be outside the scope of this work.

Thirdly, and while it is recognised that the interface with the user of a system is an important part of an operational Decision Support System (DSS), this was not a priority during this project. No literature has been consulted regarding interface design and not all systems actually have graphical user interfaces. Where interfaces have been developed, these have not been tested in consultation with the intended user community of such a system.

Finally, all numerical examples are based on Amsterdam Airport Schiphol in The Netherlands, and all presented tools have only been tested based on the layout and characteristics of this airport. Still, it is believed that the general concept and applied methodologies are applicable to other airports as well, assuming that they face similar environmental challenges and mitigation options.

1.6 Outline of this thesis

Figure 1.2 shows the outline of this thesis.




Figure 1.2: A graphical overview of the structure of this thesis

As shown in this figure, this introduction is followed by a chapter with back-

ground information:

- Chapter 2 provides more detail on airport environmental management and impact mitigation. This chapter also presents examples of important mitigation measures, previous efforts of applying optimisation and examples of applying a more integrated approach

The next four chapters each present a decision support system for integrated environmental management, in line with the levels shown in figure 1.1. Each of these chapters present the corresponding concept, design of the system, and also contain numerical results that have been generated with these systems:

- Chapter 3 presents the trajectory level system. Based on aircraft trajectory optimisation, this support system for airlines generates departure procedures that have been optimised for a particular flight.
- Chapter 4 presents the operational level system. The support system is aimed at air traffic controller actions, helping them in guiding arriving traffic near the airport in a safe and efficient matter in concert with noise exposure considerations.
- Chapter 5 presents the strategic level system. In this case, the support system is more of a management system, designed for helping an airport strategic advisor in drafting an optimal environmental impact allocation policy for the future.
- Chapter 6 presents the multi-level, hierarchical system that is based on the integration of three single-level systems.

The presentation of the four support systems and their results is followed by conclusions, recommendations and several appendices:

- Chapter 7 summarises the conclusions and provides recommendations for future work in this area of research.
- Appendix A gives background information on aircraft navigation and current developments in this area. Especially these developments are important prerequisites for the concepts on which the first two support systems are based.
- Appendix B provides additional material on the optimisation techniques and models that have been used for the trajectory level system. This appendix can be seen as an extension to Chapter 3.
- Appendix C is dedicated to reviewing Linear Programming methods, as the remaining three support systems are based on this particular optimisation technique.
- Finally, Appendices D, E, and F contain additional information on the models presented in Chapter 4, 5, and 6, respectively.

Airport Environmental Impact Management

This chapter presents background material on airport environmental impact management and as such, lays the foundation for the work in the next four chapters. Comprising two parts, the first part of this chapter provides an overview of the important knowledge or performance areas with respect to aviation related environmental impact. The first two areas, noise and emissions, and their importance to society have already been introduced and are described in the first two sections of this chapter. This is completed with an overview on third party risk, an important third field in certain countries, which concerns the risk of population around the airport in becoming a casualty as a result of an aircraft accident. Finally, Section 2.4 gives a brief overview of the remaining performance areas that may be considered part of the airport environmental management process, but are outside the scope of this dissertation.

Readers familiar with the topics in the first part may wish to proceed to the second part of this chapter, starting at Section 2.5. This section looks into several mitigation measures and strategies that can be used in airport environmental management, specifically the mitigation measures, strategies and tools that utilise some form of optimisation to minimise the impact. Finally, this chapter elaborates on previous work based on an integrated approach, as briefly touched upon in the introduction.

2.1 Aircraft noise

Compared to modern commercial aircraft, the first jet-powered airliners were much more noisy. Fortunately, decades of engineering have greatly reduced the noise levels. It should, however, be noted that a part of that reduction may have been the by-product of the drive to increase the efficiency of the engine, thereby

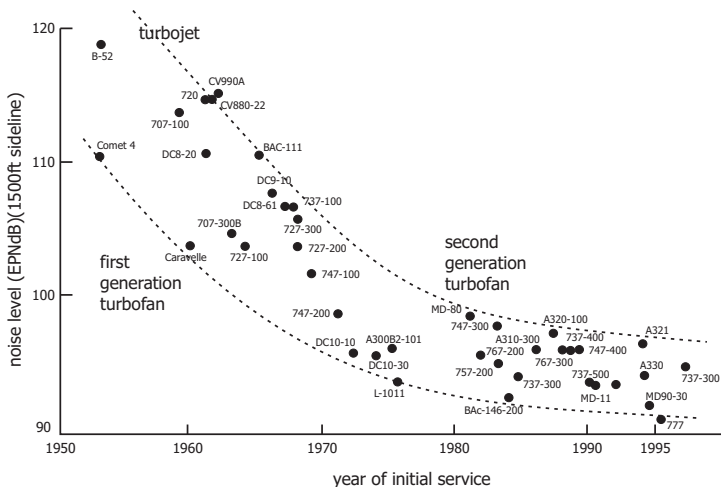


Figure 2.1: Noise levels versus year of aircraft model introduction, normalised for take-off weight^[46;102]

reducing fuel costs. By increasing the bypass ratios to let more air bypass the core of the engine, the average exhaust-jet velocity was reduced. This resulted in an increase in propulsive efficiency as well as a decreasing exhaust noise^[92]. However, a further reduction of aircraft noise was not as easy achievable by further increasing the by-pass ratio, as the exhaust-jet noise became less dominant over fan noise. Advances in fan design and the application of acoustic linings in the fan duct were applied to reduce the fan noise. Overall, the noise reduction over the years is substantial and is shown in figure 2.1.

While the aircraft’s engines are typically the primary sources of noise, they are not the only ones as the airflow over the airframe itself also produces sound. Especially during the approach, the noise generated near the extended landing gear and the high lift devices can be significant. Airframe noise may even be dominant in this phase of flight, also because the engines are generally not running at very high power settings. This means that only reducing the engine noise levels may not be sufficient to further reduce aircraft noise levels in the future.

2.1.1 Aircraft noise regulations

The international standards concerning aircraft noise levels are laid down in International Civil Aviation Organization (ICAO)’s Annex 16, titled *Environmental Protection*^[92]. The different noise standards in this document are organised through chapters, each devoted to a particular type of aircraft, depending on propulsion system, aircraft weight and certification date. Many aircraft operating

from the larger airports are currently certified based on Chapter 3 of the Annex, often referred to as Chapter 3 aircraft. Chapter 2 refers to older aircraft, which are currently banned in large parts of the world. Chapter 4, previously devoted to supersonic aircraft, is the 2006 standard for newly certified airliners.

For a Chapter 3 noise certification, the aircraft must fly both an approach and a departure reference trajectory in specific atmospheric conditions. During the test, noise is measured at several locations to provide three noise values called departure, approach and sideline. Depending on aircraft weight and the number of engines, maximum values are specified for each of the three points. Both the measured noise levels as well as the posed limits are expressed in a time-integrated noise metric, called Effective Perceived Noise Level (EPNL).

Chapter 4, the new standard, is more stringent, but is still based on Chapter 3. The main difference is that the aircraft should comply at all three points, where Chapter 3 has compensation rules. Furthermore, a Chapter 4 aircraft should have a combined margin of at least 2 EPNdB at any two points and a combined margin of at least 10 EPNdB at the three points together^[86].

2.1.2 Airport noise regulations

Noise regulations do not only apply to the aircraft itself as more and more airports have to deal with operational noise regulations as well. These regulations differ widely from airport to airport and can be based on quota, bans and other types of restrictions^[56].

Except for the mentioned Chapter 2 ban, the ICAO noise certification classification system as discussed in the previous section has little impact on the day-to-day operation of the aircraft. Indirectly however, the EPNdB-based certification levels are used to limit community noise exposure around airports. An example is the Quota Count (QC) system as used by several British airports^[15]. The participating airports all have a number of available credits for the night period: the quota. How much credits each operation requires is computed directly from the noise certification levels of the corresponding aircraft.

More often, however, airport noise monitoring activities and limits are not based on the EPNL, but on the family of A-weighted noise metrics. An A-weighting, as depicted in figure 2.2(a) is a correction procedure for sound levels that corrects for the different sensitivities of the human ear with respect to different frequencies. With respect to a single flyover, there are three important A-weighting based metrics, as shown in figure 2.2(b). First of all, the L_A provides the instantaneous A-weighted sound level. The maximum L_A that is reached during this a flyover, or the peak level, is designated the L_{Amax} . Finally, the total acoustical energy of the flyover can be expressed using the metric L_{AE} , or Sound Exposure Level (SEL). The SEL is computed by referencing the total event to a fixed duration, typically one second.

The peak level L_{Amax} is sometimes used in relation to operational restrictions. Heathrow, Stansted and Gatwick for example enforce peak level limits for depart-




Figure 2.2

period of day	Limit dB(A)
06:00 - 07:00	89
07:00 - 23:00	94
23:00 - 23:30	89
23:30 - 06:00	87

Table 2.1: $L_{A,max}$ based limits for departing aircraft at Heathrow, Stansted and Gatwick at 6.5 km from the start of the takeoff roll^[40]

ing aircraft, measured at 6.5 km from the start of the takeoff roll by using several noise monitoring terminals placed at strategic locations. The current limits are provided in table 2.1. A violation of this limits results in a financial penalty of up to £1000 for the responsible airline^[53].

While this example shows that noise levels for a particular movement can be restricted, it should also be recognised that it is more common to limit the total noise exposure of all aircraft movements together within a given period. Cumulative noise exposure is expressed using numerous different metrics, but the basic principle is to either count, sum, average or determine total duration of all (significant) aircraft noise events in a particular period, while sometimes applying penalties for noise-sensitive periods of the day. Both the single-event $L_{A,max}$ and L_{AE} metrics may be used as basis.

From an historical perspective, several countries in the European Union used their own local metric(s) for aircraft noise exposure, sometimes only differing in notation. Today, noise metrics have been harmonised according to the Environ-

mental Noise Directive of 2002^[13]. This directive states that the indicators L_{den} ¹ and L_{night} should be used to express environmental noise, not only in relation to aircraft noise, but transportation and industrial noise in general.

The L_{night} is the average or equivalent level for the night period and can be calculated as follows:

$$L_{night} = 10 \log_{10} \left[\sum_{i=1}^k 10^{\frac{L_{AE,i}}{10}} \right] - 10 \log_{10} \frac{T}{\tau} \quad (2.1)$$

where T is the duration of the night (28800 seconds) and τ the reference period that was used to calculate the L_{AE} levels, typically one second. The level for a 24h period, L_{den} is also an average level. However, it applies a 5 and 10 dB(A) penalty for evening and night events, respectively. Based on the directive, it is defined as follows^[13]:

$$L_{den} = 10 \log_{10} \frac{1}{24} \left[12 \cdot 10^{\frac{L_{day}}{10}} + 4 \cdot 10^{\frac{L_{evening}+5}{10}} + 8 \cdot 10^{\frac{L_{night}+10}{10}} \right] \quad (2.2)$$

where the computation of L_{day} and $L_{evening}$ is similar to that of L_{night} in equation 2.1.

The current Dutch regulatory system for Schiphol airport is indeed based on both L_{den} and L_{night} limits². In a number of locations placed strategically around the airport, the annual noise load is limited in terms of these two noise metrics. The details of this regulatory system will not be discussed here, but are presented in Section 2.5.3. On top of these location based maxima, the airport also faces other limitations related to the L_{den} and L_{night} levels, such as the number of dwellings exposed to a particular noise load. Section 2.1.4 will go into detail on these aggregated indicators for community noise exposure.

2.1.3 Noise modelling

With respect to noise modelling, at least two major strands of noise models can be identified. The first strand concerns the *integrated* or *segmentation* models^[39]. A very popular example belonging to this group noise models is the Integrated Noise Model (INM)^[62]. INM is the FAA's official methodology for noise impact assessment in the vicinity of civilian airports since 1978. Because its popularity, it can almost be considered an international standard, even though many countries still have their own national model.

Models like INM depend heavily on tabulated data, often referred to as Noise Power Distance (NPD) curves. These curves provide noise levels as a function of engine power and distance between the aircraft and the observer on the ground,

¹'den' is an acronym for day-evening-night.

²As of November 1st 2010, the system involving the enforcement points is inactive to allow for an experiment concerning new aircraft noise regulations. The experimental rules are now translated into new legislation.

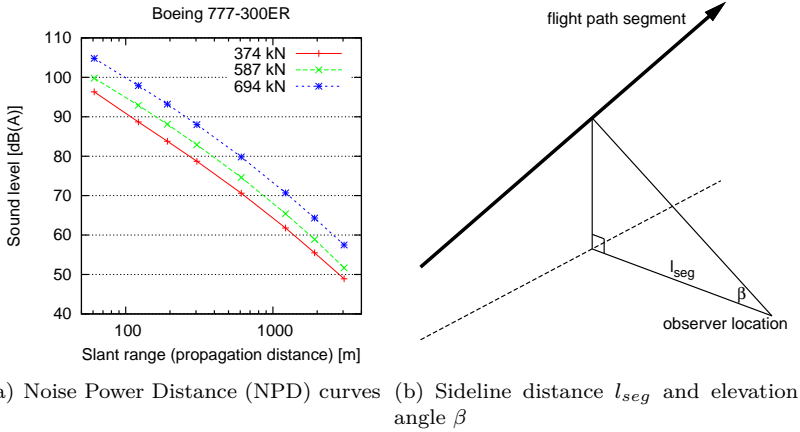


Figure 2.3

see figure 2.3(a) for an example. Basically, these curves translate source levels into levels observed at a distance by applying attenuation due to spherical wave spreading and due to standard-day atmospheric absorption and do so for multiple engine settings^[17]. Interpolation and extrapolation is used to obtain noise levels for power settings and distances not directly available from the data. After the interpolation process, several adjustments can be applied to obtain a more realistic noise level. An important correction is the *lateral attenuation adjustment*, which depends on the geometry between the flight path and the observer, see figure 2.3(b). The adjustment corrects for excess attenuation attributable, directly or indirectly, to the presence of the ground surface^[17].

The model by the European Civil Aviation Conference (ECAC) as laid down in ECAC document 29 (Doc.29) has the potential to become the European standard for airport noise assessment and has a lot of similarities with INM. In fact, both models rely essentially on the same standard^[4]. Concerning their classification as *integrated* or *segmentation* models, both models tabulate SEL data for infinite flight segments at a reference speed and correct the values as obtained from the interpolation process for actual flight segment length and duration. This means that these models not have to integrate the sound level histories to obtain the SEL values.

The second strand of noise models concerns the *simulation* models. These noise models calculate the sound contribution from consecutive aircraft positions along the flight path, thereby reconstructing the sound level time history^[39]. The current Dutch national model^[101] is such a simulation model. Still, overall, the model is very similar to INM and Doc.29. The major difference is that the Dutch national model only uses NPD tables for L_{Amax} . This means that SEL

values, required for computing L_{den} and L_{night} , are obtained by time-integrating the calculated sound histories of a single flight.

The noise simulation models are also available in higher-fidelity, more detailed variants. Typically, the noise source specification used by these models is highly directional, both in lateral as well as in longitudinal direction, and provided using 1/3rd octave bands instead of the A-weighted level. The propagation models used for these noise models are more complex as well and generally allow for more acoustical effects (e.g. shielding) and may be able to cope with non-standard atmospheric conditions (e.g. inversion³). Examples of models that use these more complex noise modelling techniques include NOISIM^[72], IESTA^[38] and FLULA^[88].

When comparing the different kinds of noise models, one should realise that the detailed simulation methods can be expected to provide more accurate results than the NPD-based models, especially under certain non-nominal atmospheric conditions and non-nominal flight conditions. Also for terrain that is non-standard (i.e not flat, acoustically soft terrain), simulation models can be more reliable. However, the simulation models are typically more computationally intensive than the segmentation models, and this is especially true for the complex, high-fidelity simulation models. This makes the segmentation models more suitable for analysing the noise load around airports, which involves computing the individual contribution of all annual flight operations, or at least all relevant contributions.

2.1.4 Community noise indicators and annoyance

Noise models are typically used to compute the noise levels in thousands of observer points around the airport, where these points are arranged in a grid with a rectangular mesh. The raw data as produced by the models is not suitable for decision makers to interpret. Therefore it is common practice to generate contours, lines of constant noise levels, and to present these contours overlaid on a map. It may however be desirable to also compute aggregated indicators in order to express the noise situation using a single number.

There are several commonly used indicators that are based on the noise level contours, such as the area that a certain contour level encloses. This indicator is very easy to compute since no airport-related information is required and is for example suitable to compare the noise performance of multiple aircraft types flying the same procedure. Contours are also used to count the number of dwellings or people enclosed within a certain contour. This obviously requires detailed demographical information, but results in indicators that are more suitable to express the actual noise load on the community.

³Normally, temperature decreases with increasing altitude in layers of the lower atmosphere. During a temperature inversion period, temperature increases with altitude, impacting sound refraction and therefore propagation.

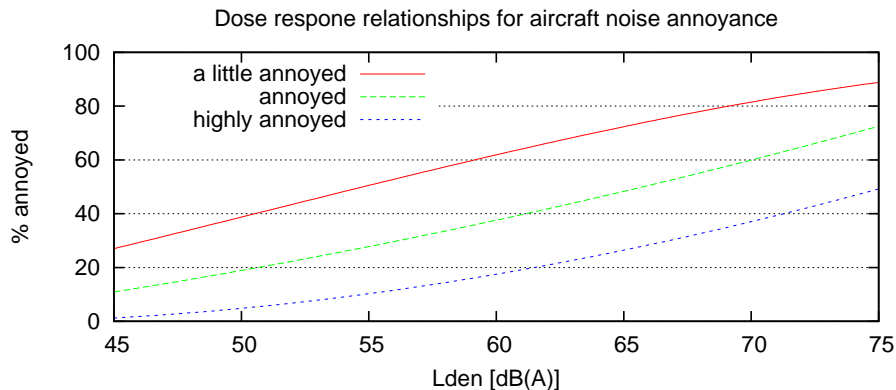


Figure 2.4: *Percentage of people feeling a little annoyed, annoyed, and highly annoyed with aircraft noise as a function of noise dose in L_{den} ^[84].*

What can be seen as a downside of indicators based on a contour is that they provide feedback with respect to a single noise level only. Even when choosing this level carefully, this still means that developments in the noise situation at higher or lower levels will not impact the chosen indicator. Apart from using multiple indicators, dose-response relationship can be used to prevent this situation. These relationships couple a particular noise load to an estimated response over a range of noise loads. By integrating the response over the total area of interest, this again yields a single number.

An example of a dose-response relationship is presented in figure 2.4. These curves present the percentage of the population that is feeling annoyed with aircraft noise as a function of noise load, either provided in L_{den} or L_{dn} ⁴, and are based on the long-term average annual dose. The curve fitting was performed on data from almost 50 different transportation noise studies in Australia, Europe and North America^[84]. It should however be noted that the provided response is an average for the three mentioned areas, but that the sensitivity to noise may differ within these areas. Apart from that, it should be also realised that actual annoyance can be influenced by non-acoustical factors^[53;97].

2.2 Aircraft gaseous emissions

As aircraft engines burn fuel, the combustion products are emitted behind the engine. Ideally, the combustion of hydrocarbon based fuels would lead to the formation of carbon dioxide and water vapour only. While these two gasses have

⁴ L_{dn} is similar to L_{den} , but does not recognise a separate evening period.

been identified as contributing to climate change, their effect on the local air quality around airports is generally not considered relevant. However, other undesired combustion products are formed, such as carbon monoxide, nitrous oxides, sulphur oxides, Unburned Hydrocarbons (UHC), and soot. These substances are formed either as by-products of complete combustion (NO_x), products of incomplete combustion (CO, UHC and soot), or as products of fuel impurities (NO_x and SO_x)^[93].

2.2.1 Emissions regulations

Similarly as for noise, aircraft pollutant emissions regulations are laid down in ICAO's Annex 16. In this case however, the compliance standards are not formulated for airframe-engine combinations, but for the engine itself. Four different pollutants are monitored through this mechanism, namely CO, UHC, NO_x and smoke, where smoke is defined as *carbonaceous materials in exhaust emissions which obscure the transmission of light*^[6].


Emission rates for these substances are highly thrust setting dependent, as can be seen in figure 2.5(a). Considering that historically the primary interest was to control engine pollutants in the vicinity of airports and the fact that aircraft use a wide range of thrust settings near the airport, the concept of the Landing and Take-off (LTO) cycle was introduced in relation to emission measurements and certification limits. Depicted in figure 2.5(b), this simplified cycle assumes five operational phases (approach, taxi-in, taxi-out, take-off and climb out), with corresponding engine settings (30%, 7%, 7%, 100% and 85% of the rated thrust). The results of the emission measurements of the different airliner engines are available through the ICAO engine emissions databank^[8]. For practically all large commercial turbo-fan engines, this databank provides the emission rates for CO, UHC, NO_x as well as the smoke number for each of the four LTO engine settings, together with the corresponding fuel flow. The latest version of the databank is hosted online by the UK CAA⁵.

From an historical perspective, technology improvements have greatly reduced aircraft CO and UHC emissions in the vicinity of the airport. Soot emissions have also been reduced significantly. NO_x emissions however remain a challenge, also because improving the fuel efficiency of the engine tends to increase NO_x emissions^[30:93]. Currently, attention also seems to be shifting towards NO_x during the cruise flight, because of the suspected climate change contribution, by influencing ozone and methane concentrations at typical cruise altitudes^[87].

2.2.2 Emissions modelling

A very common method in modelling aircraft engine emissions near the airport is an inventory method that is based on the values as obtained for the ICAO

⁵<http://www.caa.co.uk/default.aspx?catid=702>



(a) Emission level as a function of thrust setting (b) Graphical representation of the LTO cycle

Figure 2.5

reference LTO cycle^[5]. For this method it is assumed that aircraft, on average, spend a particular time in each of the four engine modes (approach, idle, take-off and climb out) for each cycle. Only the time below a certain altitude is taken into consideration, where 3,000 ft above ground level is the default value. It is assumed that pollutants emitted above this height disperse and have no significant ground level effect. The total cycle emissions are calculated by multiplying the emissions indices and fuel flows for a particular aircraft type with the corresponding times in mode and summing for the four modes. In mathematical form:

$$E_i = \sum_{j=1}^4 EI_{ij} \cdot TIM_j \cdot FF_j \cdot N \quad (2.3)$$

where E_i are the emissions for pollutant i , EI_{ij} are the emissions indices for pollutant i in mode j , TIM_j are the times (durations) in mode j , FF_j is the fuel flow rate per engine in mode j and N is the number of engines. Total emissions that can be attributed to a specific airport are calculated by performing this routine for all aircraft movements.

Because the data available from the ICAO engine emissions databank is only valid under the (standardised) test conditions, and therefore does not provide data for varying atmospheric conditions and intermediate thrust settings, the original inventory method is somewhat limited. The Boeing Method 2^[31] model (or Boeing curve fitting model) can be seen as an extension to the original inventory method that does not suffer from these limitations. Basically, the extended model allows for alternate thrust settings by means of interpolation procedures and allows for non-standard atmospheric conditions and engine bleed and installation effects by means of several correction formulas.

Models that can offer higher accuracy than Boeing Method 2 are available as well. A downside is that these kinds of models typically rely on several combustion process parameters that are considered proprietary to the engine manufacturers. This property makes these kinds of model less suitable for emission modelling, except to the parties that have access to the required data.

2.3 Third party risk

In aviation, Third Party Risk (TPR) evaluation is used to gain insight in the risks for the community of direct fatalities arising from aircraft accidents. This risk is designated as *third party*, or sometimes *external*, because the people exposed to the risk are not involved in flying, in contrast to the passengers or the crew of a flight. This also means that the exposure to the risk is involuntary.

Although the probability of an aircraft accident on a per flight basis is very low, the high number of annual flight movements close to busy airports leads to a concentration of this risk in a particular area. Contributing to this phenomenon is that about 70% of the aircraft accidents actually take place close to the airport [58]. Unfortunately, the benefits of increasing levels of safety on a per flight basis, often do not compensate for the increase in traffic, resulting in an increasing risk of an aircraft accident on a yearly basis [100].

Two metrics are in use to express the TPR levels [58]. The first one is individual risk and can be computed for a location. Individual risk is defined as the chance per year that an individual living permanently at that location will be killed by an aircraft impact. Because of the location dependency, individual risk is usually presented by drawing risk contour levels on a map, similarly as for noise. Alternatively, group risk (or societal risk) is not location dependent and is the annual probability that a group of N or more people are killed due to an aircraft accident. Group risk is presented using F-N curves, where the probability F is plotted as a function of group size N , see figure 2.6.

2.3.1 Third Party Risk regulations

Third party risk related to airports is not a field that is assessed worldwide. There are also no internationally agreed standards as could be identified for noise and emissions. However, especially in the UK and in the Netherlands, the issue of TPR is studied and even regulated. Therefore this section will provide an overview of TPR policy in these two countries.

The United Kingdom

The UK uses the concept of the Public Safety Zone (PSZ), areas of land extending from the runway ends within which development is restricted to limit the number of people exposed to high TPR levels. The size of the zone is primarily based

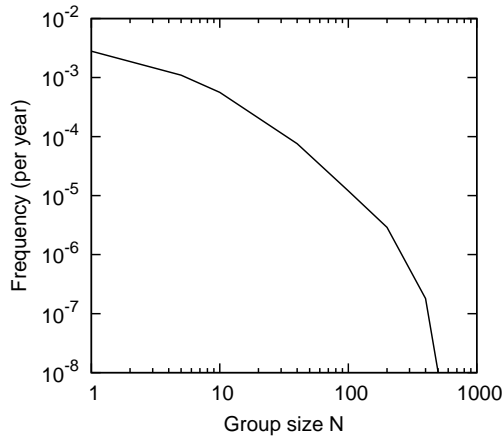


Figure 2.6: *F-N curve (Illustrative purposes only)*

on the calculated 10^{-5} individual risk contour, using a traffic forecast for 15 years ahead and a modelling methodology to be discussed in the Section 2.3.2. However, the PSZ itself is a simplified geometrical representation of the contour, usually in the shape of a triangle or an elongated pentagon extending from both ends of the runway^[23].

The general policy is to prevent any new or replacement development within the PSZ, both residential as well as non-residential. There are however several exceptions that may lead to approval, such as extensions to existing houses believed not to increase the number of inhabitants. Likewise, certain forms of new low density commercial activities may be approved, such as long term parking facilities or a golf course.

Especially for larger airports, there might also be residential or commercial property enclosed within the 10^{-4} individual risk contour. In that case, the airport operator is obliged to make an offer to purchase such a property. If the owner agrees in selling, the airport operator is expected to demolish the property afterwards.

The Netherlands

The Dutch land use policy with respect to TPR is primarily based on the calculation of individual risk contours around the larger airports. Two zones have been defined that correspond to the location of the 10^{-5} and the 10^{-6} contour. Within the inner zone, all existing dwellings are to be demolished eventually. However, people already living within the area at the time the zone was established cannot be forced to move to another location. Within the outer zone, new development is restricted, somewhat similar to the PSZ.

Apart from land use planning considerations, Schiphol Airport also faces operational limits based on the total risk level. This limit is expressed using a metric not discussed before: the total risk weight. For a particular flight, the risk weight is the product of the statistical crash rate and the maximum take-off weight. The results are summed over all movements to obtain the total risk weight. For Schiphol airport, the limit is currently 9,724 tons.

2.3.2 Third Party Risk modelling

Although there are different TPR models, the method that is employed to calculate the risk levels seems to be similar. Three main elements can be identified when analysing this method^[1]:

- First, the crash frequency is determined, that is the probability of an aircraft having an accident near the airport. This accident rate can be determined using historical data. When this accident data is categorised, different accident rates can be obtained, depending on the type of movement (arrival/departure), and aircraft characteristics (size, engine type, etc.).
- The second element involves the crash location probability. Generally speaking, the probability of an accident increases in the proximity of the runway and the arrival and departure routes. This can be modelled using crash location distribution functions, based on the historical data of accident locations.
- The final main element concerns the consequence of the impact. This is modelled as the area of the crash site and the lethality of the crash within that area, typically as a function of aircraft parameters, take-off mass in particular.

When combining these three modelling elements for all annual flights, the individual risk levels around the airport can be computed. When the model is also provided with population density information concerning the airport under consideration, the group risk graph can be constructed as well.

2.4 Other airport environmental issues

The previous three sections introduced three important performance areas with respect to the airport environmental management process. What these three subjects have in common is that the associated environmental impact can be attributed to particular aircraft movements. For the airport however, there may be several other environmental issues, more related to the presence of the airport as a whole. While these remaining environmental management areas are deemed outside the scope of this work, this section provides a high level overview of three more subjects: air quality^[100], water quality^[46] and wildlife management^[46].

Air Quality

The air quality in the vicinity of airports is influenced by more factors than aircraft emissions alone, which means that the subject of air quality is actually broader than discussed in Section 2.2. Other important sources of pollution are apron activities, including aircraft refuelling and ground transport, not only at the airport itself, but also in the area surrounding the airport. Especially if the airport faces current or future absolute air quality standards, managing the air quality may become a very important issue. However, not all of the pollutant sources are easily manageable from an airport point of perspective. While it may have the authority and opportunity to reduce the emissions from apron activities, reducing ground transport emissions however is more complicated, especially if it concerns traffic outside the airport boundaries.

Water Quality

Airports may face the risk of contaminating the ground water, surrounding ditches and connected waterways. A typical contaminant, at least in colder climates, is glycol. Glycol-based fluids are used for de-icing and anti-icing, the process of removing and preventing snow and ice accumulation on the surfaces of departing aircraft, which is critical for a safe operation. As these fluids are sprayed onto the aircraft, some of the fluid will reach the ground. Without proper measures, such as dedicated drainage and water treatment systems, large quantities of these fluids are likely to reach the ground water system eventually. Apart from de-icing fluids pollution, spills and leaks of other fluids, such as fuel and oil, may also pose a contamination risk.

Wildlife Management

Flocks of birds flying through the aircraft arrival and departure paths form a serious safety risk to aviation, sometimes even resulting in fatalities on top of the far more common economical damage. Especially heavy birds can damage aircraft surfaces and result in engines failures when ingested⁶. Airport often try to scare off birds using a wide range of tactics, including sounds, laser light and falconers. Apart from collisions with flocks of birds, airports also have to prevent collisions on the runway with large animals. This is typically less complicated than bird control, as this can be achieved by enclosing the area using fences and/or ditches.

Finally, what might be important to mention here is that the list of six environmental management subjects presented here is first of all not universal. For each airport in the world, which areas are deemed to be important will depend on local

⁶A recent and extensively covered bird strike accident concerned the January 15, 2009 US Airways Flight 1549 that made a successful water landing on the Hudson river after a collision with multiple Canada Geese resulted in an almost complete loss of thrust shortly after departing LaGuardia Airport, <http://www.nts.gov/publictn/2010/AAR1003.htm>

regulations, land use in the surrounding area, size of the airport, public opinion, etc. Furthermore, the list is not complete, as some airports also work on further subjects, such as reduction of waste and their climate change contribution.

2.5 Impact reduction based on optimisation

Sections 2.1 through 2.3 summarised three important performance areas with respect to airport environmental management. Based on several of the mentioned examples, it can be deduced that for each of these areas there are regulations to limit the negative side-effects of aviation on the community. Certainly not all airports in the world are facing the same regulations, as these regulations may range from a global to a regional, or even airport specific level.

For situations where airports do not meet current or near future environmental limits, or where airports are restricted in their current number of operations, the rationale behind further impact mitigation is clear. However, even for situations where the actual demand for air traffic can be accommodated within existing limits there are multiple motives for still attempting to reduce environmental impact further. For example by reducing the impact per flight or passenger, airports and airlines can anticipate for an expected future growth in traffic and accommodate that growth within the current environmental impact limits. Furthermore, reduction of impact may lead to other economical benefits, for example when it enables housing development in areas that were deemed unsuitable before. As a final example, demonstrating commitment with respect to mitigation may also be beneficial in terms of public opinion.

This section will present seven examples of environmental impact mitigation efforts. These examples have in common that they utilise some form of optimisation in order to minimise impact. They are classified according to the three levels of aggregation in mitigation measures as defined in Section 1.4.

2.5.1 Trajectory level

At the trajectory level, the primary instrument is to alter an aircraft's trajectory such that it results in a beneficial change with respect to the corresponding environmental impact. Optimisation techniques can be used to find the alternative trajectories that are most efficient with respect to the objectives that apply while simultaneously meeting all constraints, for example constraints with respect to air speed. This section will briefly introduce three examples of aircraft trajectory optimisation studies that use environmental objectives.

NOISHHH

The first trajectory optimisation example concerns a non-linear programming tool for the optimisation of arrival and departure trajectories. This tool is called

NOISHHH and was developed at TU Delft^[106;107]. The tool generates 3D or 4D trajectories that minimise environmental impact in the residential communities surrounding the airport. The objective function can be freely chosen as a mix of multiple objectives, including various noise related criteria and fuel use. Several numerical studies have shown the capabilities of this tool with respect to the reduction of environmental impact in the vicinity of the airport.

An extended version of NOISHHH is one of the tools that has been conceived in the context of this thesis. Therefore the tool, its capabilities, underlying optimisation techniques and developed extensions will be discussed in more detail in Chapter 3.

Lexicographic egalitarian optimisation of departure trajectories

A second example of a study into aircraft trajectory optimisation for noise abatement purposes investigated several optimisation techniques^[89;90]. Based on identified strengths and weaknesses for each of the studied approaches, the final recommended optimisation technique is a combination of various approaches and is called lexicographic egalitarian optimisation. The egalitarian component is used to guarantee a form of fairness among different noise-sensitive locations, by minimisation of the modelled noise annoyance deviation at the worst-off noise-sensitive location. The lexicographic component, which involves optimising multiple objectives iteratively, is used to eliminate weak solutions that may be the result when only using the egalitarian method.

The application of a fairness criterion is one of the more notable properties of this study and highlights a dilemma that applies to most of the optimisation applications discussed in this work. This dilemma concerns the question whether it can be justified to increase impact levels for some in order to achieve a reduction for others, as long as it results in an overall, community-wide reduction. For now it is assumed that this can indeed be justified and this will be discussed further in Section 4.5.

Genetic optimisation of arrival trajectories

A third and final example trajectory optimisation studies concerns the the use of a genetic optimisation algorithm for generating optimised trajectories from the en-route phase of the flight to the designated landing runway of the destination airport^[108;109]. Genetic optimisation is a solution search technique that is based on the process of natural evolution and uses terminology including *generations*, *inheritance*, and *mutation*. One of reason that this technique was selected for this study is that genetic algorithms conveniently allow for the optimisation of multiple criteria simultaneously.

Using fast-time simulations, this study indeed confirms the benefits of trajectory optimisation. It shows that optimising the arrival trajectories can not only increase airport throughput but also has the potential of reducing the noise

impact on the community. The work does not only focus on the optimisation of the trajectory itself, but also studies pilot and air traffic controller task demand load. It is concluded that both pilot and air traffic controller task demand load may increase when optimised arrival trajectories are used instead of standard arrival trajectories. For pilots, it is shown that new guidance displays when made available on the flight deck may alleviate the additional task demand load and it is suggested that improvements in the human-machine interfaces for air traffic controllers may help to reduce their task load.

Although this study does use trajectory optimisation, the second part of the described activities, such as the fast-time simulations and determining air traffic controller task demand load, do no longer meet the definition of the trajectory level as set in Section 1.4. Instead, this should be classified as the operational level. This clearly shows that mitigation efforts sometimes span multiple levels.

2.5.2 Operational level

Based on the stated definition, mitigation at the operational level comprises activities related to the actual day-to-day operation of the airport, in particular the provision of air traffic control in relation to environmental performance. The two examples that will be used to illustrate the operational level are both decision support tools intended to aid air traffic controllers in making decisions that help reduce environmental impact. However, before presenting the *Noise Avoidance Planner* and the *Expedite Departure Planner* concepts, the operational context for these tools is provided first.

CTAS and FAST

The United States National Aeronautics and Space Administration (NASA) and Federal Aviation Administration (FAA) have been working on a set of decision support tools to help air-traffic managers and controllers in improving flight efficiency and airspace capacity^[74]. Together, these tools are known as Center-TRACON Automation System (CTAS). They are designed for en-route air traffic management at area control centers and Air Traffic Control (ATC) at terminal control centers typically controlling approaching and departing around larger airports⁷. An example of a tool is Final Approach Spacing Tool (FAST). It is designed for operation in the terminal control centers and provides landing sequences and landing runway assignments. It generates speed and heading advisories that help controllers manage arrival traffic and achieve an accurately spaced flow of traffic on final approach.

⁷In the US, an area control center is referred to as Air Route Traffic Control Center (ARTCC) and a terminal control center is referred to as Terminal Radar Approach Control (TRACON)

Noise Avoidance Planner

Some of the CTAS tools feature *environmental awareness*, in the sense that they try to reduce community noise impact and/or emissions. The proposed concept of the Noise Avoidance Planner (NAP) is such a tool, although it can also be seen as extended functionality for FAST^[42]. As mentioned, FAST is responsible for sequencing and scheduling arriving traffic in the terminal area. In performing this process, FAST generates multiple possible future trajectories for constraint resolution, called the trajectory space. The NAP performs noise calculations for the different trajectories and generates a noise figure of merit for each of them^[41]. In computing this figure it does not only consider the noise impact of the current event, but also past and future events (aggregated noise), all in relation to population data. Using the noise rating from the noise sensitivity computation, it should be possible to turn the current single objective efficiency optimisation into a multi-objective one, optimising for both noise and efficiency. However, because of the current CTAS constraint resolution architecture, noise considerations and efficiency could not be addressed simultaneously.

Expedite Departure Planner

The Expedite Departure Planning (EDP) is the second tool to which the concept of the NAP can be applied. The tool should help controllers in efficiently sequencing, spacing and merging departure aircraft into the en route traffic flow. The proposed concept of NAP applied to the EDP is similar to that of FAST. There is, however, a second community noise reduction effort the departure planner may help to achieve. One of the functions of the EDP is to provide the departure controllers with the ability to perform unrestricted climbs, where currently departures are mostly kept below arriving traffic streams in the congested terminal areas (procedural separation). For this function, the tool uses accurate climb trajectory predictions and checks for conflicts with arriving traffic. If no conflicts are expected, it informs the controller that an unrestricted climb is advisable. This reduces time to climb, fuel burn and related gaseous emissions and also community noise impact^[75].

2.5.3 Strategic level

This section concludes with two examples of noise load minimisation at the strategic level. Since both examples are related to the regulatory system in place for Amsterdam Airport Schiphol, this system is discussed first.

Schiphol regulatory system

Since February 2003, the Dutch government controls the maximum annual noise load around Schiphol Airport through two sets of enforcement points⁸. The enforcement points are geographical locations near residential locations around the airport for which a maximum noise load has been defined. The first set consists of 35 points depicted in the left side of figure 2.7 and limits apply to the annual L_{den} values in those 35 points. While each point has its own limit value, most limiting values are near 58 dB(A) L_{den} . The second set of 25 points is used to limit L_{night} , which is influenced by night times operations only. Again, all points have their own limit value, typically near 50 dB(A) L_{night} . Because the airport can be sanctioned for exceeding the limit values, it tries to influence the development of the noise load in the enforcement points throughout the year such as to remain within the legal limits. Indeed, the primary objective is to minimise the chance of exceeding any of the limits. The noise load development can be influenced by making small changes to the preferential runway allocation rules. In practice, the airport, in consultation with other stakeholders changes the order of a list of preferential runway configurations. This list is used by the Air Traffic Services (ATS) provider to select a runway configuration for a particular weather and traffic situation. Changes in the order of the list results in changes in runway use, which in turn results in a change in the distribution of the noise load.

Genetic optimisation of the noise load

The airport has developed an optimisation method for determining the best configuration order throughout the year^[54]. This method is based on the concept of modes and phases. A mode is a way of operating the airport, according to a particular preferential order. A phase is a time period within the operational year, typically a month. For each mode and phase combination, the expected noise load contribution in each of the 60 enforcement points is calculated, using the traffic schedule for that phase and the corresponding historical weather data as inputs. A genetic optimiser then searches for the best mode for each phase in order to minimise the risk of exceeding the noise limits. Additional constraints are used, for example to limit the number of mode changes during the year. The method can be used before the start of a particular operational year, but also for the remaining part of the current operational year, while taking the current realisations into consideration.

Stochastic dynamic programming as alternative approach

A later study performed at the Air Traffic Control the Netherlands (LVNL) pointed out that the method developed by the airport can be improved further by taking into account possible changes in future weather conditions, by using weather

⁸Starting November 2010, a new regulatory system is being experimented with. This new system focusses more on noise impact, rather than noise load.

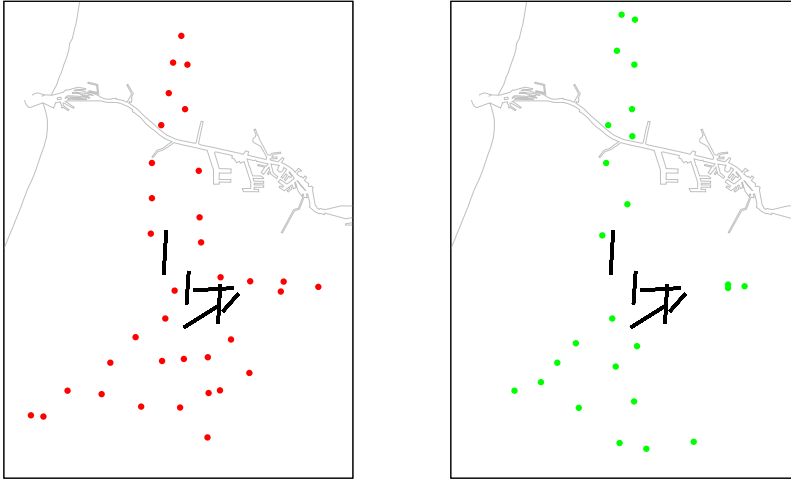


Figure 2.7: The L_{den} (left) and L_{night} (right) enforcement points located around the airport

development probability distributions instead of average weather and the possible adaptation of the preference list due to these weather developments^[82]. It is proposed to use the stochastic dynamic programming method rather than the genetic optimisation approach. A feasibility study based on only three enforcement points showed that the developed method indeed outperforms the original method. Unfortunately, the runtime of the problem increases exponentially with the number of enforcement points and evaluation of all points was not possible within a reasonable computation time.

2.6 The integrated approach

As already mentioned in Section 1.3, there have been efforts toward a more integrated approach in the impact analysis part of airport environmental management by combining multiple environmental impact models into a single system or toolbox, such as the development of the Aviation Environmental Design Tool^[35;85]. Apart from an expected benefit with respect to modelling efficiency, having multiple models use the same set of input may also contribute to more consistency in the modelling of environmental measures. Furthermore, it can be expected that it becomes easier to evaluate the impact of for example a particular noise mitigation measure on the resulting gaseous emissions or vice versa, which is also a benefit.

The introduction mentioned a second meaning of an *integrated approach*, based

on the idea that all environmental impact mitigation efforts taking place at the different levels of aggregation should be managed concurrently. The remainder of this section will provide an example of this explanation of an *integrated approach*, based on the noise impact management strategy of Luik Bierset airport in Belgium.

2.6.1 The Luik Bierset example

The Luik Bierset airport has a very clear strategy concerning noise impact, which is called an integrated approach for noise impact^[43]. The airport has opted for implementing a noise concentration policy, through defining routes that are identical for at least the first 15 nautical mile. The routes are optimised such that there is a minimum of buildings within several of the L_{den} contours of the airport. The complete set of actions that have been taken for this integrated approach is: route optimisation, official route definition, permanent monitoring of the routes, permanent monitoring of the noise at the ground, sound proofing based on noise zoning and rules for urban development and new building construction. Concerning the sound proofing, a zoning system is used, which is based on four categories according to L_{den} level, but a set of peak level limits also apply for these zones. Noise insulation requirements differ for the different levels and guarantee that peak levels indoors will never exceed 55 dB(A) L_{Amax} in living rooms and 45 dB(A) L_{Amax} in bedrooms.

The actions taken related to this integrated approach can indeed be categorised into the three different levels: the trajectory, operational and strategic level. Although it is a very good example of integral noise management where actions at one level are supporting the effectiveness of the actions at other levels, it is important to realise that this involves a rather small airport. This airport not only has a layout featuring a single runway, it is also running well below physical capacity because of an operating restriction of twenty operations per hour. This makes it much easier to actually implement results from trajectory optimisations efforts in practice. Furthermore, at the strategic level, control options are limited because of the single runway orientation and the policy choice to concentrate traffic as much as possible. All in all, while it can be concluded that this is a meaningful approach for a small and not too busy airport, it is not an obvious solution for complex and congested mainports.

2.7 Discussion

The first part of this chapter provided an overview of important performance areas with respect to airport environmental management, particularly aircraft noise, gaseous emissions and third party risk. The second part continued with examples of airport environmental impact mitigation measures and strategies reported in the literature that are considered relevant in the context of this thesis.

As discussed in the introduction, it is expected that some of the current inefficiencies of mitigation measures can be removed when a process is used that is based on three main principles: (i) to use mathematical optimisation in order to select the best mitigation options, (ii) to evaluate multiple performance areas simultaneously, and (iii) to evaluate multiple mitigation options at multiple levels of aggregation simultaneously. The examples of existing mitigation efforts presented in this chapter show that each of these principles has already been applied before. However, the question remains whether these three principles can also be applied simultaneously, by implementing them in a support system, and whether such a system would indeed contribute to additional benefits with respect to airport environmental impact.

Parts of the work presented in this thesis, especially in Chapter 3 and 4, will be based on mitigation efforts that have been discussed in this chapter. Chapter 3 will utilise an adapted version of the NOISHHH tool, as introduced in Section 2.5.1. Compared to the original version, the adapted version is slightly simplified, yielding a considerably faster and more robust tool. These two characteristics allowed for the generation of a database containing over 500 optimised trajectories in a relatively short period of time, and this database enabled additional research opportunities.

Furthermore, the work of Chapter 4 will also be based on a mitigation concept previously reported, as it is based on the concept of the Noise Avoidance Planner, described in Section 2.5.2. Unlike the NAP, the tool presented in Chapter 4 is not suitable for an operational environment and has been built for research purposes only, but on the other hand, it is able to address noise and efficiency objectives simultaneously.

Advanced noise abatement departure procedures

This chapter^[65;66] presents an aircraft trajectory optimisation application geared towards reducing environmental impact near airports. This application is presented here as a stand-alone concept or measure, but as the trajectory level system, it will also be the lowest level building block for the multi-level integrated environmental management system presented in Chapter 6.

First, the concept of custom optimised departure profiles is outlined and this chapter then explains the rationale behind this specific implementation form of trajectory optimisation. Based on this concept, the different actors and their responsibilities related to this problem are introduced. Section 3.2 presents the trajectory optimisation method chosen for generating the custom optimised departure profiles. It also describes the optimisation tool, which is based on a previously existing trajectory optimisation tool and was adapted to match the concept. Section 3.3 contains the numerical results of optimised departure profiles that have been obtained with the developed tool. The results focus on the differences compared to the current departure profiles, primarily with respect to speed, altitude, fuel burn and noise impact. Based on these results, this chapter also presents an additional study into the expected effects of the optimised profiles on runway capacity, followed by a discussion in Section 3.5.

3.1 The concept of custom optimised departure profiles

A common operational measure to limit the noise exposure in residential areas is the use of noise abatement procedures^[50]. For example, two well known types of noise abatement departure procedures (NADP) are the so-called distant and

close-in NADP's^[18]. Procedures that belong to one of these types are designed to bring noise relief either close to or somewhat more distant from the airport.

A characteristic that is shared among these types of procedures is that they are generic in nature, i.e. they are not optimised with respect to the local situation^[105]. This means that although they have been optimised for a particular shift in noise load, they do not take the actual population distribution around the airport into consideration.

These procedures are generic in a second sense as well: they are typically identical for all aircraft under all circumstances, irrespective of aircraft type, loading conditions and atmospheric conditions. All of these parameters do have a significant influence on the resulting trajectory and therefore on the distribution of the noise load. So in terms of noise abatement, these standard procedures are not optimised for a specific flight.

The remainder of this chapter presents the concept of custom optimised departure profiles. Under this concept, it is foreseen that aircraft can fly a departure profile (i.e. the trajectory in the vertical plane) that is optimised for each particular flight. At the same time, the routing (i.e. the trajectory in the horizontal plane or ground track) is still based on fixed and published routes based on Area Navigation (RNAV) principles¹.

3.1.1 The rationale behind the concept

Trajectory optimisation has already been introduced in Section 2.5.1 using three example applications. What the three examples have in common is that all at least optimise the routing, sometimes combined with optimising the vertical profile.

Although these fully optimised trajectories typically offer substantial benefits with respect to environmental impact, several problems arise with respect to a potential implementation. For example, the proposed trajectories are definitely more complex from the pilot's point of view, and not always compatible with today's navigation and guidance principles. A second problem that can be identified is a considerable increase in airspace complexity^[108]. With each aircraft flying its own optimised trajectory, the regular traffic patterns as known today will no longer be present. This not only leads to a far more complex situation for air traffic controllers, the situational awareness of the pilots could be impacted as well, especially with respect to nearby traffic.

It seems reasonable to assume that the navigation and guidance difficulties can be overcome by near future levels of flight deck automation. That means that the airspace complexity would be the dominating problem. The concept of custom optimised departure profiles presented in this chapter does not suffer from the increase in complexity, because this concept is still based on published departure routes. At the same time, it is still capable of reducing the environmental impact by means of optimisation of the departure profile. If indeed the

¹For background information on RNAV versus the traditional form of instrument-based aircraft navigation, please consult Appendix A

identified complexity problem would prevent or delay a future implementation of fully optimised trajectories, the optimised profile concept would be an interesting (interim) alternative.

It is recognised that not optimising with respect to the horizontal plane reduces the benefits that can be expected from trajectory optimisation. The performance penalty in terms of e.g. fuel burn and noise reduction that is expected is regarded as the price to be paid for a reduced complexity. However, the proposed concept does not completely rule out all lateral optimisation opportunities. Optimal routes can still be computed, as long as there is a single result that is acceptable for all possible movements and the published departure routes are updated accordingly^[60].

3.1.2 Actors involved and their responsibilities

The next three subsections describe which parties would be involved in the concept of custom optimised departures profiles and what their responsibilities would be. A primary role is identified for the airline and the ANSP and an optional role is foreseen for a governmental body.

Airline responsibility

Although the actual profile optimisation and selection for a specific flight could be made the responsibility of the ANSP, for this concept it is assumed that it remains with the airlines. There are several reasons for this task allocation:

- Realistic trajectory optimisation or even simulation requires detailed and accurate knowledge with respect to the aircraft status. Aircraft weight for example is currently unknown to the ANSP, but is known to the airline, at least at takeoff.
- Airlines have detailed information with respect to the performance of their aircraft, as provided by the manufacturers. As this data is considered proprietary, ANSPs would probably have to use higher-level performance models, resulting in less accuracy
- Placing the profile optimisation responsibility with the airlines also seems to fit the concept of the Business Trajectory (BT), an important element of the Single European Sky ATM Research Programme (SESAR)^[21]. The BT is the trajectory that best represents the airspace users intentions for a specific flight.
- Related to this principle and typical for the departure phase is the selection of some sort of reduced thrust mode for a specific departure (de-rate, flex, etc.). For the optimal departure profiles concept, it is still the airline that is free to perform the trade-off for this decision, just like it is in the current situation.

For the airline, there are several kinds of benefits that can be expected from flying optimised profiles. First, there are opportunities for costs savings. For example, the possibility to fly fuel optimal trajectories can be exploited in order to realise fuel savings. And with carbon emissions trading schemes that are being introduced, savings from fuel burn reduction are likely to exceed those from the fuel cost alone. Furthermore, airlines executing noise optimised departures could use this fact in the currently popular endeavour of convincing their (prospective) customers of their *greenness* as a company. Similarly, and especially relevant for carriers that are dominant at one of the more noise sensitive airports, they can present it to the surrounding communities as one of their efforts of being a good neighbour. For both examples, the airline will see no direct financial savings from operating in an environmentally less harmful way. This could change however, if negative externalities² are priced as well, based on the *polluter pays principle*.

ANSP responsibility

The additional responsibility for the ANSP with respect to the optimal departure profile concept would be to define the constraints for the trajectory optimisation and to present them to airlines wishing to determine an optimal profile for a certain departure. In the simplest form, these constraints only consist of minimum altitudes with respect to obstacle clearance, together with an altitude restriction on joining the airway at the end of the SID. Together with these restrictions, it might be convenient to distribute the actual or forecasted meteorological conditions to be used for the profile optimisation. Airlines could still decide to actually use this data, or to use similar data from their own sources.

Governmental responsibility

Apart from safety related issues, it seems that government involvement and responsibility with respect to the presented concept is not strictly necessary. It would be limited by providing economic incentives to correct market failure in the management of natural resources and the control of pollution. The term Environmental Fiscal Reforms (EFR) is often used in relation to this process^[49]. With respect the custom optimised departure profiles concept, the benefits of applying EFR can probably best be shown using a numerical example, based on the trajectory optimisation tool that also forms the basis of the concept discussed in this chapter. This tool typically has the possibility to significantly reduce the number of people at risk of awaking due to fly-over events. Apart from the previously mentioned effects on goodwill and image, there is, however, no incentive for the airlines to take this into consideration. In fact, compared to a strictly fuel optimised trajectory, an airline would even have to accept a slightly increased fuel

²negative externalities, or external costs *arise when the social or economic activities of one group of persons have an impact on another group and when that impact is not fully accounted, or compensated for, by the first group*^[16]

burn^[107]. Comparing the differences in fuel burn with the number of potential awakenings, it shows that most avoided awakenings are very inexpensive, with *exchange rates* sometimes as low as one thousand awakenings to a gallon of jet fuel^[70]. This means that the overall societal trade-off is clear and EFR could be used to help the airline to make the trade-off.

The main government responsibility in this respect would be to set the price on environmental performance, for example by applying discounts and surcharges. It is expected that this would direct the airline, not only in making a decision for a particular flight, but also in deciding on investing in the technology required for enabling a concept like custom optimised departure profiles.

Local community groups can also be seen as an actor in the process of managing airport environmental impact and could therefore also be involved in the concept of custom optimised departures profiles. For now, it is assumed that objectives and demand of these groups have been noted by the responsible governmental body and that these have been taken into account at setting the resulting policy.

3.1.3 Important assumptions regarding the concept

First of all, the concept is based on aircraft capabilities in terms of flight performance. As mentioned, it is assumed that any existing navigation and guidance difficulties can be overcome by near future levels of flight deck automation and that this will result in a concept that is workable for the flight deck. This has however not been confirmed by any form of experiment, as this was regarded to be outside the scope of this work.

Similarly, it is expected and therefore assumed that the concept of custom optimised departure profile is workable for air traffic controllers. The reason that this is expected is because in the current situation, even when all aircraft fly the same departure procedure, numerous causes can be identified that would lead to differences in the resulting vertical profile. Examples include performance differences between different types of aircraft, aircraft instantaneous weight, and possibly even crew habits. This means that adding flights that fly optimised custom profiles does not necessarily lead to a greater variety in departure profiles for the different flights. Section 3.4 provides further insight into the validity of this assumption.

Based on the previous assumption, it is also expected that a partial implementation seems possible, resulting in a situation where airlines are free to decide on taking part in the optimised departure profiles mechanism. Airlines choosing not to join in can simply continue to fly the standard procedures, as they do today. At the same time it is conceivable that airlines that do participate would not do so for all of their aircraft types. Some types could be more suitable for this concept than others, for example dependent on the navigation capabilities of the aircraft.

Finally it should be mentioned that passenger comfort has not been looked

into. Although the resulting trajectories are within the normal capabilities of the modelled aircraft, it is conceivable that for example numerous flight path angle or thrust setting changes during a single departure procedure would result in a situation that is less comfortable, compared to the current standardised departure procedures.

The remainder of this chapter does no longer deals with the proposed concept, but rather focusses on the tool that was developed to compute the custom optimised departure profiles. It also shows numerical examples of optimised profiles and looks into the possible implications for runway capacity.

3.2 Aircraft trajectory optimisation model

As introduced in Section 2.5.1, a tool called NOISHHH is being developed at the Delft University of Technology to facilitate the design of advanced noise abatement procedures^[106;107]. The tool can generate routings and flight-paths for both arrivals and departures for which the single event environmental impact in the residential communities surrounding the airport is minimised, while satisfying all imposed operational and safety constraints.

NOISHHH has been used in many different configurations for a wide range of applications. This section presents the configuration that was developed to match the concept of the custom optimised departure profiles. This configuration consist of a new optimal control formulation (Section 3.2.1), existing optimisation objectives (Section 3.2.2 and 3.2.3), combined with an existing optimisation technique and solver (Section 3.2.4). The final two subsections of 3.2 summarise the limitations of the model and, based on those limitations, also present an alternative profile synthesis method that is expected to be more reliable in an operational environment, albeit at the cost of some environmental performance.

3.2.1 Aircraft equations of motion

The aircraft dynamics are modelled using a slightly simplified point-mass model, the so-called intermediate model. The underlying assumption for the intermediate model is equilibrium of forces normal to the flight path. Typically, for a three-dimensional trajectory optimisation problem the equations of motion are defined as:

$$\begin{aligned}
 \dot{x} &= V \cos(\gamma) \sin(\chi) \\
 \dot{y} &= V \cos(\gamma) \cos(\chi) \\
 \dot{z} &= V \sin(\gamma) \\
 \dot{\chi} &= \frac{g_0 \tan(\mu)}{V} \\
 \dot{E} &= \frac{V(T-D)}{W}
 \end{aligned} \tag{3.1}$$

where the states x , y and z define the aircraft position in an Earth-fixed Cartesian coordinates system, χ is the heading and E is the energy height³. The aircraft is controlled using the flight path angle γ , the bank angle μ and the normalised thrust setting parameter Γ . This parameter determines the actual thrust T using the following relationship:

$$T = T_{min} + \Gamma(T_{max} - T_{min}) \quad 0 \leq \Gamma \leq 1 \quad (3.2)$$

where the maximum thrust T_{max} and idle thrust T_{min} are aircraft type, altitude, flight phase and the Mach number dependent. Equation 3.1 also features the aircraft weight W and drag D . The aircraft weight is assumed to be constant during the entire departure procedure and the drag is computed using aircraft type dependent lift-drag polars. More details on this specific part of the model, also in relation to the selection of flaps and slats, can be found in Appendix B.4. The constraints on the states and the controls will be discussed in Section 3.3.1.

To summarise, the model for three-dimensional trajectory optimisation is based on five states (x , y , z , χ , E) and three controls (γ , μ , Γ). However, as mentioned in Section 3.1, the concept of the custom optimised departure procedures is based on fixed, pre-determined routes. When assuming that the aircraft will not deviate from this route, it appears that the optimisation problem addressed here is in fact only pseudo-three-dimensional. This provides an opportunity to reduce the number of state and control variables.

First of all, the original states x and y are eliminated in favour of a new along-track distance coordinate s :

$$\dot{s} = V \cos(\gamma) \quad (3.3)$$

The relation between x , y and s is also provided in figure 3.1.

For the noise calculations however, the flight track still needs to be specified in x and y coordinates. This incompatibility has been solved by tabulating x and y versus s for each route up front. During the optimisation run, interpolation on this lookup table yields the actual position for any given s . Through this mechanism, x and y have become implicit functions of s . The lookup table also includes the heading and the turn radii of the different turns. This means that the state variable χ is no longer required. Furthermore, the control variable μ can be eliminated as well, as it can be computed from the turn radius R and the corresponding ground speed:

$$\mu = \tan^{-1} \left(\frac{\dot{s}^2}{g_0 R} \right) \quad (3.4)$$

This means that the original five states, three controls problem has been reduced to a three states (s , z , E) and two controls (γ , Γ) problem. First of all,

³energy height is the sum of the potential and the kinetic energy and is computed as the sum of the altitude and the energy equivalence of the speed expressed in altitude




Figure 3.1: *Along-track distance coordinate s and turn radius R*

this reduction in state and control variables reduces the problem complexity and therefore leads to reduced computation times. This reduction in computation time has not been studied extensively, but the new tool has been observed to be between two and five times faster than the original version. However, the reduction leads to second important advantage was well. In the old situation, a new initial solution had to be generated for each new runway and route combination in terms of a history of position coordinates x and y . This was a tedious and time-consuming exercise. For the fixed-routing version of the software, which only uses the single position coordinate s , this is no longer required. As long as the boundary conditions do not change too radically, any previous solution will typically be accurate enough for the solver to compute the new optimum. This advantage can be exploited to increase levels of automation, when a large number of solutions needs to be generated, as will also be discussed later on.

3.2.2 Sleep disturbance as an optimisation objective

Aircraft noise from a single flyover is typically expressed by presenting contour levels on a map, based on either SEL or L_{Amax} . The noise levels or the shape of the corresponding contours are, unfortunately, no suitable parameters for optimisation. Instead, they can be used to calculate noise indicators, expressing a noise measure using a single number, which can be directly used as an optimisation criterion. Examples include generic indicators like the contour area, or site-specific indicators like number of people or dwellings within certain contours.

As already discussed in Section 2.1.4, these noise indicators only respond to the changes in their specific noise level and cannot express the overall noise situation in a single number. As this is desired, both for a more balanced result as well as for the stability of the optimisation tool, a dose-response relationship for single

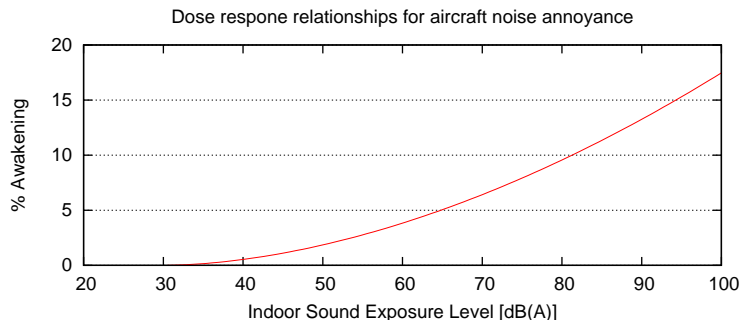


Figure 3.2: *Sleep disturbance dose-response relationship as proposed by FICAN^[9]*

flyovers has been used for this optimisation problem.

A well known single-flyover dose-response relationship has been proposed by the Federal Interagency Committee on Aircraft Noise (FICAN) and expresses the response in sleep disturbance^[9]. The curve, shown in figure 3.2, is based on the combined data and conclusions of three different field studies and represents a worst-case bound on the number of people likely to awake. The study also provides the expression for the percentage of the exposed population likely to be awakened (% A) due to a single event noise level:

$$%A = 0.0087(SEL_{indoor} - 30)^{1.79} \quad (3.5)$$

where for this study the indoor SEL is assumed to be 20.5 dB(A) lower than the computed outdoor levels^[111]. The sound levels itself are calculated using the methodology as employed by INM version 7.0^[36].

A weakness of this relationship is that it predicts the response of an average person for a single event only. It does not predict what happens to a community during a night with multiple aircraft events. This has been addressed in later research, resulting in a method to calculate the chance to sleep through a night with several noise events^[2]. This method is however not perfect either: it results in a relationship that assumes that a person awoken once cannot be disturbed any further during the remaining part of the night. At the same time, the method also assumes that all noise events are independent, acknowledging at the same time that this might not be the case.

Eventually, the limitations of these kind of dose-response relationships do not have to be problematic, as long as the results are seen as indicators (or scores) and not as the actual number of awakenings. For comparative purposes, they can still be very useful and a significant reduction in a certain indicator can still be interpreted as an improvement. Along these lines of reasoning, the awakenings function is used as a general indicator for noise impact for all flights in this chapter,

regardless whether a flight under consideration may actually take place during the day or during the night.

Concerning the optimisation tool, minimisation with respect to the number of awakenings is achieved by making it part of the objective function. Generally speaking, fuel burn is always part of this function, not only because of its environmental and economic relevance, but also because it improves the stability of the optimisation tool. When awakenings are considered as well, the performance function becomes the weighted sum of the fuel burn and the number of awakenings:

$$J = m \cdot \int_{t_0}^{t_f} \sigma_{fuel} dt + k \cdot A \quad (3.6)$$

where the integral from the initial time (t_0) to the final time (t_f) over the fuel flow σ_{fuel} represents the total fuel burn, A is the number of awakenings and k and m are user-defined multiplication factors ($k, m > 0$). Please note that eventually it is the k/m -ratio that is important for the result, and not their absolute values.

3.2.3 Gaseous emissions as an optimisation objective

The NOISHHH tool has been fitted with an emissions inventory model and is therefore also able to perform trajectory optimisation with the purpose of minimising aircraft emissions^[70]. Because the tool is used for arrival and departure trajectories, the focus is on local emissions, not global emissions. This means that only pollutants emitted below the mixing height of 3000ft altitude are considered, also see Section 2.2.2.

The modelling is primarily based on the ICAO Engine Exhaust Emissions Data Bank^[8] and computes the emission of carbon monoxide (CO), hydrocarbons (HC), nitrogen oxides (NO_x). Sulphur oxides (SO_x) can be computed as well, but are modelled to be proportional to the fuel burn. Because the data available from the Emissions Databank is only valid under the test conditions, and therefore does not provide data for varying atmospheric conditions and intermediate thrust settings, a correction has been applied. This correction is based on the Boeing Method 2^[31]. This model allows for interpolating and correcting the Emissions Databank data for non-reference conditions, also see Section 2.2.2.

If all four pollutants are to be added to the performance index of the optimisation problem, their mutual significance needs to be determined. An obvious way to do this, is to look at the damage caused by gaseous pollutants to the human environment in terms of monetary values (externalisation). This approach has been followed for NOISHHH, based on the European ExternE project^[16]. It is then found that for a typical departure procedure, at least in monetary terms, all other remaining emissions are insignificant compared to those of NO_x. Furthermore, and also in monetary terms, the cost of emissions are very low compared to the cost of fuel^[61].

Based on this conclusion, local emissions have not been regarded any further

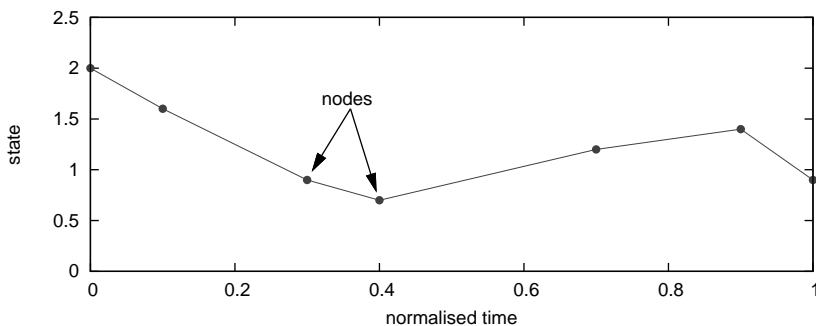


Figure 3.3: *The discretised trajectory and the nodes*

for the concept of custom optimised departure profiles. However, this does not mean that this could not be of interest for a future study. Very strict regulations could for example make local emissions of primary importance. In that case it is indeed possible to reduce local emissions, albeit at the expense of an increase in fuel burn^[61].

3.2.4 NLP profile optimisation model

The numerical optimisation method employed to solve the dynamic trajectory optimisation problem is the direct optimisation technique of collocation with Nonlinear Programming (NLP). The collocation method essentially transforms an optimal control problem into a NLP formulation by discretising the trajectory dynamics. To this end, the time interval of an trajectory solution is divided into a number of subintervals, as shown in figure 3.3. The individual time points delimiting the subintervals are called nodes. The number of nodes and their distribution over the time interval can be chosen freely.

The values of the states at the nodes are treated as a set of NLP variables. The same applies to the values of the (piecewise constant) controls between the nodes. The system differential equations are discretised and transformed into algebraic equations (implicit integration). The path and control constraints imposed in the original optimal control problem (to be specified later in Section 3.3.1) are treated as algebraic inequalities in the NLP formulation.

The transformation of the dynamic trajectory optimisation problem into an NLP formulation is performed by a software package called EZopt. EZopt also provides the NLP formulation of the problem to a numerical solver in order to generate the solution. This solution comprises the values of the states at the nodes and the values of the controls between the nodes. More details on this solver and the optimisation method in general are provided in Appendix B.1 through B.3. The model verification efforts are listed in Appendix B.7. This section continues

with the limitations of the method.

3.2.5 Limitations of the method

Although the developed tool is suitable to be used in a research environment, it is probably less suitable for an operational environment. One of the main problems stems from the nature of the optimisation method. For an arbitrary initial solution, it cannot be guaranteed that the final optimal solution can always be reached. And even if the solver will reach a converged solution, it cannot be guaranteed it will do so within a specified amount of time. Therefore the presented tool itself should not be relied upon in an operational environment.

Assuming that the solver reports to have found the optimal solution, the solution is not guaranteed to be the global optimum. This is a typical problem for non-convex optimisation. A method to get an impression about the local minima problem is to start a single problem using multiple distinct initial guesses and to compare the outcomes. Finding multiple solutions in such a case that are reported to be optimal is proof of the existence of local minima. Unfortunately, not finding this behaviour on the other hand does not guarantee that the optimum found is indeed the global optimum.

The original 3D version of NOISHHH has been tested in this way and the local minima problem has indeed been observed in a few cases. Inspection showed that the solver sometimes seems unable to make radical changes in the lateral part of the trajectory. The 2D version used for the custom optimised departure profiles has also been tested, and no local minima have been observed. However, as already mentioned, this is no guarantee that all of the found solutions are indeed global optima.

3.2.6 A fast and robust alternative to trajectory optimisation

A possible solution for the first limitation as identified in the previous section is to create a database of optimal profiles, such that a suitable profile for a specific flight is always available when required. However, there are many parameters involved in the profile optimisation problem. Some are discrete and lead to a limited, or at least controlled number of combinations like aircraft type, runway and departure route to be used. Others are continuous over a certain interval, like ambient temperature and take-off weight. These intervals need to be discretised when using this *database-of-solutions* approach, either leading to an impractical situation of an almost infinite number of solutions, or a rather coarse discretisation approach. And even though the latter approach may be improved by using interpolation in the n-dimensional space of the continuous parameters, there is still a certain risk present. Synthesising a trajectory from a number of other trajectories does not automatically guarantee that the result lies within the performance capabilities of the aircraft.

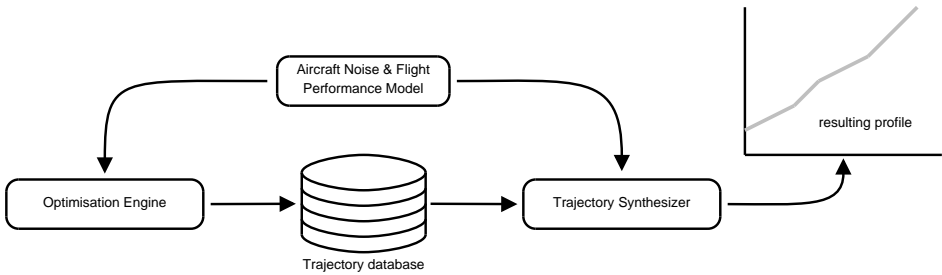


Figure 3.4: *The alternative process of optimal trajectory generation*

Therefore, the alternative approach proposed here is to construct a control vector by means of interpolation from the solutions database instead of using an interpolated state vector. The resulting state vector (the actual profile or trajectory) can subsequently be computed almost instantly using not the approximated but the actual flight parameters and the original optimisation tool point mass model. This effectively eliminates the risk of generating a profile that is outside the capabilities of the aircraft, although now it has to be confirmed that the trajectory meets all of the state-constraints. The overall process is illustrated in figure 3.4.

At this point, it is stressed that NOISHHH is not regarded as the actual software solution for this concept, but a possible implementation, primarily used to illustrate the concept. With that in mind, it is also important to realise that the reliability, stability and computing effort may differ among different optimisation methods, solvers, tools and computing hardware. This means that the need for an alternative solution generation method as proposed in this section may be dependent on the chosen implementation.

This need is also influenced by requirements of the (in-time) availability of optimal solutions for all flights. In other words, for what percentage of the flights a dispatcher would be satisfied with using an alternative in case the optimal profile is not available or late. Similarly, allowing for the use of sub-optimal solutions might also be an option, where sub-optimal solutions are solutions that show improved performance compared to standard procedures but are not guaranteed to be optimal. It may very well be conceivable to develop a reliable tool that is capable of generating such sub-optimal solutions within a predetermined amount of time. This has however not been investigated further.

3.3 Numerical examples of departure procedures

This section presents some results as generated by the optimisation tool for custom optimised departure profiles. Before showing results however, first it discusses the

constraints that the optimised departure profiles have to satisfy and discusses how the results will be compared to the current standard departure procedures. Then Section 3.3.3 shows results for a few particular flights, followed by an high-level overview of all results in the next section. The final three section discuss the influence of the final altitude constraint, the presence of wind fields, and the grid step-size chosen for the noise calculations.

3.3.1 Constraints, aircraft types and routing

The optimised departure profiles used in this study are based on the description as provided by a FAA advisory circular^[7]. Based on this description related to the so-called distant community NADPs, the primary constraints for the optimisation problem statement are:

- Acceleration for configuration clean-up should be commenced before thrust cutback is initiated
- Thrust cutback should not take place before reaching an altitude of 800 ft
- Acceleration for configuration clean-up is not allowed before reaching an altitude of 400 ft
- The aircraft is not allowed to descend at any time
- The aircraft is not allowed to decelerate at any time

The last three constraints are not mentioned in the circular, but added for reasons of safety and to obtain more realistic results. For reasons of simplicity, the model does not simulate the take-off roll phase of the departure. In light of the above mentioned constraints, this does not influence the result: optimisation is not started before an altitude of 400ft is reached. However, it does mean that the tool needs to be provided with non-trivial initial conditions. The initial conditions that are chosen are those that correspond to the situation where the aircraft reaches an altitude of 50 ft, shortly after take-off, as provided by INM. The velocity and position down the runway provided at this point for the optimised procedures do not differ in any way from the normal take-off procedures.

With respect to the final conditions, the procedure ends at a specified (route dependent) point at 6000 ft altitude and a calibrated airspeed of 250 kts. Because the final altitude and speed are fixed, all results have an equal energy state at the final point which allows for a fair fuel burn comparison. The speed limit itself is a typical indicated air speed limit below flight level 100. The 6000 ft altitude constraint is based on an initial flight level 60 restriction used at Amsterdam Airport for departures. Section 3.3.5 will look into the effects of this final altitude constraint on the results.

For NOISHHH, an aircraft model consist of an aerodynamic model, a thrust model, a fuel flow model, a noise model and optionally also an emissions model. Four different aircraft types are available for the optimisation tool:

-
- Boeing B737-300
 - Boeing B737-700
 - Boeing B737-800
 - Boeing B747-400

Naturally, much more aircraft types visit Amsterdam airport, but not all of the required models are available for these aircraft types. These four are however sufficient to show examples of optimisation results and also to show that optimal solutions do indeed differ from aircraft type to aircraft type.

Solutions have been generated for multiple departure routes based on actual departure routes from Amsterdam Airport. All existing departure routes for runway 24 have been modelled for use with the optimisation tool. Next to these seven, the four most south-easterly routes for runway 18L are available as well. Together these eleven routes cover the typical departure routes from Amsterdam airport when traffic is handled in southern direction. This will be discussed in more detail in Chapter 6.

3.3.2 Standard departure procedures as a reference

Apart from the freely optimised profiles, the optimisation tool has been configured such that it can replicate the current standard⁴ ICAO-A departure procedure at Amsterdam Airport as well. When using this procedure, the aircraft performs the initial climb at V_2 plus 10 knots^[36]. At 1,500 ft, the thrust setting is reduced from take-off thrust to climb-thrust, while maintaining the same calibrated airspeed. At 3,000 ft and still at climb thrust, the aircraft is accelerated and the transition towards the clean configuration is initiated.

For further comparison purposes, the counterpart of the ICAO-A procedure, called ICAO-B, has been modelled as well. The initial climb is again at V_2 plus 10 knots, but ends at an altitude of 1000 ft. At that point, the aircraft is accelerated such that the flaps can be retracted and thrust is reduced from the take-off setting to the climb setting. From 3000 ft, the aircraft is accelerated further towards the maximum permissible airspeed, typically 250 kts.

In principle, NOISHHH is allowed to optimise both of these procedures as well. However, the detailed specification of the procedures precludes most optimisation opportunities. Only for the part of the trajectories that is above 3000 ft, NOISHHH can determine the optimal trade-off between the thrust settings, the path angle and rate of acceleration.

The main advantage of having NOISHHH to generate the standard departure procedures is that the comparison of the results of standard and optimised procedures is more reliable. When using this approach, initial conditions, final

⁴As of April 2014, the ICAO-A procedure has been replaced by the NADP2 procedure as recommended procedure.




Figure 3.5: *The first part of the ANDIK1S departure from runway 24 and the ANDIK2E departure from runway 18L*

conditions, the flight mechanics model, the fuel burn model, and the noise model are all the same. This means that any differences in the results can positively be attributed to the differences in the departure procedures, and not to modelling issues.

3.3.3 Trajectory optimisation results of specific flights

To show the capabilities of the model, this section shows the profile optimisation results for a specific flight. The flight chosen here from all of the optimisation runs that have been performed is the ANDIK departure for a Boeing B737-300 at a take-off mass of 57500 kg. The ANDIK departure is available both from the 24 as well as from the 18L runway, and both results will be shown here. The first part of both routes is depicted in figure 3.5.

For reference purposes, both the ICAO-A and ICAO-B standard procedure will be shown, and compared to the optimised results. Three different optimised profiles have been computed⁵:

⁵for the definition of k/m, see equation 3.6 if required

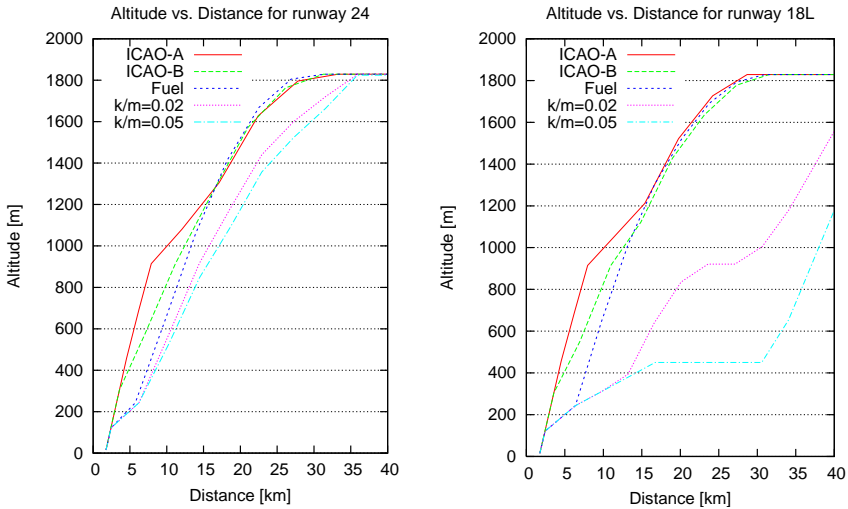


Figure 3.6: Altitude profiles for ANDIK departure from runway 24 and runway 18L

- one optimised for fuel only
- one optimised for both fuel and noise, $k/m = 0.02 \text{ kg}^{-1}$
- one optimised for both fuel and noise, $k/m = 0.05 \text{ kg}^{-1}$

Figure 3.6(a) shows the results with respect to the altitude profiles for the five different procedures for runway 24. The ICAO-A procedures shows a steep climb in the first part of the trajectory, where all excess power is used to increase altitude. The ICAO-B profile follows the same initial path, but start to accelerate shortly after reaching an altitude of 1000 ft, resulting in a less steep flight path angle. The fuel-optimised profile starts to accelerate at an even lower altitude, and this results in reduced climb performance, which is clear from the figure. However, all three profiles seem to converge for the later part of the trajectory, with the ICAO-B and fuel-optimised departure slightly outperforming the ICAO-A departure. The two noise optimised profiles clearly follow a lower path. This is caused by thrust cutbacks in order to reduce engine noise levels.

Figure 3.6(b) shows the same results, but now for runway 18L. When comparing the ICAO-A, ICAO-B and fuel-optimised profiles to those of runway 24 in figure 3.6(a), it appears that they are nearly identical. This is not unexpected, since the optimisation objective is not site-specific for these three procedures. The small variations that can be identified can be explained by differences in turns because of the different routing and possibly also small differences in flight path

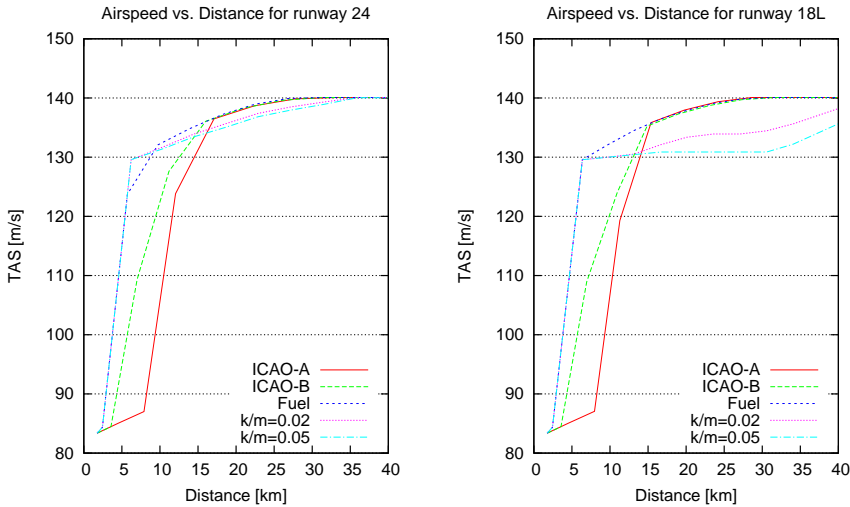


Figure 3.7: Speed profiles for ANDIK departure from runway 24 and runway 18L

segmentation⁶. The noise-optimised profiles however are clearly different for this runway, showing the effect of the site-specific criteria. Both noise-optimised procedures even have a horizontal part to reduce the required thrust and resulting noise level as much as possible, before resuming the climb to the final altitude (outside the range of the figure).

Figure 3.7(a) shows the speed profiles for the five different procedures for runway 24. From the figure it can be seen that the ICAO-A profile maintains a constant (calibrated) airspeed, while the three optimised profiles start acceleration almost immediately. The ICAO-B profile is somewhat in between. After approximately 17 km, all procedures reach their maximum speed of 250 KCAS, showing a slight variation in the depicted true airspeed because of differences in altitude at that point. Figure 3.7(b) for runway 18L shows similar results for the first part of the departure procedure, but the two noise optimised profiles have a lower true airspeed during the middle part of the trajectory. The maximum speed in terms of true airspeed is impacted by the much lower altitude of these two departures.

The absolute noise performance in terms of the number of awakenings (the noise performance index) and the relative performance compared to the baseline

⁶The number of nodes, and the resulting number of trajectory segments is the same for both routes. Because the total length differs among the different routes, the length of the segments will vary as well. This may lead to minor differences in results.

Procedure type	Awakenings for runway:		% reduction for runway:	
	24	18L	24	18L
ICAO-A	3135	3679	-	-
ICAO-B	2822	3419	-10.0%	-7.1%
Fuel	2345	2890	-25.2%	-21.4%
k/m = 0.02	2261	2032	-27.9%	-44.8%
k/m = 0.05	2248	1852	-28.3%	-49.7%

Table 3.1: *Noise results for example flight*

ICAO-A procedure are provided in table 3.1. The table contains the results for the ANDIK1S departure from runway 24, as well as the ANDIK2E departure from runway 18L. For both cases, the ICAO-A result leads to the highest noise index, with the 24 runway performing better than runway 18L. The ICAO-B departure reduces the noise performance indicator slightly and the fuel-optimised result in turn outperforms the ICAO-B procedure. With respect to the noise-optimised profiles, an interesting observation can be made. For runway 24, these two procedures perform only marginally better than the fuel-optimised departure. For 18L however, the numbers of awakenings can almost be reduced by 50% compared to the baseline.

This difference can also be explained when looking back at the figures 3.6(a) and 3.6(b). For runway 24, the noise-optimised solutions follow the fuel-optimised profile, albeit at a slightly reduced thrust setting to reduce noise impact. For runway 18L however, the strategy is different. Here both noise optimised profiles show a rather extreme measure: horizontal parts at low altitude, caused by very deep thrust cutbacks, before resuming the climb towards the final altitude.

The reason behind this difference in approach might be explained when looking back at figure 3.5. As can be seen from this figure, both routes overfly water after making two right turns. For the route from runway 18L, this is approximately 31 km after the start of the take-off roll. Now when looking at figure 3.6(b), this is also the point where the climb is resumed. Apparently, the optimal strategy is to keep the thrust setting as low as possible, until there is an uninhabited area where the thrust can be increased again without awaking much more people. For runway 24, this strategy does not seem to work. The distance to that same point is about 43 km instead of 31 km. It can be concluded that flying at low altitude for that distance either results in even more awakenings or a significant increase in fuel burn.

Finally, the ten different procedures are compared in terms of fuel burn. Table 3.2 provides the results, both in absolute numbers as well as the relative perforce compared to the ICAO-A departure. First of all it is readily clear that the overall fuel burn for runway 24 is higher. This can be explained by the fact that the ANDIK departure from runway 24 is more than 10 km longer. When comparing

Procedure type	Fuel burn [kg]		% reduction	
	for runway:		for runway:	
	24	18L	24	18L
ICAO-A	653	580	-	-
ICAO-B	641	566	-1.8%	-2.4%
Fuel	616	542	-5.7%	-6.6%
$k/m = 0.02$	616	549	-5.7%	-5.3%
$k/m = 0.05$	616	555	-5.7%	-4.3%

Table 3.2: *Fuel burn results for example flight*

the ICAO-B departure to the ICAO-A, it can be seen that the fuel burn is approximately 2% less. The most obvious explanation is that for the ICAO-B procedure, the aircraft attains a clean configuration earlier in the procedure, leading to less drag and therefore less fuelburn. This effect is even stronger for the remaining three solutions. When comparing the optimised results for both runways, another interesting observation can be made. For runway 24, there is no significant difference in fuel burn between the three optimised procedures, as there was little difference in awakenings. For runway 18L however, the low-altitude stretches clearly result in a fuel penalty for the noise optimised profiles when compared to the fuel-optimised one.

From the results it is clear that even though both examples use the same aircraft at the same mass and towards the same Terminal Area (TMA) exit point, especially the noise-optimised solutions are clearly different. In this case, these differences are caused by the differences in population density in areas that both aircraft overfly. In the first place, this example emphasises the importance of using site-specific criteria instead of generic criteria when evaluating noise impact. Secondly, it is a clear example of why it is useful to use custom profiles instead of generic Noise Abatement Departure Procedure (NADP)s. Such generic procedures will never be able to exploit the actual local circumstances in order to reduce noise impact.

3.3.4 Overview of trajectory optimisation results

The examples shown in the previous section are only ten of more than 500 trajectories generated for this study. Runway, departure route, aircraft type, take-off mass and performance index are varied among all of these results. Appendix B.6 gives a complete overview of all computed departure profiles. What the two examples point out is that the results and especially the effectiveness varies between different departure routes. This also holds for the other parameters that have been varied for this study. Therefore table 3.3 provides a high level overview of the results in terms of noise performance.

From the table it can be seen that for both runways, the ICAO-A procedure leads to the highest number of awakenings, followed by the ICAO-B procedure.

	ICAO-A	ICAO-B	Fuel	Noise (0.02)	Noise (0.05)
Runway 18L	5001	4574	3482	2944	2752
Runway 24	5203	4780	3982	3417	3197
Total	5135	4712	3815	3259	3049

Table 3.3: Average number of awakenings for all profiles

	ICAO-A	ICAO-B	Fuel	Noise (0.02)	Noise (0.05)
B737-300:					
ANDIK2E (18L)	3504	3238	2775	1984	1819
ANDIK1S (24)	3082	2769	2237	2169	2155
B737-800:					
LOPIK2E (18L)	3266	3014	2321	1830	1694
LOPIK1S (24)	2740	2426	1782	1632	1444
B747-400:					
ARNEM2E (18L)	8830	8077	5867	5615	5392
ARNEM1S (24)	7292	6522	5122	5067	5018

Table 3.4: Average number of awakenings for three selected route and aircraft type combinations

The three optimised procedures present a better score, completely in line with examples presented in Section 3.3.3. Also remarkable is the fact that the noise index is lower for runway 18L, while runway 24 is generally regarded as a noise-preferential runway.

A closer inspection of the underlying data can explain this observation. It turns out that one of the departure routes (the SPYKER departure from runway 24) leads to a much higher number of awakenings than all other routes, irrespective of the aircraft type. As this route is only available for runway 24, it distorts the comparison. Simply excluding this route reduces the averages for runway 24 by approximately 20%, resulting in a noise performance that is indeed better than that of runway 18L. For a better comparison, table 3.4 presents the noise performance for a selected number of route and aircraft combinations for both runways.

This table confirms the two earlier observations concerning the procedure type and resulting awakenings, and also the lower impact when using runway 24. Table 3.4 also shows a much higher impact for the ARNEM B747-400 departures. Inspection of other results shows that this is not route-related, but is caused by the type of aircraft: the B747-400 results in much higher awakening scores than the other types. This is not unexpected, as the noise levels of this aircraft are indeed higher.

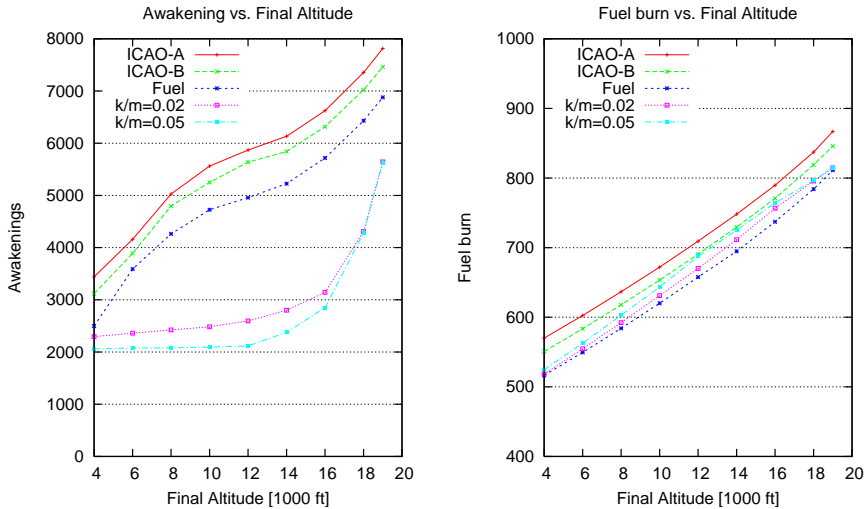


Figure 3.8: *Awakenings and fuel burn as function of the final altitude for the VALKO1S departure*

3.3.5 Effects of the final altitude constraint

As already mentioned, all optimised profiles use a final altitude of 6000 ft. It is expected that the final altitude influences the fuel burn and may also have an effect on the number of awakenings. To study these effects, the final altitude is varied between 4000 and 19000 ft⁷ for the VALKO1S departure from runway 24 for a Boeing B737-800 at 71000 kg. The awakenings and the fuel burn are plotted versus the final altitude in figure 3.8(a) and 3.8(b) for all five types of procedures.

When analysing the number of awakenings, two different patterns can be distinguished. The three profiles that are not optimised with respect to noise show a steady increase of awakenings with an increase of final altitude. In order to minimise fuel burn, all three have the tendency to climb to the final altitude as fast as possible and from there to fly to the end of the route. Increasing the final altitude for these types of profiles leads to a longer period with a high thrust setting, explaining the increase in awakenings. The two noise optimised profiles however hardly show any increase in awakenings, unless the final altitude is higher than 12000ft.

When looking at the route itself, it can be seen that the VALKO1S departure has a part over land, followed by a part over the North Sea. As it turns out,

⁷Starting at 4000 ft, the final altitude was increased in steps of 2000 ft. At 20000ft, the problem becomes infeasible, so 19000 ft was added as an additional data point




Figure 3.9: *Altitude profiles for different final altitudes*

it is possible to increase the final altitude to 12000ft without changing much to the part of the profile that is over land: most of the additional climbing can be performed when already over sea. This can also be seen in figure 3.9. The figure in fact shows the application of a location-based thrust cutback, that is tailored to the constraints and objective function of the problem. This is a similar strategy as seen for the ANDIK2E departure from Section 3.3.3 and allows the solver to keep the number of awakenings down. However, as soon as the final altitude is higher than 12000 ft, thrust over land must be increased to reach the final altitude and the number of awakenings starts to rise more and more rapidly.

The influence of the final altitude on the fuel burn can be seen in figure 3.8(b). Naturally, the fuel burn increases with an increase in final altitude, as more power is required for the climb. The ranking of the 5 different procedures with respect to fuel burn is however not dependent on the final altitude, and equal to the ranking seen in the previous section. However, a close inspection reveals that the noise optimised profiles show a fuel burn that is close to that of the fuel optimised profiles for very low and very high final altitudes. However, for the final altitudes around 12000 ft, the fuel burn is significantly higher. Apparently, the very high reduction in number of awakenings that can be achieved in this area does also increase the fuel burn.

The effect of changing the final altitude has not only been investigated for the VALKO departure, but also for the ARNEM departure from the same runway. This route has been chosen because there are no large uninhabited areas underneath this route. The route is also much shorter and additionally, the take-off mass is increased from 71000 kg to 74500. All three differences are expected to reduce the opportunities for a noise-beneficial thrust regime. The results are displayed in figure 3.10(a) and 3.10(b). The final altitude is varied between 4000 and 14000 ft. At a final altitude of 15000 ft and higher, the problem becomes infeasible.

The results show that there is indeed very little room for reducing the number

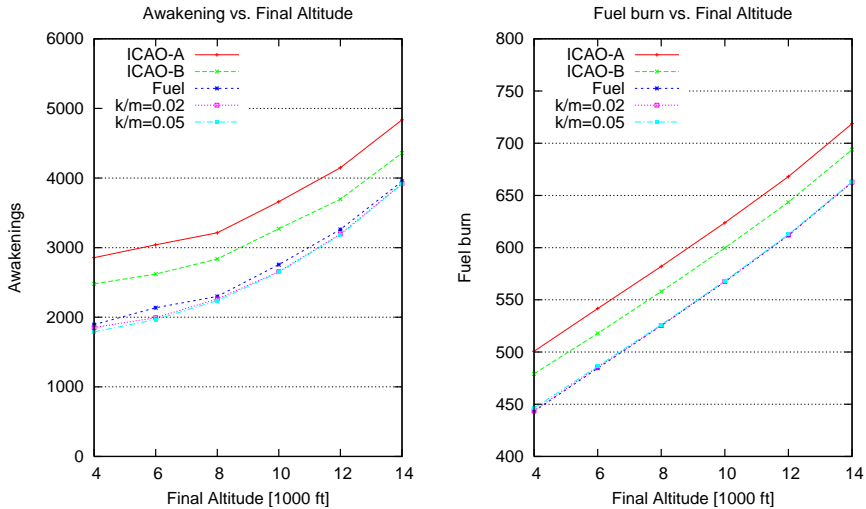


Figure 3.10: *Awakenings and fuel burn as function of the final altitude for the ARNEM1S departure*

of awakenings relative to the fuel-optimised profile. Not surprisingly, there is also hardly any difference in fuel burn between the three optimised profiles. In fact, the three optimised profiles are just very similar in this case. More in general, the fuel burn and the number of awakenings rise with an increase in final altitude, both as expected and also seen with the VALKO example. However, the offset between the different types of procedures is fairly constant with the increase in final altitude.

Comparing the results for both routes, it can be concluded that the possibility of reducing the noise impact relative to the fuel-optimised departure is highly dependent on the local conditions. The results also show that this relative reduction in noise impact may indeed be influenced by the final altitude constraint, as seen for the VALKO departure, but also that this is not always true. Once again, it is clear that there is no general strategy that can be applied for maximum noise abatement.

A final note on the increase of awakenings with increasing final altitude: from the figures, it could be concluded that it is beneficial for noise considerations to keep the final altitude as low as possible. However, the results shown here only count the awakenings until reaching the terminal point of the departure route. An aircraft arriving at this point at a very low altitude will need to climb further beyond this point and may cause additional awakenings there. Therefore, drawing this conclusion is not recommended.

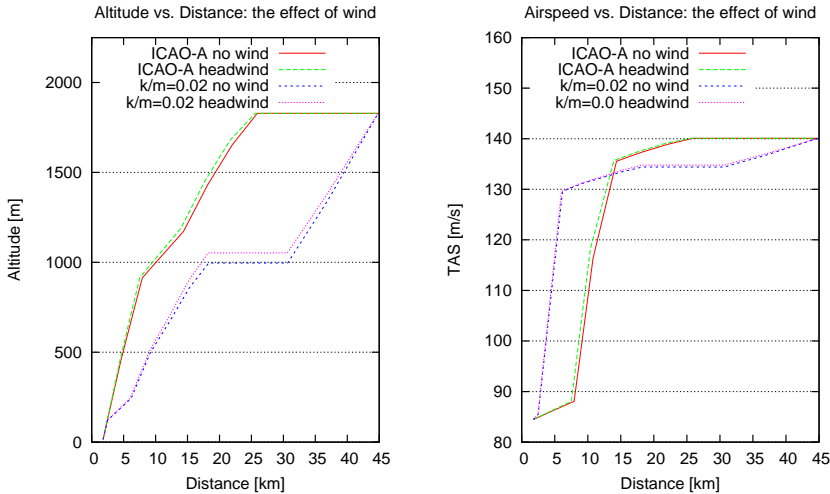


Figure 3.11: *The effect of a headwind on altitude and speed*

3.3.6 Effects of the presence of wind fields

All of the previous results are based on a no-wind situation. This is however rarely the case in reality. Because wind has an influence on for example the ground speed, climb performance and the resulting noise computation, it is expected that this will influence the trajectory optimum in particular cases. To study these effects, the aircraft performance model in NOISHHH has been adapted to allow for wind input. The adaptation itself is discussed in more detail in Appendix B.5, and this section shows the results for a particular case. Using the same VALKOIS departure as in the previous section, the default no-wind situation is compared with a situation where the wind is blowing from direction 240 (between Southwest and West-southwest) for all altitudes. This results in an almost direct headwind for this route with a wind speed that increases from five knots at ground level to 15 knots at 3000 ft and above.

Figure 3.11(a) presents the ICAO-A profile and the noise-optimised profile for $k/m=0.02 \text{ kg}^{-1}$, both for the situation with and without the headwind. The other six profiles are not shown here as this would clutter the figure too much. In both case, the profiles that are based on head-wind climb steeper, although the effect seems very limited for the ICAO-A profile. Figure 3.11(b) shows the results for true airspeed for the the same four profiles. Although the speed profiles also show some effect of wind, the differences between the profiles with and without wind are minor.

Since the profiles are influenced by the wind, it can be expected that the

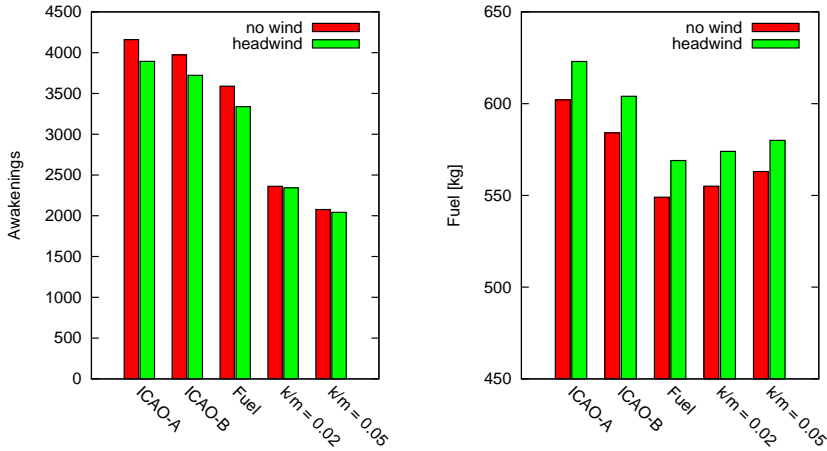


Figure 3.12: *The effect of a headwind on awakenings and fuel burn*

number of awakenings and total fuel burn will be impacted by the wind as well. Figure 3.12(a) first shows the effect of the headwind on the number of awakenings. For all five profiles, the presence of the wind reduces the number of awakenings, although this effect is very small for the two noise-optimised profiles. Apparently, the relatively large change in profile for the noise-optimised solution as seen in figure 3.11(a) does not automatically lead to a relative large change in number of awakenings. Figure 3.12(b) shows the effects on the fuel burn. As expected, fuel burn increases when the aircraft encounters a headwind.

The results show that it can be important to take the wind situation into consideration in order to obtain the actual profile optimum at any specific point in time. It should be realised, however, that the effect of wind on the aircraft performance is only one of the effects the weather may have on the location of the optimised profile^[72]. Wind for example also influences noise propagation, and temperature is a factor in both aircraft performance and the propagation of noise through the atmosphere. Ideally, all of these factors should be taken into consideration, although this probably increases computational time considerably and simultaneously also increases the number of variables for the database approach discussed in Section 3.2.6. Eventually, it is expected that some sort of trade-off between level of detail and feasibility would be required to decide on this issue.

3.3.7 Effects of noise grid step-size

Both the custom optimised profile generator and the original version of NOISHHH use a rather coarse step size of 1 km for the noise grid. Considering that the noise

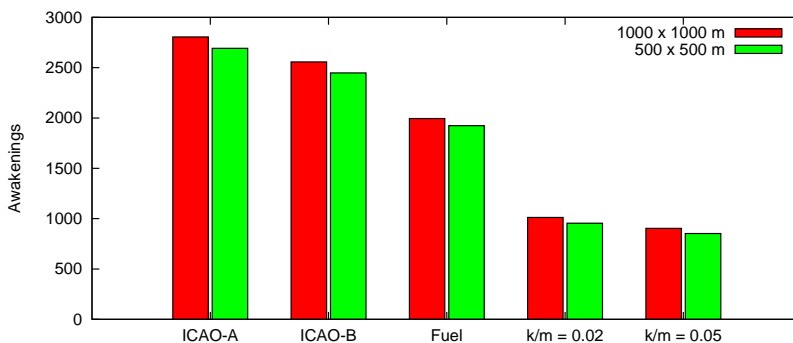


Figure 3.13: *The effect of the noise grid step size on the objective value for noise for the BERGI1S departure*

evaluation is the most computationally intensive process during an optimisation run, the decision to use this coarse noise grid reduces overall computational time considerably. It is recognised however that there may be an impact on accuracy. A sensitivity study is performed, primarily to determine to what extent the noise optimised profiles change when using a finer mesh. To this end, some of the optimisation runs have also been performed using a noise grid with a step size of 500 m. The results that will be compared here are for a BERGI1S departure from runway 24, using a Boeing 737-700 at 66500 kg.

The noise results for the fine mesh and coarse mesh are provided in figure 3.13 for all five profiles. The results indeed confirm that the chosen noise grid step size influences the results, as the awakenings are between 3-6% lower for the fine mesh. Here it is important to realise that for both ICAO procedures and the fuel-optimised procedure, this difference can only be the result of the noise calculation itself, as it cannot be caused by differences in the resulting profiles, because noise does not influence the profile optimisation process for these three procedures.

For the remaining two profiles that are noise-optimised, the change in step size could result in a different profile optimum. A first check reveals however that the transit time and the fuel burn do not differ between the two meshes. A further inspection confirms that the profiles are in fact identical. For this case this means that although a finer mesh does seem to improve the accuracy of the calculated absolute noise response by a few percent, the relative noise performance among the different profiles remains rather constant. Even more relevant, the finer mesh does not lead to a different optimised profile and therefore does not seem to justify a considerable increase in computational time.

3.4 Estimating the runway capacity effects

One of the observations made in the previous section is that the speed profile of the optimised trajectories differs significantly from the current situation, particularly in the first few thousand feet of the climb. The velocity and acceleration behaviour determines the resulting spacing between consecutive departures, and this difference could impact the runway departure capacity. Since runway capacity is an important consideration for busy airports, this section looks into these capacity effects. This includes the effects that can be expected when mixing several types of departure procedures.

3.4.1 Capacity simulation method

Runway departure capacity is ultimately a function of the required separation distance between different flights. In practice however, separation standards between two consecutive flights may not be distance-based but time-based, provided in minutes and based on differences in speed and routing^[20]. Especially in absence of radar facilities, counting down may indeed also be more practical than to aim for a certain spatial separation. However, the objective of this capacity simulation is to identify changes in airfield capacity due to the changes in aircraft operation procedures, particularly due to the changes observed in speed behaviour. This means that simply using the time-based intervals will not work, because then capacity is only dependent on the aircraft mix, not on changes in airspeed.

Instead, the capacity analysis method to study the effects of the optimised departures will use the resulting separation distance between different flights. Provided with a separation requirement, a search algorithm is used to find the minimum interval that is required between specific flights in order not to violate the separation requirement. This interval is subsequently used to calculate the theoretical runway capacity. It is important to realise that the theoretical capacity is not attainable in practice and therefore the results of the capacity analysis results should not be compared to capacity declarations. The numbers can, however, be used for a relative assessment of the different cases that are presented here.

General simulation method description

Beforehand, it was expected that variations in parameters such as aircraft type, weight and routing may influence the runway capacity, especially in conjunction with the departure sequence. A wide range of capacity simulations has been performed to identify the main drivers and the extent of the variation in departure capacity. A Monte Carlo approach was used to generate the input for all of these experiments. The overall process is depicted in figure 3.14.

First the Monte Carlo module selects a number of eligible trajectories from a database of departure trajectories. These trajectories are provided to the search




Figure 3.14: Monte Carlo procedure for calculating departure capacity

algorithm in a random order, to form a departure sequence. The capacity simulation uses a search algorithm to determine the minimum interval between each of the flights and calculates the resulting runway capacity for that specific situation. This process is repeated as often as desired.

The next three subsections provide more detail on the three elements used for this capacity simulation. First, more details are provided on the database of departure trajectories, followed by a discussion on the Monte Carlo module. The last subsection presents the search algorithm that determines the minimum safe interval and calculates runway capacity.

A database of departure trajectories

NOISHHH was used to generate a database of optimised departure trajectories. Several parameters were varied to obtain the set of distinct solutions. However, to keep the number of unique combinations in check, the variation in the parameters was restricted. First of all, four different aircraft types were used: the B747-400, B737-300, B737-700 and the B737-800. For each of these aircraft types, six different take-off weights were selected, loosely based on the available take-off weight of the respective types in the INM database. Next, the departure routes were varied based on six different SID's from runway 24 at Schiphol airport in The Netherlands. Finally, four different sets of optimisation criteria have been used, as indicated in table 3.5. For this experiment, the mixing of aircraft types, weight, and SID resulted in 54 unique combinations. All 54 combinations have been computed for each of the four optimisation criteria sets, resulting in a database of 216 distinct trajectories for the runway under consideration. A list of all trajectories that have been used can be found in Appendix B.6.

The results as generated by NOISHHH are unfortunately not directly suitable for usage in the capacity simulations. One of the incompatibilities is that the optimisation tool generates trajectories that do not include the take-off roll. As mentioned before, all NOISHHH trajectories start at an altitude of 50 ft (15 m). Because it is desirable for the traffic simulation to simulate the ground run part as well, all optimised trajectories are extended with a ground run and initial

Set	Fuel Factor m	Noise Factor k	ICAO-A Constraints
1: ICAO-A	1.0	0.0	Yes
2: Fuel optimised	1.0	0.0	No
3: Noise low	1.0	0.02	No
4: Noise high	1.0	0.05	No

Table 3.5: *Sets of optimisation criteria*

climb-out. This extension is based on the assumption of linear acceleration during the missing initial part of the trajectory. Using this assumption and given the location of the brake-release point on the runway and the location and the velocity of the start of the optimised trajectory, the duration of the take-off roll and the corresponding acceleration can be computed.

A second problem with the results generated by NOISHHH is that they only provide the aircraft states at a limited number of nodes, as can be seen in the figures from the previous section. Some of the nodes can be large distances apart and this was considered inconvenient for use with the search algorithm. Therefore the database was populated with computed trajectory data providing the interpolated state of the aircraft in 1-second intervals, based on the original trajectory data as generated by NOISHHH.

Monte Carlo experiments

Because it is expected that the runway capacity is dependent on the chosen subset of trajectories from the database and their actual sequence, a Monte Carlo experiment is designed to study the runway capacity under these varying conditions. On a high level, the Monte Carlo experiment first selects a number of trajectories (typically 20) from the database and generates a take-off sequence as input for the capacity simulation. Next, it runs the capacity simulation, saves the capacity results, and repeats the whole process as often as specified, typically a hundred times.

The Monte Carlo algorithm can be instructed to use only a subset of the database of solutions. This property is used to compare the results based on the different sets of optimisation criteria. For example, it can be specified that the algorithm should only select ICAO-A departures (set 1) from the database. Multiple subsets can also be made eligible. One of the results in the next subsection, for example, will be based on traffic from only set 1 and 2.

The algorithm also uses subsets of the database to make a distinction based on the wake vortex categories of the different flights. This property is used to get a certain traffic mix between medium and heavy aircraft for each of the capacity runs. The specified mix however is not enforced, but it is used as a probability in the drawing process. In the results presented here a 30-70 percent heavy-medium ratio was used. For the typical capacity run using 20 flights, this means that there is an expectancy of six out of twenty being heavy. The actual mix for a

certain run may of course vary and the number of heavies observed during all of the experiments varied between none and as much as thirteen.

Runway capacity calculation

The departure capacity simulation is based on a basic conflict search method. First of all the departure sequence is provided by the Monte Carlo algorithm, using an id for each departure. The corresponding trajectory data is then retrieved from the database. The first aircraft in this sequence is assumed to start the take-off at time $t=0$. The next aircraft (number two in the sequence) is positioned a short while behind the first one (e.g. take-off roll starting at $t=20s$.) and a conflict search is initiated along their tracks. If at some point in time a horizontal loss of separation is detected, the latter aircraft (number two) is delayed for one second and the conflict search is repeated, until all possible conflicts are eliminated. When the first two are conflict free, the time of take-off for the second departure is frozen and the third departure is positioned shortly behind the second one. Once again, delay is added until all conflicts are resolved. This process is continued for all remaining aircraft in the sequence.

Some of the potential conflicts may be resolved naturally when two consecutive departures have different and quickly diverging departure routes. Although this is fine, this introduces a small chance that a departure two positions behind another one may gain too much on the first one. For example, if the last of three aircraft is allowed to commence the take-off roll when separated from the middle one because of route divergence, it might still catch up with the leading one if these two aircraft are on the same route. This situation was identified during initial verification checks of the capacity model (see Appendix B.7) and should be prevented. Therefore, starting from the third aircraft in the sequence it is made sure that aircraft n is not only conflict free with aircraft $n-1$, but also with aircraft $n-2$.

When the last aircraft in the sequence has been positioned, the average inter-departure interval is calculated as the departure time of the last aircraft divided by the number of aircraft in the sequence minus one. Finally, the runway capacity in departures per hour is obtained by dividing 3600 seconds by the calculated inter-departure interval.

One final issue that has not been discussed is the separation standard to be used for the conflict search. As already mentioned, the time-based criteria cannot be used for this method. This means that distance-based criteria have to be used, but their exact definition may have a significant influence on the results. To reduce this sensitivity, three different sets of separation standards have been used for the experiments, as presented in table 3.6.

As can be seen, the first two sets use a standard separation distance of 3 nm, with the difference the required separation during the take-off roll. Set 2 uses a strict definition and does not allow the take-off roll to commence, before the preceding aircraft is already 3 nm down from that point. This minimum

Set	Separation during take-off roll	Horizontal separation
1: Airborne only	Not required	3 nm
2: Strict	Required	3 nm
3: Variable (WTC)	Not required	wake class dependent

Table 3.6: *Sets of separation standards*

separation is also maintained for as long as both aircraft are on their departure route. The strict definition of separation however results in a considerable head start for the preceding aircraft, so actively maintaining separation by delaying departure time of the follower is usually not required. Separation set 1 allows for anticipation and makes sure that the required separation is available before the trailing aircraft becomes airborne. In this situation, maintaining separation during the remainder of the trajectory is more of an issue. Set 3 also does not require separation between airborne and rolling aircraft, but in addition uses wake vortex class dependent separation distances^[94]. This means 5 nm for a medium trailing a heavy, 4 nm for a heavy trailing a heavy and 3 nm for the two remaining combinations.

3.4.2 Runway capacity simulation results

For each of the experiments, the Monte Carlo algorithm was instructed to select trajectories from a different subset in the trajectories database. These six subsets have been used for the different experiments:

- A: ICAO-A, also provides the baseline capacity
- B: Fuel optimised, noise factor k set to zero (see equation 3.6)
- C: Low noise, a combination of fuel and noise with a k/m -ratio of 0.02 kg^{-1}
- D: High noise, a combination of fuel and noise with a k/m -ratio of 0.05 kg^{-1}
- E: Mix of ICAO-A and Fuel, a combination of A and B
- F: Mix of all, a combination of A,B,C, and D

For each experiment, the MC-algorithm generated 100 departure sequences of twenty aircraft based on the active subset of trajectories and all experiments were performed three times to allow for the three different definitions of separation. This means that in total 1800 runway capacity runs have been performed and all of them are shown in the next figures. Figure 3.15(a) first presents the results for experiment A. The vertical axis shows the departure capacity, and the horizontal axis shows the ratio of heavy aircraft. Each capacity run result is shown using a single marker, based on the capacity result and the fraction of heavies of that run. Marker style is used to make a distinction between the three different definitions of separation standards. The figure next to it shows the same, but then for experiment B.

When analysing the baseline results from experiment A in figure 3.15(a), a number of observations can be made. Clearly, the first set of separation standards (airborne only) results, as expected, in a much higher capacity than the second

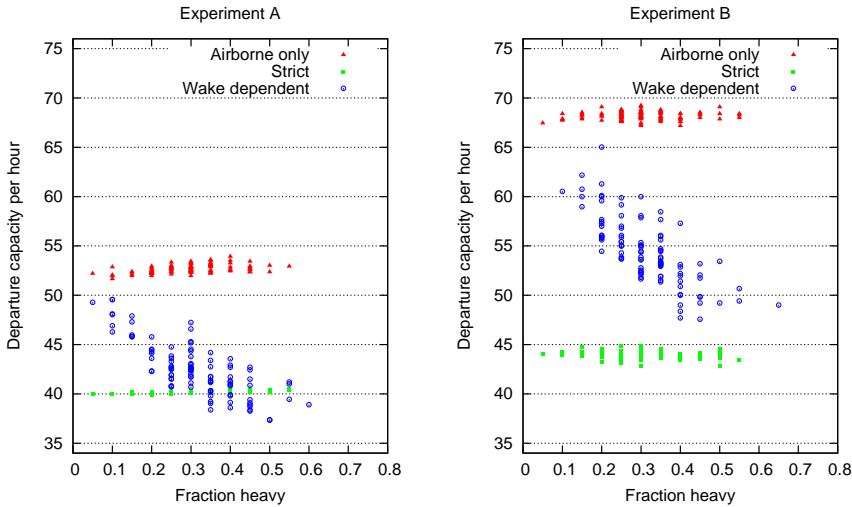


Figure 3.15: Runway capacity results for experiments A and B

one (strict). For both results, the actual mix between medium and heavy aircraft does not seem to influence capacity. This is not the case for the third set of separation standards. Again, this is not unexpected, as the separation distances are influenced by the composition of the traffic. There is clearly a negative relation between the fraction of heavies and the runway capacity. At a low number of heavies, the capacity approaches that of set one, and for a high fraction of heavy aircraft, the capacity even falls off to levels below that of set two. Apart from the (average) capacity, the spread in results is important as well, as this influences the reliability and predictability of the runway capacity. The standard deviations and averages for experiment A and all other experiments are provided in table 3.7.

As was already clear from figure 3.15(a), the variation in capacity for set 2 is extremely low. The individual data points for the capacity runs can hardly be distinguished. It can be concluded that traffic composition, departure sequence, nor route selection results in significant capacity variations. Apparently, the large initial separation as caused by the strict definition that prevents any conceivable conflict during the remainder of the trajectory. This means that it is not necessary to add further delay even if, for example, the following aircraft is faster than its predecessor. This means that the actual departure sequence is therefore not really important and the variation becomes very small.

For set one the standard deviation increases to 0.42, still a reasonably low number. The actual departure sequence and route choice becomes somewhat

Experiment	Separation standards					
	Set 1		Set 2		Set 3	
	μ	σ	μ	σ	μ	σ
A: ICAO-A only (baseline)	52.7	0.42	40.2	0.12	42.5	2.73
B: Fuel optimised	68.2	0.42	44.1	0.39	54.5	3.43
C: Noise low	68.2	0.45	44.1	0.47	55.2	3.72
D: Noise high	68.1	0.42	44.0	0.41	54.5	3.69
E: Mix of ICAO-A and Fuel	60.3	2.00	42.4	0.54	49.5	2.85
F: Mix of all available	62.4	2.03	43.1	0.54	51.2	3.33

Table 3.7: *Sets of separation standards*

more important, especially if there are trajectories with speed profiles with substantial differences. For set three, the situation is completely different and the standard deviation increases to 2.73. For this set, the departure sequence becomes a major factor. Of course, the actual sequences in combination with speed profile incompatibilities still may have some influence, but in this case the departure sequence also influences the required average separation because of the wake vortex considerations.

The results for experiment B (fuel optimised trajectories) are provided in figure 3.15(b). At first sight, the results show a similar pattern. Set one shows a horizontal group of high capacity numbers, set two shows an also horizontal group of much lower capacity, and finally group three is again in between and clearly a function of the fraction of heavy aircraft. However, if the results of experiment A and B are compared, it is readily clear that the overall capacity is much higher for experiment B. For set two, the difference in average capacity is about 4 departures per hour, but for set one and three this increases to 15 and 12 respectively. This can only be explained by looking at the speed profile as discussed in Section 3.3. All departures in experiment B start the acceleration as soon as possible, whereas all departures in experiment A remain at constant calibrated airspeed until reaching 3000ft. Higher speeds leads to covering the same distances in less time, so the earlier acceleration for the fuel optimised trajectories also reduces the inter-departure intervals.

Figure 3.16(a) and 3.16(b) show the results for experiments C and D, both optimised for fuel burn as well as for noise impact, using two different weighting factors in the objective function.

Both figures are also almost identical and both results are also similar to those of experiment B. This is also confirmed when looking at the averages and standard deviations. Differences can be identified, but these are very small. This means that increasing the k/m-ratio for the optimisation problem does not result in differences in speed behaviour that have a significant influence on the runway capacity. This observation is valid for all three definitions of separation.

Based on the results until now it can be concluded that the optimised trajectories result in a higher runway departure capacity than the ICAO-A departures




Figure 3.16: Runway capacity results for experiments C and D

and that the choice between fuel-optimised and noise-optimised trajectories does not have a significant influence on departure capacity. This conclusion is advantageous for the concept of custom optimised departure trajectories. However, these results are only valid for a situation in which all aircraft use the optimised profiles. Since the execution of these profiles will probably have to rely on onboard automation systems, this may not be very realistic, or at least not for the near future. Therefore, more experiments have been performed, in which some aircraft still use the ICAO-A departure, while the remaining aircraft use the optimised departures. The results are provided in figures 3.17(a) and 3.17(b).

For experiment E, the Monte Carlo algorithm selects at random from the ICAO-A and the fuel optimised results. This means that the average ratio between the two types should be 50-50, but again this may vary for a specific capacity run. For experiment F, all four subsets are available, so the average percentage of traditional departures should be 25% and 75% for the remaining three groups.

Generally speaking, the capacity for the two experiments that involve mixing is in between that of experiment A and the three optimised ones. For all three sets of separation standards, experiment F scores better than experiment E. Apparently the ICAO-A departures are indeed the limiting factor. The variation in results for the last two experiments is evidently higher than for the previous four, at least for the first two sets of separation standards. This is already clear from the figures, but it can also be confirmed when inspecting the standard deviation. This can most likely be explained from the presence of the two different types of procedures.

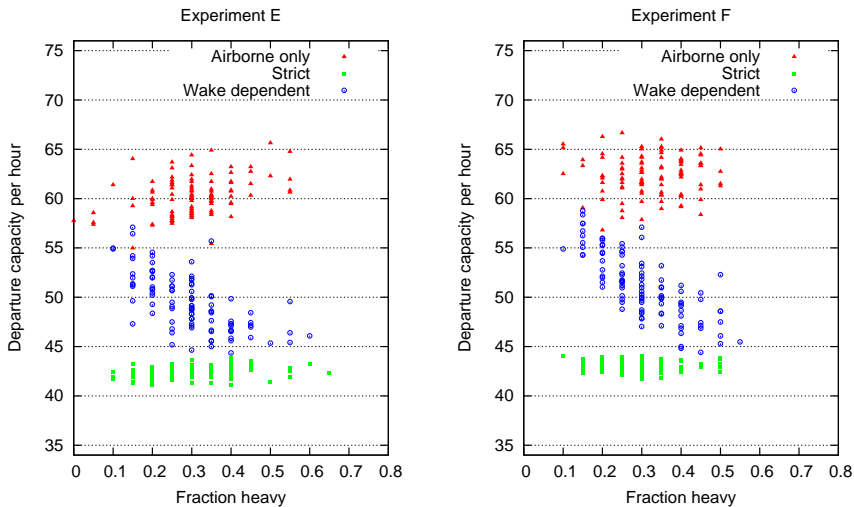


Figure 3.17: Runway capacity results for experiments E and F

First of all the actual ratio between ICAO-A and optimised departures for a particular capacity run may differ, but on top of that the departure sequence with respect to procedure type becomes important as well.

Probably the most important conclusion from these results should be that switching to optimised departures or mixing traditional and optimised departures procedures does not lead to a reduction in runway capacity when compared to the baseline situation. This means that the concept of custom optimised departure profiles does not exhibit a disadvantage in this area. Based on the observed required separation times, it could even lead to a higher runway capacity if this is desired. A second main conclusion should be that when using a mix of traditional and optimised departures, the process of optimising the actual departure sequence becomes even more advantageous than it is today for situations where runway capacity is of critical importance.

3.5 Discussion

This chapter presented the concept of the custom optimised departure profile, outlining the situation where each departing aircraft would execute a departure profile that has been optimised for that particular flight, while still using predefined and published departure routes. Although it is recognised that not optimising with respect to the routing is likely to reduce the environmental benefits, the concept does not suffer from a considerable increase in airspace complexity. Especially

this benefit could be of paramount importance for a potential implementation.

An existing three dimensional trajectory optimisation tool was adapted to generate the optimised departure profiles. Several of the presented numerical examples of these optimised profiles showed small reductions in fuel burn and impressive reductions with respect to the applied indicator for community noise when compared to two different standardised departure procedures. A summary of results for over 500 optimised profiles confirmed that the examples are representative for other cases as well. The resulting database holding all of the generated optimised departure profiles will be used for the multi-level integrated environmental management system of Chapter 6.

A runway capacity study has been performed, because the optimised profiles trajectories show a speed profile that differs significantly from the current situation. Based on the results of this study, it is concluded that switching to optimised departure profiles or mixing traditional and optimised departures procedures does not lead to a reduction in runway capacity when compared to the baseline situation and that there are opportunities to even increase throughput at the departure runway.

The tool that has been developed to illustrate the concept could certainly be extended or improved upon. As a first example, and as mentioned in section 3.3.6, the aircraft performance model was adapted to deal with the effects of wind. As a further improvement⁸, the noise model could also be replaced with a noise model that corrects for wind and temperature effects^[26]. Furthermore, all of the results based on noise-optimisation are based on the awakenings dose-response relationship as proposed by the FICAN. The model could be extended with other dose-response relationships, also to investigate the effects of using different noise objectives on the resulting optimised profiles.

Finally, it should be mentioned that this chapter was directed at departing aircraft only. This does not imply that there are no opportunities for optimisation at the trajectory level for arrivals, as will be discussed further in Chapter 6. The next chapter will not continue with the concept of the custom optimised departure profiles. Instead, it will present a second optimisation concept, corresponding tool and set of results. Based on the classification as defined in figure 1.1, the tool and the concept belong to the next higher level of aggregation: the operational level.

⁸Recently, this has already been studied, albeit for arriving flights^[3]

Reducing environmental impact as objective in arrival management

Similarly to the previous chapter, this chapter^[64] deals with a promising mitigation strategy and presents the design of a corresponding tool that can be used separately, but selected elements will later on also be used for the demonstration of the multi-level integrated system. Compared with the concept at the trajectory level shown in the previous chapter, this decision support system belongs to a higher level of aggregation, because it involves a group of flights and their interaction, instead of a single flight.

The first section starts with introducing arrival management in its traditional form, as this will be an important element of the concept presented in this chapter. Based on a suggestion made earlier, Section 4.1 also discusses how arrival management support systems can be made environment aware with the inclusion (or possibly substitution) of optimisation objectives. Section 4.2 concerns the development of such an environment aware arrival management tool. The tool itself is based on the implementation of a sequencing method selected from literature, that is extended such that it gains the ability to address community noise impact and flight efficiency objectives simultaneously. Numerical examples of results obtained with this environment aware arrival manager are presented in Section 4.3, followed by a discussion in the final two sections of this chapter.

4.1 The concept of an environment aware arrival manager

Based on the introduction of the different level of aggregation, the operational level comprises the mitigation measures related to the actual day-to-day operation of the airport. As this level can be seen as the level where the actual environmental

impact is caused, or at least geographically allocated, it is readily clear that this level is of paramount importance. With respect to this distribution of the impact over the different areas around the airport, important control mechanisms are the allocations of flights to runways and corresponding routing towards or from that runway.

Routing and runway allocation are decisions within the domain of air traffic control. This would mean that air traffic controllers have much influence on the environmental impact or at least its allocation. However, it should be realised that they are first of all responsible for safety, and have to meet objectives with respect to efficiency as well, which are often conflicting with the environmental objectives.

The issue of these three conflicting objectives can be illustrated by recapping the situation for full, three-dimensional trajectory optimisation as discussed in the previous chapter. It was assumed at that time that it is feasible to execute such trajectories, at least in isolation. However, a situation where dozens of aircraft would be flying their own trajectory in a relatively small area is assumed to be uncontrollable and therefore unsafe, at least based on the current state of technology available to air traffic controllers. For this example it would mean that in order to achieve the environmental benefits, a reduction in capacity would be required to ensure safety.

The concept of the environment aware arrival manager is based on a decision support system for air traffic controllers that will regard three conflicting objectives. The system is designed to meet requirements with respect to safety while simultaneously optimising for both efficiency and environmental objectives.

An Arrival Manager (AMAN) is a tool to help controllers in creating an efficient flow of aircraft towards the runway, eliminating delay as much as possible while meeting separation requirements. Typical functionality is to generate a landing sequence and the corresponding schedule for that sequence. Furthermore, runway assignment can be performed as well for situations where there are multiple runways available for landing. Depending on the capabilities, the support tool may also provide trajectory advisories and corresponding flight speeds^[110]. Typically, these trajectories would be small variations to existing arrival patterns, in order to achieve a certain amount of delay required for a safe and efficient flow.

As mentioned in Section 2.5.2, the suggestion to adapt an arrival manager to make it *noise-aware* has been made earlier. This Noise Avoidance Planner (NAP) was based on the idea of injecting noise related information into an already existing constraint resolution and scheduling logic, to create a system that is capable of generating advisories with respect to both delay and noise. Unfortunately, because of the underlying constraint resolution architecture, noise considerations and efficiency could not be addressed simultaneously^[41;42].

The tool presented here does not suffer from this problem, because it was explicitly developed to address these two objectives simultaneously. As such, it allows to study the interaction between noise and efficiency objectives in an

otherwise delay driven model. On the other hand, the tool presented here was not designed or build to function as an operational arrival manager, but purely as a model to generate results for this study.

Apart from balancing efficiency and environmental performance, adding environmental awareness to an AMAN can serve another purpose. For air traffic controllers, taking environmentally beneficial decisions can also be difficult because of a lack of required information. Although controllers may have access to static information, like the preferred use of certain routes and runways, they are not provided with information on the continuously developing situation with respect to, for example, the community noise exposure. This means that they cannot respond, at least not without considerable delay, to developments in the noise exposure in the past or expected developments in the near future. Nor can they evaluate the environmental effects of decision they are about to take. Automated support systems on the other can be developed to perform all of these tasks simultaneously.

Although the system that is presented in this chapter only focusses on aircraft noise, the same principle could be applied to emissions. When changing the runway or routing of aircraft, the spatial distribution of emitted emissions is altered as well, resulting in a influence on the local air quality to some extent. Although aviation as a source only has a limited share in the resulting local air quality, it may for example be possible to avoid certain areas if that area is experiencing air quality problems at that time. In this case it could be a dynamic measure, comparable to road traffic demand management strategies used to improve the actual air quality in a particular area^[98].

4.2 AMAN optimisation model

All of the elements of the environment aware AMAN optimisation model will be presented in the next few sections. First Section 4.2.1 will look into the metrics and indicators that are used to decide on the environmental performance of the alternatives the arrival manager will generate. The next two sections look into the approach procedure and the different routing alternatives that are available to the AMAN, including their environmental impact for several selected aircraft types. Section 4.2.5 presents the sequencing and scheduling algorithm, based on an existing sequencing method that was selected from literature. This section is followed by the definition of the objective for the optimisation, and finally Section 4.2.7 provides an overview of the assumptions and limitations of the model.

4.2.1 Environmental metrics and indicators

In order to allow the model to perform routing selection based on noise criteria, it requires information on the noise exposure resulting from the selection of a route. As already discussed in Section 3.2.2, the actual noise exposure for the different

L_{Amax} dB(A)	Complaint rate (per 1000 inhabitants)
$x < 50$	0.000
$50 \leq x < 60$	0.130
$60 \leq x < 70$	0.437
$x \geq 70$	1.269

Table 4.1: *Dose-response relationship for expected complaints*

trajectories and aircraft needs to be expressed using high-level noise indicators computed from the actual noise levels in all relevant areas. Four different single event indicators based on two different metrics have been used for this study.

The first two indicators are based on peak level limits. The WHO recommends to limit the number of maximum A-weighted indoor L_{Amax} events of 45 dB(A) or more during night-time in bedrooms^[32]. This corresponds to a 60 dB(A) level outdoors, when assuming a sound transmission loss of 15 dB, a modest value that allows for people to sleep with the windows open. Based on this recommendation, the number of people exposed to higher peak levels during an aircraft flyover is a suitable single-number indicator for undesirable night-time noise events. This number is used as the first indicator. For day-time noise, there is no similar recommended or commonly applied limit. However one could argue that, based on the 10 dB penalty that is typically applied to night-time events for the cumulative metrics, a daytime 70 dB(A) L_{Amax} limit is equivalent to a 60 dB(A) night-time limit. Therefore, the number of people exposed to peak levels higher than 70 dB(A) is used as the second indicator.

The third indicator is the FICAN relationship as introduced in Section 3.2.2, representing an upper bound on the percentage of people likely to awake due to a flyover. The fourth and final indicator is the estimation on the number of complaints that can be expected from a single flyover. This relation is based on Dutch research that relates the number of complaints concerning a flyover to its (computed) maximum noise level^[80]. In this research it is found that the different communities around the airport show different complaint rates. However, because the study was performed for a limited number of communities only, here the complaint rates for the most sensitive community are used for the complete study area. Therefore, the number should be interpreted as a worst bound on the expected number of complaints due to an aircraft noise event. The resulting dose-response relationship is given in table 4.1.

4.2.2 Continuous Descent Approach

Although this chapter is focussed on an optimisation application at the operational level, this does not exclude the use of a noise abatement arrival procedure for the individual trajectories, although it will be a generic procedure. As briefly introduced in Section 1.3, the general principle of the CDA is to keep aircraft




Figure 4.1: *The continuous descent approach compared to a conventional step-down approach*

higher and at reduced thrust while performing the approach. Both can be accomplished by maintaining a continuous descent, as level segments during the approach lead to higher thrust setting in order to maintain airspeed at these segments, also see figure 4.1. The procedure results in benefits with respect to noise, emissions and fuel burn if executed properly.

The CDA can be classified into two types: the Basic CDA and the RNAV CDA^[91]. When using the basic CDA, controllers retain the lateral controlling freedom through issuing vectors, quite comparable to the practice with conventional approaches. Different is that they provide the flight crew with estimates of the distance to the runway threshold. On the flight deck, these estimates are used to achieve the CDA, using rules of thumb or dedicated charts. The basic CDA is in use at several major airports, including London Heathrow^[77]. Main problem with this procedure is that the distance to go estimates may be inaccurate. This could lead to level segments if the distance is under-estimated or rapid descent or even missed approaches if the distance is over-estimated. Both situations result in reduced environmental benefits.

The second type of CDA is the RNAV¹ CDA. This CDA makes use of pre-defined arrival routes based on waypoints, which can be pre-programmed in most aircraft. Since the route is fixed, the Flight Management System (FMS) can calculate the distance to go and the corresponding ideal descent profile, leading to maximal environmental benefits. Main problem related to the RNAV CDA is its deterministic character. Controllers lose all their controlling freedom, as the route is fixed and the vertical speed profile are determined onboard the aircraft. The result is a reduced landing capacity, which explains why at some airports this procedure is currently only used during low-demand hours. It also means that the CDA is good example of conflicting interest between environmental and efficiency objectives.

¹RNAV and traditional instrument navigation are explained in Appendix A



Figure 4.2: *The three arrival routes*

The RNAV CDA is the type that will be used for the environment aware arrival manager, not only because its environmental performance is typically better. It has an important advantage for this particular application, because both the routing and the vertical plane are predefined. This means that noise levels at the ground can be computed rather accurately before the approach is executed, which is beneficial for a planning tool.

4.2.3 Arrival route allocation

As mentioned, the model can deal with multiple fixed routes to a single runway of an airport. The configuration presented in this paper is based on three routes towards one of the runways of Amsterdam Airport Schiphol. This is depicted in figure 4.2. For this scenario, it is assumed that all traffic from the West is guided to this runway. Traffic from the East is not modelled, but is assumed to land at another runway, independent of the modelled part. Of the three routes, the last part of the centre route (labelled B) is very similar to the current night-time CDA for this runway. The two outer routes (A and C) are variants of the route in the sense that they cross the coastline either further away or closer to the runway. Please note that in reality, these routes do not exist and the use of fixed arrival routes is currently limited.




Figure 4.3: *Metering fixes and starting points for the three routes*

Because all approaches are assumed to be CDAs from the points where the fixed routes start, the three different routes cross the more densely inhabited areas close to the coastline at different altitudes, before turning towards the runway over a less populated area. This is expected to result in different noise exposure, except for the final part, where all three routes are equal.

Before a flight movement starts on one of the three depicted fixed approach routes, it is assumed that it crosses one of two available metering fixes, situated north-west and south-west of the starting points of the fixed approach routes, as shown in figure 4.3. Since all three routes can be used from both fixes, this may lead to crossing traffic. The model does not regard separation before traffic is on the depicted routes. Therefore, possible conflicts in the area between the gates and the approach routes must be solved using altitude constraints.

The model is aware of the (undelayed) transit times from both fixes to the runway threshold via the three different routes. These transit times have been determined using the NLR ATC Research Simulator (NARSIM) for different aircraft types, see Appendix D.1 for more information.

Indicator	Route	Aircraft type				
		CRJ700 Dash 8-400	A320-200 B737-800	A330-300 B767-400	B777-300	B747-400
1: 60 dB(A) L _{Amax}	A	228	545	30350	24030	50335
	B	178	1045	7108	4845	24655
	C	178	1045	8448	6945	25480
2: 70 dB(A) L _{Amax}	A	35	73	298	348	700
	B	35	78	308	345	1095
	C	35	78	308	345	1093
3: FICAN awakenings	A	445	773	1860	1672	2756
	B	287	470	1896	1560	2929
	C	321	501	2483	1946	3909
4: Expected complaints	A	3.22	5.83	18.82	16.33	26.86
	B	0.97	2.66	11.84	9.55	21.72
	C	1.29	2.78	13.18	10.42	26.78

Table 4.2: Noise indicator values for all aircraft types and routes

4.2.4 Route related noise impact

Using the Dutch computation model, the noise results for different aircraft types and the three different routes have been computed^[101]. Figure 4.4 provides the results for two aircraft types (B747-400 and CRJ700) for the three routes. When analysing the noise contours, it is readily clear that both the type of aircraft and the selected route have a large influence on the size and the location of the impacted area. Furthermore, it can be seen that in absolute terms the influence of the route selection on the total noise exposure is not equal for both aircraft types. For the small CRJ the noise reallocation caused by a route change is relatively small compared to the shift in noise load that follows from a re-routing of a B747-400.

The results for all aircraft types and the three different routes have been computed for the four single event indicators as introduced in Section 4.2.1. The results are presented in table 4.2 for a limited number of aircraft types. Please note that aircraft types are categorised in the Dutch noise model based on their noise performance. This may lead to the same results for different types, like for the Airbus A320 and the Boeing 737-800.

When analysing the different results, the first observation is that the *inhabitants within the 70 dB(A) contour* metric can hardly make a distinction between the three different approach routes, except maybe for the Boeing 747. This means that it is unfit for this purpose and therefore discarded. Not very surprisingly, the remaining indicators show that exposure increases with aircraft size, with the exception of the B777-300 category performing slightly better than the smaller A330-300 / B767-400 category. Concerning choice of routing, the B-route generally performs best, with a second place for the C-route. Again some exceptions can be observed, especially for the B747-400.

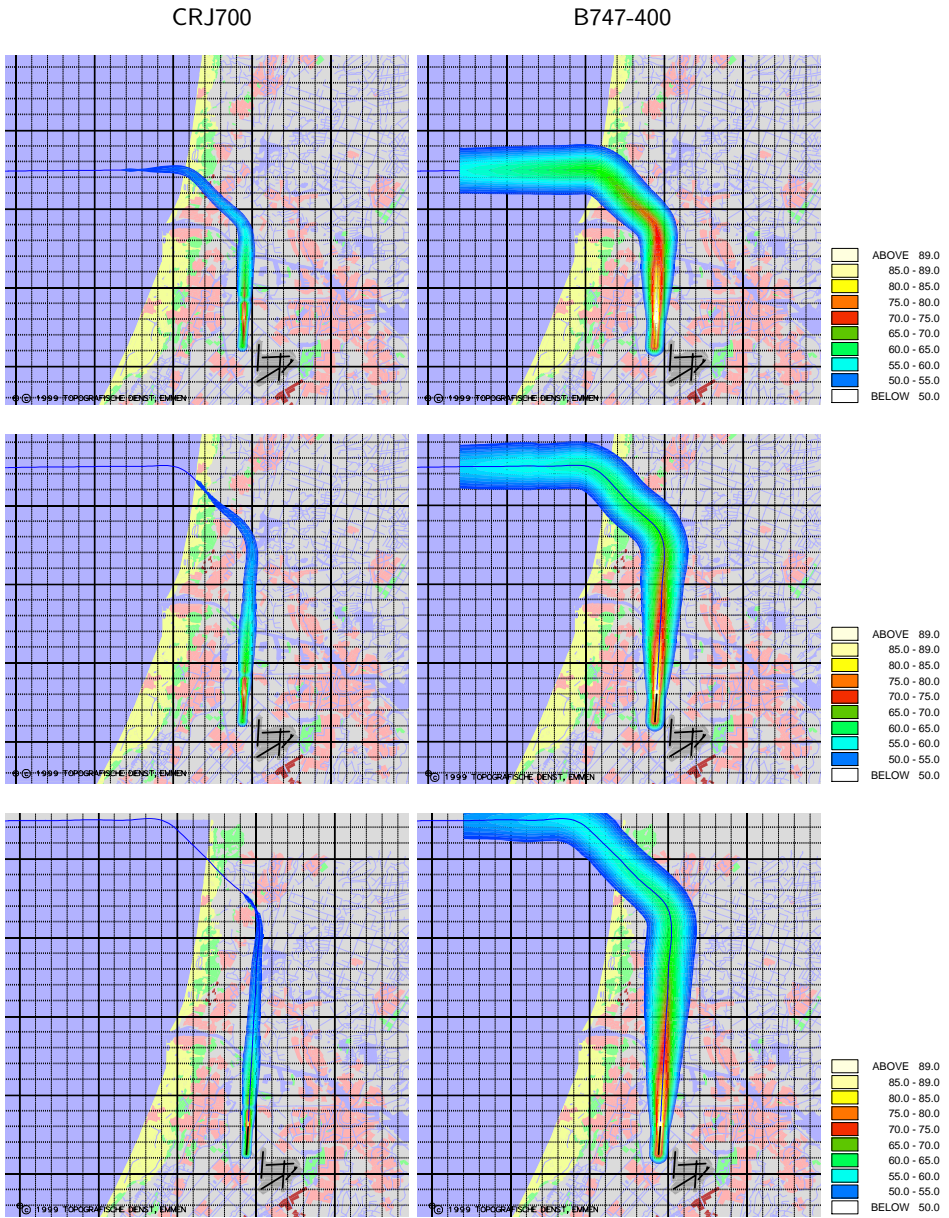


Figure 4.4: Noise results (L_{Amax}) for two different aircraft types and three different routes




Figure 4.5: Decision tree for the example problem based on Constrained Position Shifting

It should be mentioned here that instead of selecting one of the single event metrics for the arrival management model, it is also possible to use a composite function of the three remaining indicators, by summing them up. However, because of the very large differences in absolute numbers, some sort of scaling is deemed necessary to prevent one indicator from dominating the other ones.

4.2.5 Sequencing and scheduling

The sequencing is based on the existing principle of Constrained Position Shifting (CPS), where an aircraft is allowed a difference of n positions between the First Come First Serve (FCFS) order and the actual landing order^[27;29]. When using FCFS, all aircraft land in order of their scheduled arrival times at the runway. When using CPS- n , an aircraft that is for example fourth in the FCFS sequence, is allowed to take the landing positions 3,4 and 5 when $n = 1$. For a sequence of four aircraft, this leads to a decision tree as depicted in figure 4.5. Aircraft 1 and 6 do not join the sequencing process. Aircraft 1 can be thought of as the last aircraft that already has a fixed or frozen landing position and time. It prevents the aircraft taking position 2 from landing earlier than possible, based on the landing time of aircraft 1 and the required separation. The last aircraft, number 6 in this example is not necessarily a real (future) aircraft, but is mainly used to prevent the scheduler to push heavy aircraft to the back of sequence. Without this additional aircraft, the scheduler might do so because it does not regard the required separation behind the last aircraft in the sequence. Adding the dummy aircraft automatically adds the required separation behind the last real aircraft.

The CPS problem is formulated as a Mixed Integer Linear Programming (MILP) problem. This approach allows the problem to be formulated in generic form using algebraic constraint equations instead of designing a dedicated algorithm. Apart from the clarity of using equations, it also allows for easy changes or additions to the model, such as changing the performance index. The details on the modelling of this problem are provided in Appendix D.2 through D.5.

Apart from scheduling, there is also the route selection process. The scheduler is forced to choose exactly one of the three offered approach routes for each flight. The route selection determines the noise score for a specific flight as discussed

previously, as well as the earliest possibility to land. For example, a Boeing 737-800 approaching via the southern metering fix cannot land earlier than the time required to reach the fix plus 737 seconds when using route A, 807 seconds when using route B or 867 seconds when using route C.

Some additional constraints are required to model this problem correctly. Aircraft are prohibited to land earlier than their predecessor in the sequence and are also required to respect a minimum separation time based on the wake vortex category of the pair under consideration. The separation values used are 95 seconds minimum following a medium aircraft, 125 seconds minimum for a heavy aircraft following a heavy one and 155 seconds minimum for a medium following a heavy.

Due to the constraints described above, it is very likely that an aircraft is scheduled to land later than its earliest possible landing time. The difference - that is the delay - needs to be absorbed at some point. The model does not look into how the delay is absorbed; it only calculates the required amount. In practice, delay should probably be accommodated before crossing the metering fix. Because of the location of the metering fix, this will also not influence the noise impact.

4.2.6 Optimisation objective

The objective function for this problem is defined as:

$$Min : \sum_{j=1}^m LT_j + k \cdot \sum_{j=1}^{m-1} NE_j \quad (4.1)$$

where LT_j is the landing time of the j^{th} aircraft, k is the noise cost multiplier and NE_j is the noise exposure of the j^{th} aircraft. The landing time is expressed in seconds from the instant the schedule is created, and is used as a proxy for delay. The noise exposure itself is formulated as:

$$\begin{aligned} NE_j &= a_j \cdot NE_{j,r=A} + b_j \cdot NE_{j,r=B} + c_j \cdot NE_{j,r=C} \\ a_j + b_j + c_j &= 1 \\ a_j, b_j, c_j &\in \{0, 1\} \end{aligned} \quad (4.2)$$

where $NE_{j,r=A}$ is the noise exposure of aircraft j when using approach route A, etc. As can be seen, the noise cost for the last aircraft is excluded from objective function, since this is not a real aircraft to be scheduled. Its flying time is included on the other hand, because of the reason the aircraft was added in the first place. The noise cost and flying time for the first aircraft are taken into consideration, although the scheduler will not be able to optimise for these values, since they are already fixed. Noise cost multiplier k determines the importance of the noise related performance relative to the delay related performance. When k equals zero, noise exposure is not regarded at all, turning the optimiser into a traditional, delay driven tool. When k is very large, the optimiser will still

generate an optimal landing schedule, but the routing process is expected to be completely dominated by noise considerations.

The problem itself is generated by a script that reads the input variables, and writes the mathematical formulation for the problem. The problem is then solved by a solver, as discussed in Appendix C.3. Finally, the solution as returned by the solver is post-processed for ease of interpretation. In the post processing, the solution is also converted to an input file for the same air traffic control simulator that was used to determine the transit times. Using this file, the simulator can be instructed to playback the solution on a radar screen, making it very convenient to visualise, verify and interpret the results. For more information on the verification of the model, see Appendix D.6.

4.2.7 Model Assumptions and limitations

Several assumptions have been made with respect to the design of the optimisation model. These assumptions, which may lead to limitations depending on the intended use of the model, are summarised here:

- As mentioned, the fixed route system does not allow for much delay absorption and it is assumed that, if required, delay is absorbed before starting on any of the fixed routes
- It is assumed that all transit times are deterministic. The research model does not have any replanning functionality in case of unexpected events
- It is assumed that aircraft are sufficiently separated when approaching on different routes. The model only provides in-trail separation for the common part of the trajectory and for subsequent aircraft that approach over the same route

4.3 Numerical examples of an environment aware arrival management

Scheduler results and the trade-off between average delay and noise exposure are shown in figures 4.6, 4.7 and 4.8 for 20 arrivals in a mix of 30% heavy and 70% medium aircraft. Figure 4.6 is based on an arrival rate of 45 aircraft per hour, which is higher than the runway capacity. Figure 4.7 is based on the same traffic sample of 20 aircraft, but arriving at a rate of 36 aircraft per hour and figure 4.8 is based on an arrival rate of 30 aircraft per hour. All figures show the average delay per operation and the resulting Noise Exposure Index (NEI), both against the noise cost multiplier k .

During the experiments, this multiplier was varied between 0 and 500. However, from the results it was clear that the solution does not change any further

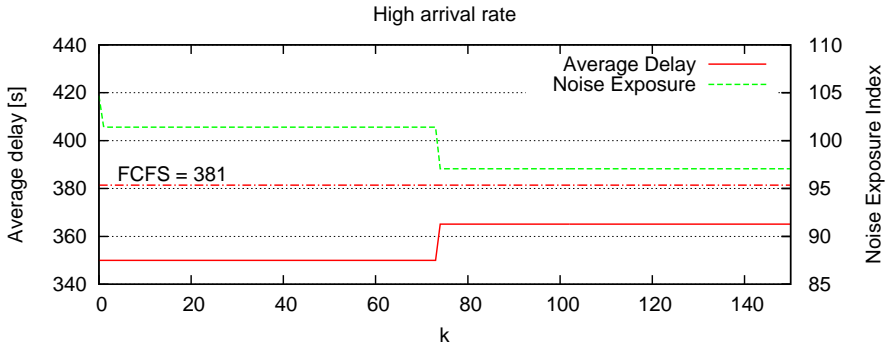


Figure 4.6: Scheduler results for an arrival rate higher than the runway capacity

when k is increased to values higher than 150. Apparently, noise is already dominant at that point and therefore the results do not display results for higher values of k . For the figures presented here, the NEI is the sum of expected complaints for all aircraft in the sequence. The average delay that is shown is also compared to the average delay that is achieved when using FCFS. The FCFS solution is based on time optimal routing only, so it does not regard noise exposure at all.

From the results, it can be concluded that adding noise considerations to the model does indeed reduce the noise exposure indicator. Furthermore, increasing the importance of the noise objectives relative to the efficiency objectives, the noise indicator value can be reduced further, at the cost of increased delay, possibly leading to a solution that is worse than the reference FCFS solution. Of course, in such a situation, the noise exposure of the optimised solution is lower than that of the FCFS solution. Interesting to note is that a (small) noise improvement can be achieved without an increase in delay. This can be seen in all three figures by looking at the differences in solutions resulting from $k = 0$ and $k = 1$. Apparently, routing can sometimes be changed in favour of noise without affecting the sequence and the schedule.

When comparing the three figures, the effect of the arrival rate can be seen. In the situation where the arrival rate is below the runway capacity, it can be seen that reducing the noise exposure indicator easily leads to solutions that are worse than the FCFS solution in terms of efficiency. At arrival rates higher than the runway capacity, the situation is clearly different. In this situation, all aircraft need to be delayed. For efficiency, it does not matter whether the required amount of delay is absorbed before the metering fix or during the approach by using a different route. This allows the scheduler to assign longer, noise preferential routes without effecting landing times and runway throughput.

Apart from the results shown here, which are based on the expected complaints

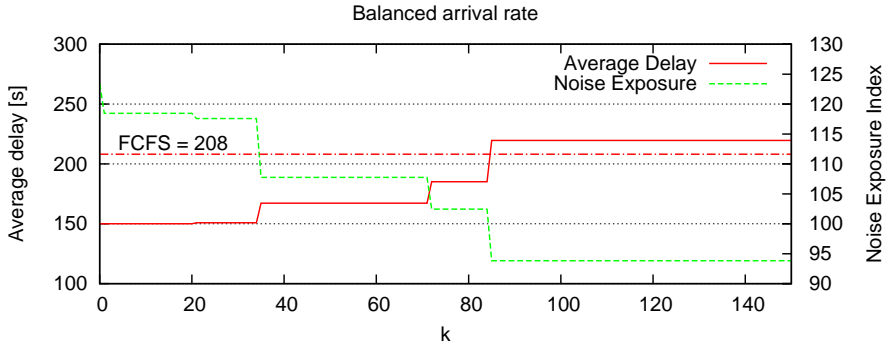


Figure 4.7: Scheduler results for an arrival rate near the runway capacity




Figure 4.8: Scheduler results for an arrival rate lower than the runway capacity

indicator, additional results have been generated for the other optimisation criteria. Similar results are obtained when using the other two indicators, as well as with an indicator based on a combination of the three. More results have also been generated using higher and lower arrival rates and different traffic samples, all showing different results of course, but similar trends.

4.4 Cumulative exposure and fairness

So far, only single event noise metrics have been used for the environment aware arrival manager. This section will first discuss an alternative formulation based on cumulative noise exposure. Based on the outcome of this analysis, yet another objective function for noise is proposed, based on the notion that it is also feasible to distribute noise as equitably as possible over the different areas.

4.4.1 Cumulative exposure

Rather than on single event metrics or indicators, true community noise exposure is often based on a cumulative exposure metric, such as the the day-evening-night level L_{den} . However, when assigning routes to aircraft based on the single event noise indicators only using a method as showed above, it is likely that one particular route is used predominantly, especially if the noise cost multiplier is set high. When only using single event indicators, a route that is optimal for a particular flight, is also optimal for all similar flights. Basically, this means that the model concentrates noise in the least noise-sensitive region. While this is indeed considered optimal considering the objective function, such a concentration of noise may be unacceptable.

An obvious solution would be to add a cumulative exposure indicator to the problem, like the annoyance dose-response relation as presented in Section 2.1.4. However, irrespective of using a single event or cumulative exposure metric, the scheduler will still need to make decisions on a per flight basis. A possible approach would be to use the cumulative exposure of any past period, and calculate the increase in exposure due to the flight currently under consideration. Based on the difference, the rise in total population annoyance can be computed, resulting from the marginal contribution of that movement.

The first problem with this approach is that the dose-response relations have been established for long term average and stabilised exposure. As such, the additional annoyance calculated from a single flight, is certainly not guaranteed to be near the actual increased annoyance due to single flyover, assuming such an increase could be quantified in the first place. However, when aware of this limitation, the result can still be used to compare different alternatives.

A more fundamental problem lies in the nature of the dose-response relationships. One of these functions, as established by Miedema and shown in Section 2.1.4 is plotted again on the left side of figure 4.9^[84]. In this figure, the population




Figure 4.9: *The dose-response relationship for annoyance, plotted versus L_{den} values as well as versus the number of daily events*

annoyance percentage plotted versus the day-evening-night-level (L_{den}), expressed in dB(A). In this figure, the function appears to be convex. This would be desirable for our purpose, because when minimising for annoyance, the increasing slope would result in traffic being directed away from the areas where exposure is already high.

However, the function can also be plotted versus a number of noise events, say the number of daily noise events of 90 dB(A) SEL each, as indicated in the right part of the figure. Here the function turns out to be concave, resulting in exactly the opposite behaviour: as soon as a certain area is experiencing high noise levels, annoyance is hardly increased by additional flights. This is not a property exclusive to this particular dose-response relationship. Other, similar functions^[52] show the same behaviour. Apparently, at least according to these dose-response relationships, overall community annoyance can be minimised by maximising exposure in the least sensitive area.

At this point it is concluded that it seems feasible to use indicators based on cumulative exposure for the optimisation objective, based on the notion of the (fictitious) marginal contribution of a particular flight to total community annoyance, and this will be investigated more thoroughly in Chapter 6. However, the identified problem of the potential unacceptably high noise levels due to the noise concentration behaviour cannot be solved using this alternative objective formulation.

Alternatively, it is possible to set maximum allowable cumulative levels. When enforcing these maxima by adding corresponding constraints to the environment

aware arrival manager, traffic will automatically be redistributed in order not to violate the limits.

4.4.2 Fairness

Apart from the options discussed above where the goal is to allocate the noise such that overall community exposure or response is minimised, a different approach would be to distribute the noise based on the principle of fairness. This idea of fairness has already been mentioned in Section 2.5.1 in relation to trajectory optimisation^[89;90]. For this particular application, the fair allocation of the noise relates to single-event noise allocation. However, fairness is also studied in relation to cumulative exposure. When part of the noise allocation policy, it can as a non-acoustical factor (see Section 1.5) help to reduce the noise annoyance under certain conditions^[81].

Concerning a practical example, the airport of Sydney, Australia has a long term operating plan that is based on a noise-sharing principle. For as far as noise exposure cannot be avoided, for example by routing flights over water or non-residential areas, noise impact is shared fairly. This is achieved by rotating runway use, limiting the number of hourly aircraft movements, maximising hours with no or little overflights and preventing that residential areas are not simultaneously exposed to noise from both arrivals and departures from the same runway^[34].

This noise sharing principle has also been tested at the TU Delft in relation to arrival management. A similar, although multi-runway route and runway allocation optimisation model for arriving traffic was extended to take the noise load distribution into consideration. Instead of a multi-objective approach and trade-off analysis as shown in this chapter, the noise sharing functionality was implemented using constraints that limit differences in noise load at a number of specific locations. For an example problem, it was concluded that by adding constraints that limit the difference in noise load, a noise spreading allocation strategy can indeed be achieved. This results in relatively wide areas that experience peak noise loads, with overall peak loads lower compared to the situation where noise is not forced to be spread. In terms of efficiency, adding the noise spreading constraints leads to an increase of arrival delay. The increase was, however, less than one percent for the sample problem^[63].

4.5 Discussion

This chapter presented the concept of an environment aware arrival manager (AMAN). The concept is based on an extension of a typical arrival manager, a decision support tool for air traffic controllers, used to assist in creating an efficient flow of aircraft towards the runway. Based on a combination of fixed arrival routes and continuous descent approaches, the arrival route selection process is adapted to enable a particular noise allocation strategy. To this end, the objective function

regarding delay is changed into a multi-objective formulation based on delay and computed noise impact.

In its current form, only single event noise metrics have been used for the environment aware arrival manager. The result is that the model will concentrate noise in the least noise-sensitive region, which could be deemed unacceptable. It is argued that a noise objective based on annoyance due to cumulative noise exposure is no solution for this phenomenon. Instead, noise distribution policies based on fairness and/or noise sharing can be used, at least from a technical perspective. Eventually, the selection of a noise allocation strategy is a sociological-political problem considered to be outside the scope of this work.

Irrespective of the discussion regarding which noise allocation strategy would be best, two important conclusions can be drawn based on the results of the current state of the model. Firstly, the results show that in some cases some reallocation of the noise load can be achieved even without sacrificing efficiency. Secondly, if desired, the noise load can be redistributed further, at the expense of an increase in delay. For low traffic situations, this can easily lead to situations that are worse than the FCFS solution in terms of delay, but for heavy traffic situations, the trade-off is more advantageous.

Based on these results, it is recommended to incorporate noise information in the arriving traffic scheduling process, especially if the arriving traffic operational concept is based on fixed but alternative arrival routing options as assumed in this chapter. Even when sacrificing efficiency is deemed unacceptable, there seem to be opportunities for noise allocation according to the selected allocation policy.

The optimisation application presented in this chapter was classified as a system at the operational level, one level of aggregation higher than the trajectory level system presented in the previous chapter. The next chapter will increase the level of aggregation further, and will demonstrate a system at the strategic level.

It should also be mentioned that this chapter demonstrated an operational level application targeted at arriving flights only. There are, however, also opportunities for departing flights at the operational level, as will be demonstrated when addressing the multi-level system in Chapter 6.

Chapter 5

Environmental impact allocation in airport strategic planning

This chapter^[67] contains the work related to the third and final stand-alone DSS presented in this thesis. The previous two chapters dealt with examples originating from the lower two levels of aggregation: the trajectory level and the operational level (see figure 1.1). In this chapter, the focus shifts to the next and also highest level of aggregation: the strategic level. The resulting strategic level support system will also form the third and final cornerstone for the multi-level system to be presented in the next chapter.

The strategic mitigation strategy outlined in this chapter is based on the long term average spatial allocation of the negative impacts due to flight operations, noise and Third Party Risk in particular. The DSS developed for this purpose uses the allocation of flights to runways for a full operational year as the main instrument for achieving a particular distribution in environmental impact.

The first section of this chapter provides more detail on this concept. Then, Section 5.2 presents the runway allocation optimisation model that has been developed. This section also describes all underlying, pre-existing models that have been used to generate input for the main optimisation model. Finally, Section 5.2 also shows the user interface design for this DSS that helps the user to achieve a satisfactory trade-off between various performance criteria through the use of an interactive procedure. The next section contains a collection of numerical examples generated with the runway allocation optimisation model, followed by a discussion on the results in Section 5.4.

5.1 The concept of strategic environmental impact allocation

For airports that have multiple runways, the actual distribution of the flights over the different runways is a problem concerning a multitude of often conflicting interests. Typically, one of the runways may be beneficial for noise reasons, where another one may be preferred in the sense that it is perceived to be safer for the surrounding community. Simply choosing one of these runways for all flights is usually not an option. Not only could it lead to huge delays in most cases, but under certain weather conditions it may lead to unacceptable risks for the passengers on board the aircraft. In 1997 an accident occurred at Amsterdam Airport Schiphol that illustrates this problem. A passenger aircraft was blown off the runway during landing by a severe crosswind. Later investigations showed that the aircraft was assigned to a runway because of noise considerations, where another runway would have been more appropriate under the occurring wind conditions at that time^[10;11].

This example shows that runway allocation can be of vital importance at the operational level. The process can be simple, such as an air traffic controller using rules of thumb, or choosing a runway that leads to the shortest path or choosing one that is in another way convenient. However, it could be that these simple methods do not yield the desired results, for example in terms of efficiency. In that case, it can be considered to use a DSS to aid controllers in making the final decision for arriving flights^[37;73].

For departing flights, runway assignment is usually a rule-based activity too, that for example depends on the destination of the flights in order to eliminate the need for crossing departure flows. However, also for departures, support systems have been designed that consider other objectives when appropriate, such as balancing the load on the active runways^[28].

Based on these examples, it is evident that the actual runway allocation process takes place at the operational level. However, this does not mean that there are no opportunities to study or optimise runway allocation from a strategic perspective, especially in relation with environmental impact management. This notion, the concept of strategic environmental impact allocation, forms the foundation of this chapter. The concept itself is based on the annual (or even longer term) distribution of flights over runways with the purpose of achieving a particular spatial allocation of the environmental impact. Here the strategic variant is also at an advantage, since the environmental impact levels and corresponding limits are mostly based on annual doses rather than hourly or daily, thereby matching the planning period of the strategic level.

Strategic runway allocation results can be used for different purposes. First of all, strategic runway allocation results that have been optimised with respect to the allocation of environmental impact can be used as input for a runway assignment policy for the operational level. Two optimisation studies^[54;82] aimed

at generating an optimal allocation policy for Schiphol Airport in order not to violate annual noise limits have already been discussed in Section 2.5.3. Compared to those models, the strategic allocation model presented in this chapter will not only consider noise as objective, but will also allow a trade-off between noise, third party risk and delay. Furthermore, it does not explicitly consider noise limits, but rather aims for noise impact minimisation. On the other hand, the developed model is at a disadvantage in terms of applicability, at least based on the current operational practice. The exact difference and its implications will be discussed in Section 5.4.

With respect to the second application, runway allocation and resulting environmental impact may also be considered in airport strategic planning activities. When an airport wishes to construct a new runway for example, the location of this new runway together with the new flight allocation regime determines to a large extent the future situation with respect to noise and third-party risk. Because of the complexity of such problems, the use of a DSS can help to identify solutions with maximum efficiency with respect to the scarce available resources. HARMOS, developed at the TU Delft, is an example of such a system^[112]. Since the model presented in this chapter was designed as a module for HARMOS, it is also able to produce results for these type of strategic problems.

Irrespective of whether the system is used in the design process of an operational policy or for exploring master planning options, the user of the DSS would typically be a person either working directly for or by order of the airport authority. Alternatively, the system could also be used by regulators responsible for setting the environmental regulations the airport has to comply with. The remainder of this chapter describes the development of the system, a runway allocation optimisation model for the strategic level.

5.2 Runway allocation optimisation model

The goal of the runway allocation optimisation model is to provide the user with a balanced runway usage scheme for a generic airport, based on a provided annual traffic schedule. The optimisation model is based on three objectives. Considering the first two objectives, the optimisation model uses the results from a noise model and a third-party risk model in order to reduce total noise and risk, both with respect to the defined population density distribution. The third objective is to minimise the delay that arises with the runway allocation solution that is chosen. This third objective is chosen, not only because it represents runway efficiency, but also because delay has a negative effect on the economic performance of the airlines. As constraints, the operational runway usage, wind conditions and runway capacity are taken into consideration during the multi-objective optimisation. Finally, an interactive method is used to reach a unique final solution that not only meets the requirements but also satisfies the preferences of the user.

This section presents the runway allocation optimisation model in some detail.

However, before discussing the details of the problem and the resulting model, Section 5.2.1 first explains four important underlying concepts:

- runway configurations,
- wind limits
- wind rose analysis,
- runway allocation.

Readers familiar with these subjects may wish to proceed to Section 5.2.2.

5.2.1 Introduction to underlying concepts

Runway configurations

For airports with multiple runways it is common to use a certain number of available runways in specific combinations. This may be because not all runways can safely be used at the same time, or not all of the runways are suitable under certain weather conditions, wind in particular. Furthermore, not all of the runways may be required in terms of capacity. A runway configuration is defined here as a specification of which runways are in use, in which direction and in which mode (arrivals, departures, or mixed mode¹). Figure 5.1 shows a few examples of runway configurations as used by Schiphol airport.

Wind limits

As already mentioned in the introduction of this chapter, there are limits to the wind speeds that can be allowed for aircraft landing and taking off on a specific runway. To check for these limits, the actual wind velocity is decomposed in a component along the runway and a component perpendicular to the runway. The latter component is called cross wind and should be checked against the cross wind limits. The component along the runway is called either tail wind or head wind, depending on the direction. Like for cross wind, there are limits for tail wind. Unlike with cross wind however, a tail wind can be turned into a head wind by using the other direction of the runway.

It is important to realise that for any given runway configuration to be available under a particular wind condition, none of the runways in the configuration should violate any of the imposed wind limits.

Wind rose analysis

A wind rose is a graph showing the distribution of both wind speed and wind direction for a particular location. The underlying data that is presented in the

¹In mixed mode, a runway is used for both arrivals and departures, all moving in the same direction




Figure 5.1: Example runway configurations for Amsterdam Airport Schiphol, all four with two arrival runways (red) and one departure runway (green)




Figure 5.2: A wind rose. The orientation of the bars show the wind direction while the length of the bars show the corresponding frequency. Colours represent the wind speed distribution for each direction.

polar coordinates system is typically the long term average annual wind situation, based on thirty to fifty historical years. Figure 5.2 shows an example of a wind rose. When a wind rose is generated for the location of the airport, it can be used to calculate how often the different runway configurations are available within a year, given limits for cross and tail wind.

Runway allocation

The final concept discussed here is runway allocation itself, the process of assigning a runway to each flight movement. However, in relation to runway configurations and wind limits, runway allocation is in fact primarily the task of selecting one of the available runway configurations. Only for configuration with multiple runways available for either arrivals or departures, the runway allocation process also involves the distribution of the flights over the different runways.

Please note that when allocating a flight to a certain runway, the spatial allocation of the environmental impacts related to that flight is fixed to a large extent as well. Of course, the selection of the runway determines the general area the aircraft will overfly during its approach or departure. In most cases however, the destination or origin of the flight, in combination with the selected runway, also leads to a particular routing to or from the airport. It is exactly this property that makes runway allocation such an important control parameter with respect to environmental impact allocation.

5.2.2 Problem formulation

The optimisation model seeks to minimise two important environmental impacts of airport operation: noise annoyance and the safety risk for the surrounding community. Both are used in the objective function of this optimisation problem. However, without the appropriate constraints, a minimisation of these objectives would lead to the trivial solution of no more flight operations. Therefore, it is assumed that the optimisation will have to allocate a prescribed amount of traffic. This means that the model should be provided with a specification of the traffic demand for the considered period, such as a flight schedule.

Even with the additional constraint of meeting the actual number of flights, this optimisation would still not lead to a satisfactory result, if no runway capacity and delay characteristics are taken into consideration. To solve this, the ultimate (or theoretical) runway capacity is added as a constraint in the optimisation. This capacity is calculated using an analytical airfield capacity model. This model also provides an estimate for delay occurrences, based on the capacity used relative to the ultimate capacity. These delay results are used as the third element of the objective function of this problem. This means that the user will have the possibility to perform a trade-off between the two environment-related objectives and the resulting delays.

The operational runway usage should also be taken into consideration. All modes for which the airport may wish to operate its runways are specified beforehand as runway configurations. The availability of these configurations with respect to cross- and tailwind is calculated based on posed operational limits and statistical wind data of the airport under consideration.

5.2.3 Environmental and capacity models

As mentioned previously, the runway allocation optimisation model uses three external models to evaluate the impact related to all potential allocation decisions it could take. These models are discussed here briefly.

Airfield capacity and delay model

The FAA airfield capacity and delay model for independent runways has been selected to perform the required calculations^[71]. It is a well known, but not very extensive model, that requires very little computational effort since it is based on an analytical solution technique.

As the name already suggests, the model first determines the capacity of the runway. For computing the arrival capacity, the model assumes that there are three different categories of aircraft, all with their own approach speeds and mutual separation requirements. Based on the length of the common approach path, the model can compute the minimum interarrival times for each of the nine leader-trailer combinations. Next, it computes the average minimum interarrival time, based on the statistical occurrence of the nine pairs. The ultimate, theoretical arrival capacity then follows from the reciprocal of this value.

For departures, the model assumes fixed inter-departure times for the same nine possible combinations. Again by calculating the expectancy value, the model determines the minimum interdeparture time, that also yields the ultimate departure capacity.

The delay calculation is similar for both arrivals and departures and is based on queuing theory². Under the assumption of a steady state and a demand that does not exceed the capacity, the mean delay of the system is calculated using the Pollaczek-Kinchine formula^[69]:

$$W_q = \frac{\lambda (\sigma^2 + 1/\mu^2)}{2(1 - \lambda/\mu)} \quad (5.1)$$

where W_q is the waiting time in queue for each aircraft, λ is the mean arrival (or departure) rate, μ the mean service rate based on the calculated runway capacity and σ the standard deviation of the mean service time.

²According to queuing theory terminology, the system is classified as M/G/1, which indicates the interarrival time distribution (Markovian or exponential), the service time distribution (General or arbitrary) and the number of servers (one)

Community noise model

For the noise calculations, INM has been used^[57]. As introduced in Section 2.1.3, INM is the US FAA's standard tool for determining the predicted noise impact in the vicinity of airports, and its use is widespread in the rest of the world as well. For this application however, it is important to realise that INM uses its own Graphical User Interface (GUI) to collect its input and present the results. This severely limits the possibility to integrate this model into the runway allocation optimisation model. However, INM does use a database structure for all input and output data. This property can be exploited by the optimisation model to generate the required input for INM, i.e. to populate INM's database structure without using the program. The output data from the noise model can also be retrieved directly from the INM database, again without using the program.

To further limit the dependency on the GUI, INM is set up to perform the noise calculations for all possibly occurring movements on a per flight basis. The results for the single flight event computations are provided as Sound Exposure Levels (SEL). In this way, INM ultimately only needs to be used during the preparation phase of the optimisation problem. All further noise calculations, like computing the impact of a certain allocation decision or the final noise result can be performed by the optimisation model itself.

The dose-response relation that earlier has been introduced in Section 2.1.4 has been used as key performance indicator for noise for this optimisation problem^[84]. This relation, based on almost fifty research projects in Europe, the US and Australia, estimates the percentage of people feeling annoyed when exposed to a certain aircraft noise exposure level. The relationship is only valid from 45 to 75 dB(A). To cover the remaining levels, it is assumed there will be 100% percent annoyance at levels over 75 dB(A). The introduced error will be negligible, because the number of people living at locations with noise levels over 75 dB(A) is small. At the same time, it is assumed that below 45 dB(A) there will be no annoyance at all, which may lead to an underestimation of the number of people feeling annoyed.

Third party risk model (TPR)

For the individual risk calculations, the model developed by the UK National Air Traffic Services (NATS) has been selected^[45]. This model calculates the probability that an individual living permanently at a particular location near an airport will be killed by an aircraft impact in any given year. The method that is used by this model corresponds to the general method for calculating TPR, as discussed in Section 2.3.2. Similarly as for noise, a single number is preferred to be used as a key performance indicator. For external safety, such a number can be calculated by aggregating the total individual risk for all people living within the evaluation area. This results in an annual casualty expectancy value.

5.2.4 Objective functions

As mentioned before, three different objective functions are used for this problem, one for delay, one for noise and one for external safety. For delay, the solver is tasked to minimise the total delay (or average delay per movement, which would yield the same result). For external safety, the objective is similar. For a single movement, the total risk is calculated by multiplying the individual risk at a certain location with the number of inhabitants. This results in an expectancy value for casualties for a single movement. This number is used for the objective function for external safety.

The objective function for noise is less obvious. The annoyance dose-response relationship is for cumulative noise exposure only, so it cannot directly be used for making the allocation decision for one or just several flights. Instead, as a proxy, the SEL contour levels for every possible movement are used to calculate the number of people living within certain levels, starting at 60 dB(A). Multipliers that multiply the computed number of people are used to apply penalties for higher levels. The default multipliers are 1 for 60, 5 for 70 and 25 for 80 dB(A), but these values can be changed through the user interface. Afterwards, when the solver has reached a solution, the resulting cumulative noise levels are calculated and the dose-response relationship is evaluated in order to present the results to the user.

5.2.5 Optimisation model design

Based on all available data and the previously prepared computations on noise and external safety, the optimisation model can now be generated. A schematic overview of the design of the optimisation module is given in figure 5.3.

The allocation decisions that the model has to take are very much like the runway usage decisions the air traffic controllers would make during a complete year. This does not only include the decisions on which runway configuration will be used at a certain moment in time, but when applicable, also the distribution of the flights over the runways of a configuration.

The model uses the Mixed Integer Programming (MIP) optimisation method, discussed in Appendix C. However, before the MIP-model itself can be generated, a few preparation steps have to be completed:

- The flight schedule for the selected year is scanned to identify periods in the schedule where traffic demand is similar.
- The historical wind data is scanned too in order to identify periods of similar wind conditions.
- The results from both scans are combined and yields a number of situations, where each situation represents a small part of a certain year. The duration of these situations is related to the statistical occurrence of both similar traffic patterns and similar wind conditions.




Figure 5.3: *The design of the optimisation module. The figure shows all inputs to the optimisation model in terms of objectives and constraints. Based on the interaction with the user, the optimisation model outputs the result: a runway usage scheme that results in the optimised environmental impact allocation.*

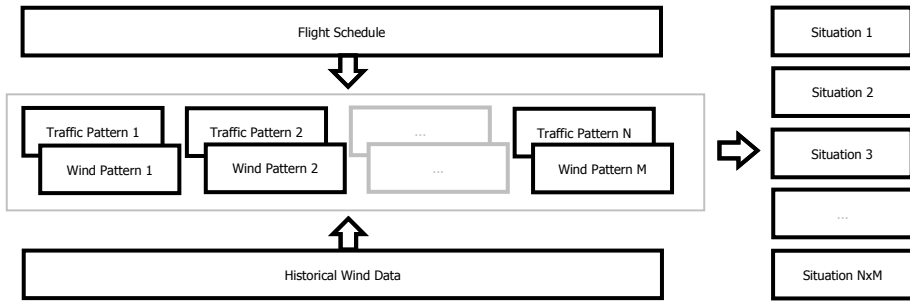


Figure 5.4: *The flight schedule and historical wind data are scanned to identify periods of similar traffic demand and similar wind patterns respectively. Based on all possible combinations, this leads to a number of situations where both wind and traffic conditions are similar.*

These first three preparatory steps are shown in figure 5.4. The process then continues:

- For each situation, it is determined which runway configurations are available, given the wind conditions and required capacity for each situation. The available runway configurations are called options.
- Delay, noise impact and risk impact are calculated for each option and are called the consequences of the options. Depending on the number of runways for a particular option, the resulting delay and impacts are either a static number, or a function of the distribution of the flights over the available runways.

At this point, all preparatory steps are completed and the actual MIP-model is created. Constraints are used to force the solver to select one, or a certain combination of some of the options offered for each situation. This automatically means that all traffic will be allocated. Because of an analysis preceding the generation of the optimisation model, only feasible runway configurations will be offered for each situation. This means that solutions returned by the solver will automatically be feasible with respect to the wind situation and operational runway use. A more elaborate description of the model itself is provided in Appendix E.1 through E.3. The verification of the model is discussed in Appendix E.4.

5.2.6 Interactive multi-objective optimisation procedure

In order to reach a final solution for the multi-objective optimisation, an interactive procedure has been developed, which is very similar to a procedure developed to solve a production planning problem with four objectives^[83]. This procedure is based on the default weighted sum method, where the actual objective is a combination of the three objective functions:

$$\textit{Minimise} : \quad m_1 \cdot w_1 + w_2 \cdot w_2 + m_3 \cdot w_3 \quad (5.2)$$

where m is a vector consisting of the weighing factors for the different objectives w_i . Often, the decision maker determines the weights beforehand, which requires a priori knowledge of the optimisation. Especially for non-experienced users of the optimisation module, this can be a problem, making it hard to reach a satisfactory result. To eliminate this weakness, the default weighted sum method is slightly adapted and used interactively.

When using this interactive method, the weights are determined automatically such that all three objectives become equally important at the point where they reach their absolute minimum. The objective function for the multi-objective problem becomes:

$$\textit{Minimise} : \quad \frac{1}{m_1} \cdot w_1 + \frac{1}{m_2} \cdot w_2 + \frac{1}{m_3} \cdot w_3 \quad (5.3)$$

where m is now the normalised vector based on the minima of the three single-objective problems. To reach a satisfactory solution, the decision maker is presented with a range for the different objectives together with three possible solutions. Starting from one of these three initial solutions, a new solution can be generated that will minimise for all three objectives simultaneously. At the same time, upper bound values for the objectives can be specified explicitly. In practice, achieved values for some of the objectives should be relaxed if the remaining objective values should be reduced further. If desired, one of the objectives can be given priority over the others during one of the iterations, but this is not required.

A part of the GUI used for this optimisation is shown in figure 5.5. The upper part shows the situation with one of the initial solutions selected. This solution, which is not Pareto-optimal³, has been calculated together with two other initial solutions to provide starting points for the optimisation and to determine the ranges for the sliders. The lower part of the figure shows the situation after a single iteration, where delay is allowed to increase a few seconds, compared to the absolute minimum. This solution is Pareto-optimal and if it satisfies the needs of the current user, the optimisation is complete. If not, the user can continue with the new solution or start over from one of the three available initial solutions.

³For a multi-objective problem, a Pareto-optimal, or non-dominated solution is any solution where one of the objective values cannot be improved without degrading any of the remaining objectives^[47].




Figure 5.5: *User interface for interactive optimisation procedure: initial solution and a possible solution after one iteration*

This way of presenting a multi-objective optimisation has turned out to be very convenient and easy to use. No fundamental knowledge concerning the underlying optimisation theory is required. It is expected that users with no optimisation experience will be able to use this optimiser with only a few instructions. However, as mentioned earlier, the user interface has not yet been tested by potential users.

5.2.7 Model assumptions and limitations

Several assumptions have been made with respect to the design of the optimisation model. These assumptions, which may lead to limitations depending on the intended use of the model, are summarised here:

- All of the results are based on data for an average year. However, especially the weather situation during a specific year is not likely to be exactly equal to the long term average. This may lead to a situation where the strategic target is not attainable for a specific year because of non-average weather conditions during that year.
- Both the used wind data and the constraints with respect to the cross- and tailwind limits are based on steady wind, not including gusts. Depending on wind characteristics, ATC policies, and airline policies it could be that for some airports, a more appropriate approach would be to use both wind data and wind limits that include gust. Both definitions of wind can also be combined, leading to the third option of applying limits to both the steady speeds as well as the wind speeds including gusts. This choice is expected to have an effect on the final results. However, since only the first approach has been applied, the magnitude of this variation is unknown.
- It is assumed that the availability of runway configurations is only dependent on the time of day and the wind situation. However, visual conditions can be a factor as well. Constraints with respect to visibility and cloud base have currently not been implemented. This could be done by combining wind and visibility constraints. This would lead to more situations for the model and probably result in higher computation times, but the structure of the problem can remain unchanged.
- For the third party risk and noise calculations, an average distribution of the flights over the different tracks of a particular runway is assumed. This distribution can be specified for each runway independently, but it is not traffic situation dependent.
- When in reality the runway configuration is changed during the day, some measures may be required to transition from one configuration to the next in order not to disturb the traffic flow severely. This may even include the temporary use of a runway that is not part of both the old and the intended new configuration. While it may be important functionality for an

allocation model at the operational level to be able to plan such transitions, it has not been considered for this strategic model. Basically, this means that the strategic model assumes that configuration changes can be performed instantly and in any order.

- Finally, the model only supports independent runway usage, and runways should only be used for either departures or arrivals. For the airport that is considered in the case study that follows in the next section (Amsterdam Schiphol), this is not a very serious limitation. However, for airports that do use mixed mode or highly dependent configurations on a more regular basis, the optimisation model should be extended first in order to obtain realistic results.

5.3 Numerical examples of strategic environmental impact allocation

In this section, some results of a case study will be presented for Amsterdam Airport Schiphol. In order to evaluate the performance of the optimisation, the optimised results will be compared to a reference scenario, based on the calculated current performance. The reference scenario results are computed using exactly the same amount of traffic and the same computation models. The only difference is that the runway usage percentages⁴ are based on the usage projected by the airport authority, instead of calculated by the runway allocation optimiser.

5.3.1 Case Study Baseline Scenario

First, the results of the reference scenario are presented in figure 5.6. Both figures show the area around the airport, the runways, and coastal lines. The left subfigure shows contour lines for the noise metric LDEN. The individual risk contour lines are shown in the right subfigure. The same holds for figure 5.7, but now for an example of optimised results. This specific optimisation result was obtained by using the the minimum delay solution as the starting point. From there, the obtained value for delay was relaxed somewhat in order to allow the noise and risk values to decrease during the next iteration. The final trade-off between the three objectives was as indicated by the sliders in figure 5.5.

When graphically comparing the two noise results and the two risk results to each other, it can be seen that the images look quite similar, but also that some differences can be identified. A closer inspection reveals that for the optimised case, more traffic is handled in the North-South direction.

Besides graphically, both situations can also be compared in numerical terms using the chosen performance indicators. For the reference scenario, the estim-

⁴While runway usage percentages can be used to estimate noise and third party risk for the reference scenario, a reference value for runway delay could not be computed using this input.

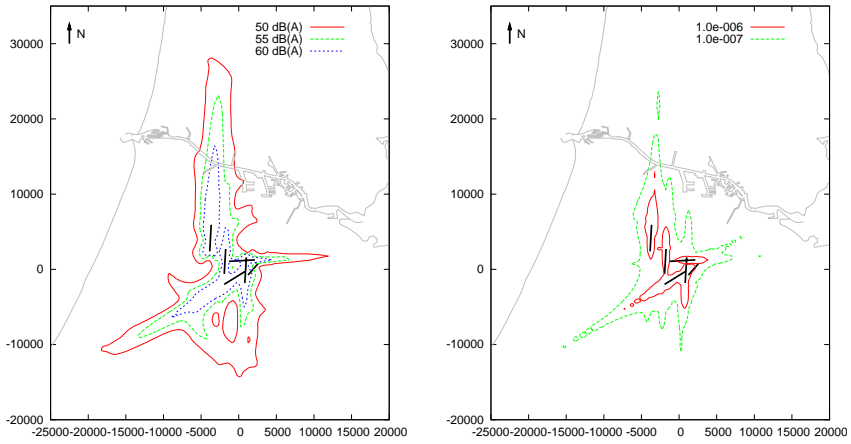


Figure 5.6: Reference scenario showing LDEN noise contour levels (left) and individual risk contour levels (right) for the reference scenario




Figure 5.7: Results after optimisation showing LDEN noise contour levels (left) and individual risk contour levels (right)




Figure 5.8: Results for the future scenario showing LDEN noise contour levels based on the current aircraft fleet (left) and based on a newer aircraft fleet (right)

ation for the number of people annoyed by aircraft noise is 98 100. For the optimised scenario, this number has dropped to 69 000. Concerning individual risk, the casualty expectancy value is 0.048 per year for the reference scenario, compared to 0.035 for the optimised scenario. This means that both performance indicators have been reduced by almost 30%.

5.3.2 Case Study Future Situation

The runway allocation optimisation model can also perform predictive calculations for future situations. This section will give an example of such an analysis, but it should be noted that the assumptions made concerning the future are not based on actual forecasting studies. The future traffic schedule is generated by increasing the original schedule by 24% in terms of traffic volume. The optimisation procedure followed to obtain the optimal solution that is presented here is the same as described before: starting from the minimum delay solution, delay is allowed to increase a few seconds in order to improve the other two objectives. The results for external safety risk are not presented here, because they look very similar to those presented in figure 5.7.

Assuming that the same fleet of aircraft will still serve the airport in the future year, the results for noise are shown on the left side of figure 5.8. The right side of the figure on the other hand shows the results if development of the fleet is

considered. This development is based on the expectation that the future fleet of aircraft will be quieter than the current fleet. For this analysis it has been assumed that there will be a decrease in SEL values of 2 dB(A) for all considered movements.

Concerning noise annoyance, the total number of people annoyed is predicted to rise to 75 700 without fleet development, but is estimated to drop even under the baseline scenario levels to 42 800 with fleet development, despite the 24% growth in traffic. A similar observation is made in a cost-benefits analysis for the expansion of Schiphol Airport with two possible new runways^[78]. When accounting for the newer aircraft, the environmental capacity with respect to noise is expected to grow faster than the physical capacity of the runway system.

5.4 Discussion

The developed runway allocation model can assign all yearly flights to the different runways of an airport, while taking into account delay, feasible runway usage combinations, wind conditions and noise and external safety risk with respect to the surrounding population. Guided by the graphical interactive optimisation procedure, a user can determine a balanced runway usage scheme, not only for the current but also for any predicted future situation with respect to airport layout, climate and traffic demand.

Concerning the results, the optimiser shows significantly improved results with respect to the reference situation. The total overall risk reduces by about 30%. The same holds for the performance indicator estimating annoyance due to aircraft noise. Unfortunately, these results cannot be directly translated to the actual situation at Amsterdam Airport Schiphol because of some key assumptions and simplifications in the design of the model already mentioned.

Furthermore, because the results are high-level and strategic, they cannot be used directly for the actual runway assignment decisions an air traffic controller would have to take. This is a difference with the optimisation studies as presented in Section 2.5.3. These models are based on generating lists of preferential runway configurations, and these lists can be used by the ATS provider. With respect to applicability, this is a disadvantage of the model presented here. At the same time, the model proposed herein is also not restricted, in terms of solution space, to optimising preference lists. Therefore it is likely that the obtained optima are more efficient with respect to the objectives, although this is not proven at this point.

It is expected that some form of guidance for air traffic controllers is required in order to implement the results from the strategic level in practice. Here it might be possible to just alter the runway allocation assignment rules such that the actual realisation would match the results of the strategic DSS. However, if this does not lead to the desired result, an operational DSS would be necessary to help controllers with their runway assignment task.

All in all, it can be concluded that runway allocation optimisation designed for minimising environmental impact yields very promising results at the strategic level. However, when it is desired that the actual runway assignment should approximate the strategic results, the use of an operational DSS is advisable as well. In this case the strategic result would become the target or input for the lower level support system. This concept will also be addressed in the next chapter, concerning the multi-level hierarchical system for integrated environmental management. This multi-level system will not only simultaneously address the operational and the strategic level, but will also make use of the results generated with the trajectory level system shown in Chapter 3.

Multi-level integrated environmental management

This chapter is devoted to the fourth and final DSS of this thesis. Contrary to the previous three chapters, the system presented in this chapter does not involve a single level of aggregation. Rather, all three levels will be addressed simultaneously, as shown in Figure 1.1. This results in a multi-level integrated environmental management system, which makes use of selected parts of the concepts, system components, and results from the trajectory level, the operational level, and strategic level system presented in Chapters 3, 4 and 5.

First, Section 6.1 will discuss the concept of the multi-level system. Following a discussion on design considerations, the resulting multi-level system presented here is categorised as an operational level DSS that interacts with the other two levels. The details and the underlying optimisation model of such a system are presented in 6.2. The next section contains numerical results that have been obtained with the multi-level integrated environmental management system. Following these results, Section 6.4 concludes this chapter with a discussion.

6.1 The concept of the multi-level integrated system

In the introduction of this thesis (Section 1.3) it has been discussed that in the ideal situation, all environmental impact mitigation efforts taking place at the different levels of aggregation are managed concurrently. In this way it can be ensured that all actions taken to minimise the nuisance caused by aircraft environmental impact will be consistent, complement each other, and make use of synergy benefits.

Subsequently, it has been proposed to design a decision support system with the capabilities to consider multiple mitigation options at different levels of aggregation, to evaluate multiple performance areas and to apply formal optimisation

for choosing the optimal combination of the mitigation measures. The resulting design problem was decomposed along the different levels of aggregation and three seemingly independent support systems have been developed. Now, some of the capabilities developed at the different levels are coupled, resulting in the integrated environmental management system envisaged originally.

Before the actual system is presented in 6.2, in the remainder of this section the multi-level system is compared to the single-level systems. The benefits of the multi-level system are discussed, together with the additional challenges that arise. It is also discussed why the design of the multi-level concept will be based on a decision support at the operational level, while involving the other two remaining levels through interaction. The last part of this section lists the most important assumptions regarding the concept.

6.1.1 Single-level versus multi-level support systems

The previously presented single-level support systems already featured two of the desired properties, namely the capability to evaluate multiple performance areas (e.g. noise impact and fuel burn) and the capability to apply optimisation in order to minimise environmental impact. The multi-level system, which is based on integrating the single level systems, adds the third requirement stipulated in the introduction: the capability to simultaneously evaluate multiple mitigation measures at different levels of aggregation.

The main benefits of adding this third capability have just been mentioned: it is expected that the overall combination of measures will be more consistent and that the overall result will improve due to synergy benefits. Accordingly, it is anticipated that where the single level systems already showed environmental impact reduction, an optimised combination of these measures leads to a further reduction in environmental impact.

While the multi-level system is expected to lead to more benefits, the major disadvantage or even risk is the increased complexity of such a system. A part of this increased complexity stems from the fact that the system will be larger, allows for more mitigation measures and may involve more environmental performance areas. Additionally, the objectives and results of the optimisations at the different levels of aggregation need to be attuned to each other.

Further adding to the increase in complexity is the multi-actor issue. As previously identified, the different actors involved in the process of managing airport environmental impact typically have different interests and objectives, which may also be conflicting. Managing these objectives with the use of a single system is expected to increase the overall transparency. This expectation is first of all based on the notion that the objectives have to be quantified and gathered in a single location. Furthermore, transparency is also anticipated to improve because the system can provide insight into the synergies and the tradeoffs. However, it should be realised that just integrating these objectives into a single system does not necessarily resolve conflicting interests of the different actors.

6.1.2 Actors involved and their responsibilities

The actors introduced in Section 3.1.2, i.e. the airlines, the ANSP and the government are also the actors that are involved in the multi-level system. For each of the individual levels, one of these actors has been identified as the principle stakeholder. For the trajectory level system this is the airline and for the operational level system this is the ANSP. At the strategic level, the responsible actor was discussed less explicitly and is also less evident. It is assumed that the government, potentially in consultation with community representatives, is the most appropriate party for setting the high-level environmental limits and objectives. This responsibility can also be delegated to the airport operator. However, for these targets to be realised in practice, the actual execution of the measures will be at the underlying levels, primarily at the ANSP and potentially also at the airlines.

For the multi-level system, each of the actors has a prominent role. As will become more clear from the description of the actual multi-level optimisation model in Section 6.2, the division of the tasks and responsibilities is similar to the situation at each of the individual levels. For example, it is the airline that is the primary actor with respect to the trajectory level component of the multi-level system, as was also the case for the independent trajectory-level system.

6.1.3 Implementation at the operational level

Although the multi-level system should regard all three levels of aggregation, the operational level is a very important level with respect to the actual environmental impact. Especially decisions that determine runway assignment, runway configurations and flight routing have a major influence on the spatial distribution of the environmental impact. This is one of the reasons that multi-level system will in fact be a DSS at the operational level that interacts with its adjacent levels.

Compared to the operational level, the strategic level is less suitable as implementation level for the DSS. This has already been discussed in Section 5.4. There it has been noted that because of the fact that the generated results are high-level and strategic in nature, they cannot be used directly for the actual runway assignment decisions an air traffic controller would have to make. Consequently, it has been concluded that some form of guidance for air traffic controllers would be required. If such guidance would be in the form of a DSS, this DSS would be classified as a system at the operational level.

An implementation at the trajectory level does seem a possibility. In that case, for the system to be multi-level and use optimisation results from all levels, the operational and strategic level can both issue constraints and targets for the trajectory level optimisation engine. Contradictions in these targets and constraints from both higher levels could be prevented by using an hierarchical procedure. A downside would be that it is not possible for such a system to account for the mutual interaction between different flights. This means that the system, for

example, cannot improve the environmental performance of one flight at the cost of another, in order to achieve a better overall result. As soon as such a system would take multiple flights into consideration, it would no longer be classified as a system at the trajectory level, according to the definitions set in Section 1.4, and become an operational level system as well.

Summarising, it can be concluded that the operational level seems the most appropriate level for implementation. As will be shown in Section 6.2 the final design of the multi-level system is a design in which the strategic level sets the objectives, while the operational level generates the solution, also incorporating solutions provided by the system at the trajectory level.

6.1.4 Assumptions regarding the concept

The assumptions for the multi-level concept are first of all based on a collection of the concept assumptions that have been discussed in previous chapters. For example, when the multi-level concept applies the custom optimised departure profiles, it is still assumed that aircraft navigation and guidance capabilities will support these flight profiles (at present or in the near future), that the concept is workable for both pilots and air traffic controllers, and that the resulting passenger comfort remains within acceptable limits, see also Section 3.1.3. Similarly, assumptions that apply to other elements discussed previously, still remain applicable to the the multi-level concept, at least for as far as these elements are used.

One important assumption has not explicitly been mentioned before but actually applies to all single-level systems presented so far, as well as the multi-level system. It is a requirement for the optimisation that the objectives can be specified explicitly. While this may be evident for a performance area like fuel burn, it is more complicated when the optimisation concerns noise exposure allocation. As discussed in Section 4.4.1 and 4.4.2, noise allocation can be based on noise concentration (optionally within limits), on noise sharing, but also on compromise solutions. The decision support systems presented in this thesis do not determine which policy option should be considered optimal. It is assumed that a particular policy has been chosen. Consequently, the support system applies optimisation to support adherence to this policy. Additionally, the system can also provide insight in the effects of different policies, for example by facilitating a what-if analysis.

In relation to this assumption, it is mentioned here that for the multi-level system, it is sufficient to specify the policy at the highest or strategic level. The hierarchical relation between the levels will ensure that the decisions made at the lower levels will automatically support the high-level policy.

6.2 The multi-level optimisation model

The elements of the multi-level integrated environmental management system will be presented in this section. First, Section 6.2.1 provides a high-level overview of the system. It also shows the overall interaction between the different levels of aggregation. More detail is provided in Section 6.2.2 by showing the allocation model that is implemented at the operational level. In Section 6.2.3, the optimisation objectives and the relation between the strategic and the operational level components of the model are discussed. Finally, an overview of the assumptions and limitations of the model can be found in Section 6.2.4.

6.2.1 High-level overview of the multi-level model

From a high-level point of view, the multi-level integrated system is an operational support system that interacts with the adjacent levels. This is shown in figure 6.1. For a user defined planning period, typically a few hours, the system is tasked to assign a runway, a route and a procedure for all planned flight movements in that period.

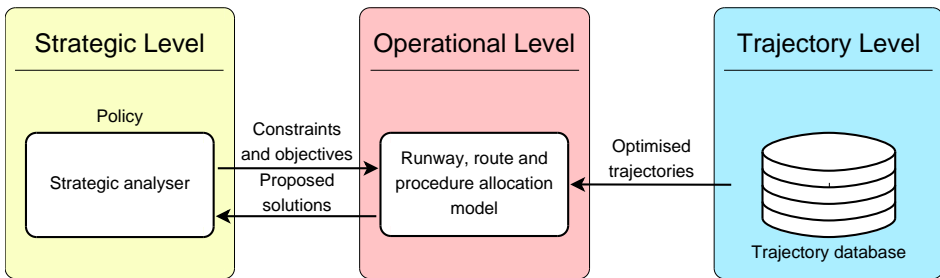


Figure 6.1: High-level overview of the multi-level integrated support system

When considering runway, route and procedure selection, the operational component can only select precomputed trajectory optimisation results from a database. The database itself is populated with trajectories that have been optimised for different criteria. As such, the solution that is computed by the operational component is in fact an optimised combination of solutions generated at the trajectory level. The operational component is responsible for the overall feasibility and optimality of the selected combination of trajectories.

The operational component also interacts with the strategic level. The main reason for this link is to ensure that the selections that are eventually made at the operational level are in line with long-term goals and developments. When considering annoyance as introduced in Section 2.1.4, it is important to realise that this dose-response relationship is based on long-term exposure and should not be

computed based on the noise exposure of the relatively short planning period. In this case, the operational component receives either feedback or guidance from the strategic level component in terms of an (updated) objective function and/or constraints. However, depending on the chosen objective, this interaction may not always be required. For example, when optimising exclusively for minimum fuel burn, the operational component has all required information to perform the optimisation without consulting the strategic component.

The different components will now be discussed in more detail, except for the trajectory level component. This component is the database that has already been presented in detail in Section 3.3.4.

6.2.2 Operational allocation model

For the examples in this chapter, it is assumed that the airport is operating in south flow, with departures from noise preferential runway 24, and if required, also from runway 36L. The planning period, typically a few hours, is divided into blocks of 20 minutes. For each block, the model is forced to select one out of a maximum of six available operational modes. These operational modes are depicted in figure 6.2. It should be noted that not all of these modes are necessarily used in practice.

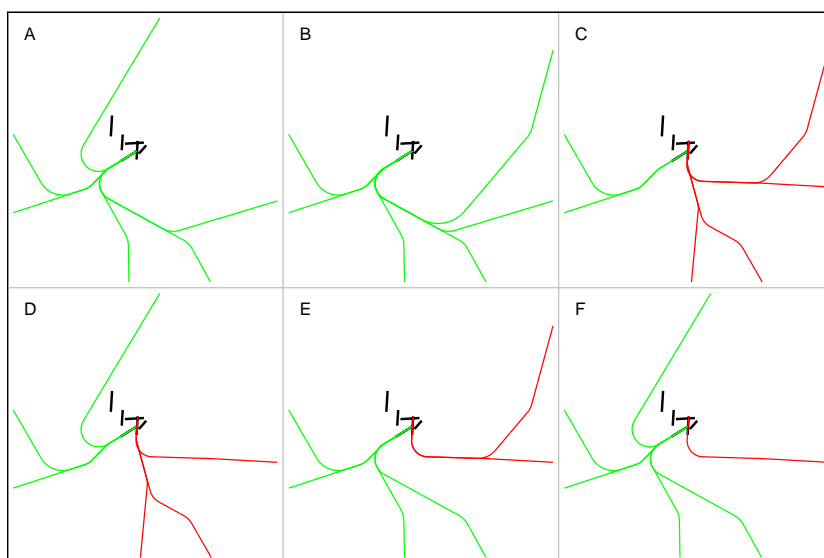


Figure 6.2: *Alternative operational modes for departures in South flow from runway 24 (green) and runway 36L (red)*

The selection of an operational mode results in three related decisions. First

of all, it determines whether departing traffic is managed from a single runway (mode A and B) or from two runways (mode C through F). Secondly, the mode also determines, or at least up to a certain extent, the distribution of flights over the two runways. This can be seen best when comparing mode C to mode F. In the latter case, the 18L runway is only used for departures towards the East, while in case of mode C, this runway is used for all departures, except for those towards the West. Finally, the selection of a mode may sometimes also result in the selection of an alternative routing towards the same TMA exit point. This can be seen when comparing, for example, mode A and B. Traffic towards the North is in mode A routed in a clockwise fashion, while in the case of mode B, traffic towards the North is routed counter-clockwise.

Overall, it can be seen that mode selection has a considerable influence on the geographical allocation of noise as it determines, at least to a large extent, the runway as well as the route allocation for a particular flight. However, mode selection also affects fuel burn, as it also determines which out of potentially three different routes is used towards a particular TMA exit point, where each route results in a different fuel burn.

When defining the problem, certain modes can be made unavailable for selection, either throughout the entire planning period, or during particular blocks. For the examples discussed in this chapter, only the required runway capacity is considered as a constraint, unless mentioned otherwise. The assumed capacity limit for each runway is 36 departures per hour for each runway. Because of this limit, all modes within each 20-minute block that would result in more than twelve departures per runway have been disabled. Modes can be made unavailable for other reasons as well, for example because a particular mode might be undesirable because of interference with the arrival flows, but such factors are not considered in the examples.

Apart from the selection of an operational mode, the model should also select a trajectory from the database for each flight. The selected trajectories should match the chosen operational mode and the considered aircraft type. The model can choose between the five different procedures as defined in Section 3.3.3:

- optimised for fuel only
- optimised for both fuel and noise¹, $k/m = 0.02 \text{ kg}^{-1}$
- optimised for both fuel and noise, $k/m = 0.05 \text{ kg}^{-1}$
- ICAO-A standard procedure
- ICAO-B standard procedure

¹ k/m is the noise-to-fuel ratio or, more precisely, the awakenings-to-fuel cost coefficient ratio as defined in equation 3.6

To summarise, the model selects an operational mode for each block and also selects a departure trajectory for each flight that is compatible with the selected mode. For more details and related equations, see Appendix F.1. For the verification of the model, see Appendix F.4.

6.2.3 Optimisation objective

The general form of the objective function for this problem is defined as:

$$Min : c_1 \cdot \sum fuel + c_2 \cdot \sum Awakenings + c_3 \cdot \sum \Delta A + c_4 \cdot \sum \Delta HA \quad (6.1)$$

Accordingly, the objective is to minimise the weighted sum of the fuel burn, the number of awakenings *Awak*, the additional annoyance ΔA and the additional high annoyance ΔHA of all departures under consideration². For more details concerning the first two elements, see also Section 3.2.2. The other two elements of the objective function are based on the notion of the (fictitious) marginal contribution of a particular flight to total community annoyance, as introduced in 4.4.1.

The total fuel burn and total awakenings can be computed directly by the operational allocation model itself, by summing the fuel burn and awakenings per trajectory as stored in the database of optimised departure trajectories. This implies that if the problem is optimised for fuel and/or awakenings only, the strategic component has no function in the process. For the two marginal annoyance elements of the objective function, the strategic and the operational component of the model interact during the optimisation run, as shown in figure 6.1.

The marginal annoyance elements are computed by the strategic component by comparing the computed overall community annoyance of a reference period to the overall community annoyance of the combined noise exposure of the reference period and the flight under consideration. This comparison is performed for each departure trajectory available in the database. The reference noise exposure is a required input for the strategic component of the model. For the numerical examples presented in this chapter, the reference noise exposure is typically the total noise exposure of all air traffic (including arrivals) for a period of one year. This is considered the minimum period for which the annoyance dose-response relationships can be used. Shorter and longer reference periods are supported as well and might be more appropriate when alternative objectives are considered. If the goal is to spread noise as even and fair as possible, a shorter reference period, e.g. a day, might be desirable.

It should be realised that the strategic component only considers the contribution of a single flight relative to the reference. The method does not consider the additional cumulative exposure of all flights in the considered planning period and

²Please note that for specific runs, one or more multipliers c_i can be set to zero, effectively excluding these elements from the objective function. This means that it is, for example, also possible to optimise for minimum fuel burn only.

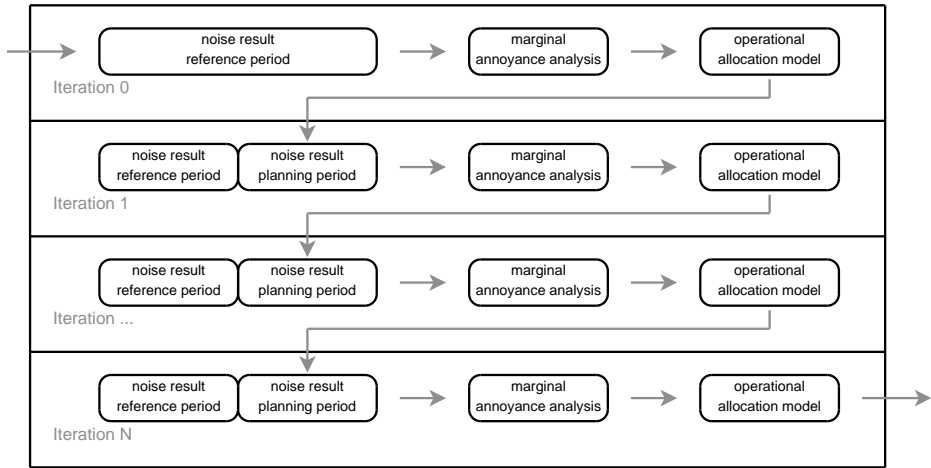


Figure 6.3: *The iterative procedure based on the interaction between the operational and the strategic component of the model*

the potential effect on the noise situation of the reference period is ignored. If the number of flights in the planning period is very small compared to the number of flights in the reference period, the effect of disregarding the additional noise load is expected to be negligible. However, if this condition is not met, it is desirable to also incorporate the noise exposure of the group of flights that the operational component is about to plan. This has been enabled by the implementation of an iterative procedure.

The overall iterative procedure is shown in figure 6.3 and is initialised by computing the marginal additional annoyance of all flights in the database based on the noise exposure of the reference period only. Next, the operational allocation model performs the optimisation, and provides the initial optimum to the strategic component. Again, the strategic component computes the the marginal additional annoyance for each flight, but now based on the noise exposure of the reference period accumulated with the noise exposure expected from the initial solution. The optimisation is performed a second time, using the adjusted values for the marginal additional annoyance. This process is repeated until the solver has generated:

- two consecutive identical solutions (i.e. convergence), or
- a particular solution for the third time.

The latter stopping criteria has been added to prevent the process from continuing indefinitely in case of cyclic behaviour.

6.2.4 Model assumptions and limitations

The operational allocation model only considers departures at this time, and assumes arrivals are managed independently. This should not be considered a limitation that follows from the design of the model. If the modes would be extended to consider the selection possibilities for arrivals too, the model could handle arrivals simultaneously. The primary limitation at this time is that the database of optimised trajectories does not hold arrival trajectories. This means that the database would have to be extended first. Similarly, the developed operational allocation model is currently only able to operate in South flow as shown in figure 6.2. Again, this limitation is not by design, but the result of the contents of the current database.

One of the objectives is to minimise fuel burn. If this objective is used, the model will not only consider the different departure procedures, but will also select a route and corresponding runway. The selection of an alternative runway will most likely result in either an increase or decrease in fuel required for taxiing towards the runway. The model does not consider taxi fuel at this time, but based on the principle of an integrated approach, it is recommended to include this as well.

The model does not consider aircraft separation requirements explicitly. It is assumed that for all modes shown in figure 6.2, air traffic controllers can maintain sufficient separation between the departing aircraft. For the examples that also use mode E and F, this may require navigation performance standards higher than currently in use.

Finally, it could be argued that the multi-level setup demonstrated in this chapter only utilises optimisation at the lower two levels of aggregation. Strictly speaking, optimisation techniques are indeed not employed at the strategic level, even though this level is used to provide guidance for the optimisation running at the operational level. Nevertheless, the current setup is suitable for optimisation at all three levels. One of results as generated by the strategic system presented in Chapter 5 is the optimal annual distribution of flights over the different runways. It would not be complicated to have the strategic component of the multi-level system enforce this optimal distribution upon the operational allocation model, effectively resulting in optimisation at all three levels.

It was decided, however, not to implement this setup. This decision was first of all based on the presumption that annual traffic distributions percentages alone seem too generic to be used as target for the operational allocation model. As there are several alternatives and the best method does not seem evident, it is recommended to further investigate this link between the strategic and the operational model. A second consideration in not using the strategic results concerns the notion that the strategic noise allocation tool shown in Chapter 5 only considers the default departure procedure, whereas the operational allocation model shown in this chapter can select from multiple procedures, including several optimised procedures. Therefore it is also recommended to first extend the strategic

noise allocation program with departure profile selection functionality, before using the results at the operational level.

6.3 Numerical examples of the multi-level integrated system

This section shows the numerical results of the multi-level integrated environmental management system, based on the design presented in the previous section. Before the results are presented, the traffic scenario that has been used to generate the results is described. This is followed by a definition of the reference case. This reference will be used to compare to the different optimisation results. In Section 6.3.3, the results of several single-objective optimisation problems are shown and the multi-objective problems and the corresponding trade-off can be found in Section 6.3.4. Finally, in the last two sections two additional aspects of the strategic-operational interaction and the corresponding impact on the results are described.

6.3.1 Traffic scenario

The traffic scenario is based on a planning period of three hours, or rather nine blocks of 20 minutes. During these nine blocks, 120 flights are scheduled for departure. The first three and last two blocks feature a relatively modest traffic flow and can be managed using one of the single runway configurations if desired. The four blocks in the middle of the planning period are busy and require the use of two departure runways and potentially also a particular distribution of the flights over these runways. Consequently, not all configurations are eligible during these four blocks.

The composition of the fleet is based on the aircraft types that are available in the database of optimised departure trajectories. This means that the traffic is a mix of the Boeing 747-400 and three variants of the Boeing 737. Overall, 80% of the departing aircraft are of one of the Boeing 737 types, and the remaining 20% of the operations are based on the Boeing 747-400. The flights are distributed over the various destination directions and corresponding routing options.

Appendix F.2 specifies the traffic of the traffic scenario, including the number of departures per block, and the distribution of flights over the different aircraft types and destination directions.

6.3.2 Reference case

A reference case has been prepared, based on the typical operational practice at Amsterdam Airport Schiphol. This purpose of this reference case is to obtain a baseline performance, to which the results of the multi-level optimisation can be compared. This baseline performance has been computed using the multi-level

optimisation model, but with severe restrictions on eligible operational modes and departure procedures. In line with operational practice, eligible operational modes were restricted to configurations B or D, depending on required capacity, and departure procedure selection was restricted to the ICAO-A procedure only. The traffic scenario for the reference case is identical to the scenario used for the optimisation cases, as discussed in the previous section.

6.3.3 Single objective optimisations

Based on the objective function shown in equation 6.1, the multi-level optimisation has been performed four times, each time with one of the multipliers in the objective function set to one, and the remaining three to zero. Combined with the reference case, this results in five solutions:

ID	Solution description
Ref	Reference
O1	Minimum fuel burn
O2	Minimum number of awakenings
O3	Minimum number of additionally annoyed
O4	Minimum number of additionally highly annoyed

Figure 6.4 shows the results in terms of the components of the objective function, or performance criteria. All values are normalised and expressed as a percentage of the result of the reference solution. The first observation is that none of the optimised solutions exceeds the reference solution on any of the performance criteria. Apparently, improvements are feasible in all four areas simultaneously. As can be expected, the minimum fuel solution O1 shows the highest reduction in fuel burn. Compared to the reference case, the reduction is 14%, considerably more than the 6% reduction reported in Chapter 3 for trajectory optimisation only³. This higher reduction is the result of selecting the fuel burn optimised departure trajectories, in combination with the optimum airport configuration. This is a clear example of the benefit of optimising at the trajectory and operational level simultaneously.

Although the minimum fuel solution also shows reduction in awakenings, the minimum awakenings solution O2 reduces this even further to only 46% of the reference score. This is a 54% reduction and again more than what should be expected of trajectory optimisation in isolation. Concerning the remaining three indicators that were not part of the optimisation objective, the two indicators for noise annoyance also show a further improvement, while the fuel burn increases slightly compared to the O1 solution.

The two minimum additionally annoyed solutions look very similar, judging from the scores in figure 6.4. While solution O3 indeed has a lower score on

³Depending on the departure runway, the average reduction was between 5.7% and 6.5%, as shown in table 3.2.

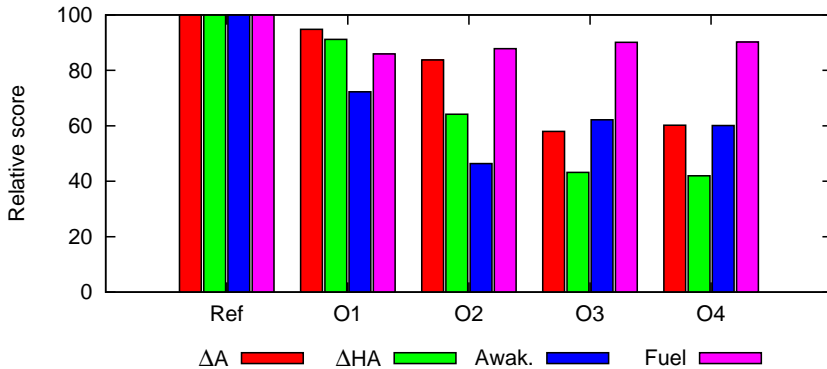


Figure 6.4: Values of the performance criteria for the four optimised solutions, compared to the reference solution

additionally annoyed, and solution O4 has a lower score on additionally highly annoyed, the differences are minimal. Compared to the first two optimised solutions, additional annoyance is reduced considerably, however, this is achieved at an increase in fuel and awakenings.

The results presented here are based on a particular composition in terms of selected configurations and profiles. Figure 6.5 shows the profile distribution for the five solutions. The reference case is restricted to the ICAO-A procedure, and this is confirmed in the figure. The minimum fuel problem is based on a free selection of departure profiles, but the optimum result is based on minimum fuel profiles only. Similarly, the minimum awakenings solution O2 is based exclusively on profiles that, from all profiles available in the database, should feature the lowest number of awakenings. Any other results for these first three solution would have been regarded as a malfunction either at the multi-level model, or at trajectory-level model.

The two minimum annoyance solutions O3 and O4 are not based on a single type of profiles, but are based on a composition of several types. Although the majority of the profiles belongs to one of the types that reduce the number of awakenings, especially solution O3 also consists of minimum fuel profiles and, hardly discernible in figure 6.5, also a single ICAO-B profile. For solution O4, the minimum awakenings profiles becomes more dominant.

Both results are noteworthy. They suggest that optimisation with respect to noise for a single event does not automatically results in profiles that, together, will also minimise the response to cumulative noise exposure. To further validate this finding, but also to eliminate the possibility of some malfunction of the system, an additional test has been performed. To this end, an extra solution has been generated using the objective function of solution O3 (minimum num-




Figure 6.5: Selection of departure procedures for the reference solution and the four optimised solutions

ber of additionally annoyed), but with profile selection restricted to minimum awakenings profiles only. This test showed that the value of the objective function increases by 5% compared to solution O3, confirming that the mix of profile types indeed yields a better solution.

The solutions can also be analysed in terms of selected airport operational modes (figure 6.2) and table 6.1 shows the results. For the reference solution, the planning period starts in single runway configuration B, switches to dual runway configuration D in the fourth period and switches back to B in the eighth period. This result is as expected, because runway use is dictated for the reference solution, as explained in 6.3.2. For the other solutions, the solver is granted much more freedom to select the ideal configurations. For the minimum fuel solution, configuration C appears to be the preferred configuration, with D as alternative only for the two blocks where C cannot provide the required capacity. When comparing these two configurations, the only difference proves to be the runway that is used for departures towards the North. An inspection of the database of optimised departure trajectories confirms that it is indeed slightly more fuel efficient to use the counter-clockwise departure from runway 18L, compared to the clockwise departure route from runway 24.

For minimum awakenings (O2), the solver starts in configuration E and maintains this throughout the entire planning period. The two minimum annoyance solutions O3 and O4 show identical runway use. Both solutions select the single runway configuration B whenever capacity allows, thereby concentrating traffic on what is considered the noise preferential runway. During the busy period, both solution use configuration D, or F if capacity allows this selection.

A final observation that can be made is that the solutions show significant differences with respect to both selected configurations as well as selected departure

	1	2	3	4	5	6	7	8	9
Ref	B		D				B		
O1	C		D	C	D	C			
O2	E								
O3	B		D	F	D		B		
O4	B		D	F	D		B		

Table 6.1: *Applied operational modes for the nine periods for the five different solutions*

profiles. Together with the observations made upon inspection of figure 6.4, this is an indication that these objectives are conflicting. The next section will look into the trade-off between two of the objectives.

6.3.4 Multi-objective optimisation

Based on the results of Chapter 3, Chapter 4 and the previous section, it can be observed that the reduction of fuel burn and the reduction of the noise response are generally conflicting objectives, at least to a certain extent. Here the trade-off between fuel burn and number of additionally annoyed people for the multi-level system is presented.

The trade-off study is based on a number of multi-objective solutions of the multi-level system. The approach that is used for the multi-objective optimisation is the weighted sum method. For the objective function in equation 6.1, the coefficients c_1 and c_3 are varied systematically over a sufficiently wide range of the ratio c_1/c_3 . The expected range of solutions is bounded, on one side by solution O3 and on the other side by solution O1, both presented in Section 6.3.3. Considering these bounds, the corresponding range of interest for the ratio has been found to be:

$$0 \leq \frac{c_1}{c_3} \leq 0.14 \quad c_3 > 0 \quad (6.2)$$

Figure 6.6(a) presents the trade-off in terms of the performance criteria. Similarly to figure 6.4, all values are expressed as a percentage of the result of the reference solution. As expected, this figure shows that for an increasing c_1/c_3 ratio, fuel burn decreases and the additional annoyance increases. It can also be observed that fuel burn decreases steeply at the left border of the figure but hardly decreases further for higher ratios. The additional annoyance indicator, on the other hand, exhibits an increase over the full range of the ratio.

The area near the left border of figure 6.6(a) shows an interesting part of the trade-off, because it shows that it appears to be possible to obtain most of the potential fuel burn reduction, without simultaneously sacrificing the potential annoyance reduction to the same extent. To further explore this, figure 6.6(b) shows a detail of the original figure, for a small part range of the original c_1/c_3 ratio, from 0 to 0.01. The arrows show the trade-off for two different values of




Figure 6.6: Trade-off between fuel burn and additional annoyance, normalised to the reference solution

the c_1/c_3 ratio. The left arrows show that it is possible to achieve 80% of the potential fuel burn reduction at the cost of only 23% of the potential annoyance reduction. Similarly, 90% of the potential fuel burn reduction is possible while sacrificing less than half of the potential annoyance reduction.

6.3.5 Effects of the reference period duration

As explained in 6.2.3, the multi-level optimisation system uses an iterative procedure that allows the strategic component to also consider the actual flights in the planning period, instead of only considering the past reference period. For the results shown in the previous two sections, the duration of the reference period was set to a full year, compared to a planning period of only three hours. Based on this ratio, the expected effect of adding the cumulative noise load of the planning period is negligible. Monitoring of the iterative optimisation procedure indeed confirms that the optimum found during the initialisation run is identical to the solution generated during the first iteration, for all results presented so far. This implies that for the current combination of annoyance objectives and reference period, it could be argued that the iterative procedure is unnecessary.

For shorter reference periods, the effect of the interactive procedure can be observed, as the system needs up to five iterations to reach the final solution. The remainder of this section looks into the effects on the generated solutions.




Figure 6.7: Objective value as function of the applied reference period duration

	1	2	3	4	5	6	7	8	9
8 hours	E								
1, 3 days	B			E				B	
10, 30, 90, 365 days	B			D	F	D	B		

Table 6.2: Applied configurations during each of the nine periods for different reference period durations

For this purpose, the problem of minimisation for additionally annoyed (O3) has been optimised for six more reference periods, ranging from 8 hours to 90 days. Figure 6.7 shows the values of the objective function relative to the reference solution. These results show that for a reference period of 10 days or longer, the result is similar to the result for the full year. For shorter reference periods on the other hand, the objective value is clearly lower, suggesting a better performance for the optimised solution versus the reference.

Before looking into this performance improvement, the solutions have also been compared in terms of selected profile types. This showed that although there is some variation, the differences between the solutions are minor. The applied runway and route configurations on the other hand are less similar. The results are summarised in table 6.2. This table shows that for a reference period of 10 days or longer, the result is as presented in the previous section. For shorter reference periods the optimiser applies less variation and uses configuration B for the periods where a single runway provides sufficient capacity and configuration E when two departure runways are required. If the reference period is reduced further, configuration E is applied during the entire planning period.

These solutions suggest that the optimal strategy for shorter reference periods is based on applying less variation in runway and route configuration. An analysis

based on closely monitoring the iterative procedure shows that the initial solution⁴ in terms of applied configuration is always the same and equal to the results at the bottom line in table 6.2. This suggests that the shift towards using configuration E is exclusively caused by taking the noise exposure of the proposed solution into consideration. As the duration of the reference period is reduced, the relative influence of the proposed solution increases, apparently leading towards a different optimum. In the next section, this behaviour is explored in more detail.

6.3.6 Effect and implications of the strategic-operational link

The design of the strategic-operational interaction in the multi-level integrated system has been discussed in Section 6.2.2. From a high level point of view, the link serves two purposes. First of all, it allows the operational component to use cost coefficients for the optimisation which have been calculated at the strategic level, based on long-term objectives. The second element of the strategic-operational link is the iterative procedure. This can be viewed as a feedback mechanism with the purpose of verifying the impact of the proposed operational optimum against the strategic objectives, before a final decision is made.

The examples in the previous section revealed that for a very short reference period, the iterative procedure has a major impact on the final solution. Simultaneously, the objective value is reduced significantly, as indicated in figure 6.7. Although this further reduction of the objective value is desirable, the nature of the solution gives rise to a suspicion that the applied strategy might not be ideal. This suspicion is based on the observation that the solver applies a uniform route and runway configuration throughout the entire planning period. If the configuration is not varied, there is also no variation in the geographical allocation of the noise. This appears to be a concentration policy, previously warned for in Section 4.4.1. If the multi-level optimisation system indeed applies noise concentration to improve the objective value, this could lead to unacceptably high noise levels in particular areas.

A few additional runs have been performed to test the hypothesis that there is a clear tendency towards a uniform solution, when using a very short reference period. These runs are based on adapting the set of eligible runway configurations and analysing the new optima in terms of applied runway and route configuration. The results of these additional runs indeed support the hypothesis⁵. This means that the improvement in performance due to the iterative procedure is achieved at the cost of a potentially unacceptable concentration of aircraft noise. An obvious solution, as already mentioned in Section 4.4.1, is to limit total noise exposure. Such constraints could be added to the system described in this chapter and the strategic level seems the most appropriate level for implementation. The strategic level already monitors long term noise exposure and can influence the operational component by setting additional constraints or applying penalties.

⁴Shown as iteration 0 in figure 6.3

⁵For a full description of the runs and corresponding results, please refer to Appendix F.3

Although the concentration issue is identified as a concern, it is not considered a major issue, at least not for the setup presented in this chapter. This is first of all because the uniform solution is a dual runway solution⁶. Using two runways, the noise load is already divided over multiple areas. A second and more fundamental observation is that the issue of concentration only arises for very short reference periods. In this case, the planning period at the operational level is three hours, while the reference period is less than three times as long. With such a short reference period, the strategic component does not seem to meet the definition of the strategic level and is not able to monitor what could reasonably be defined as long-term noise exposure. This means that, although concentration still might be a concern, in this case it is most of all an artefact of the sensitive study performed in Section 6.3.5.

6.4 Discussion

In this chapter, a multi-level integrated decision support system for airport environmental management is presented. The implementation that has been chosen is an operational decision support system that interacts with the adjacent levels. The operational component of the system first of all optimises environmental impact and the corresponding geographical allocation by selecting departure procedures for each specific flight. Simultaneously, it also influences aircraft routing and runway usage, by selecting airport operational modes.

The interaction of the operational level component with the trajectory level component is based on selecting pre-processed, optimised departure profiles from a database. Selected trajectories do not only meet the constraints imposed by the operational level, but also support the operational level in reaching its optimisation objectives. At the same time, it should be realised that this support could probably be improved. An inherent weakness of the chosen mechanism is that the actual trajectory optimisation has been executed using trajectory level objectives instead of higher-level objectives. Formulated alternatively, the operational component can only select the best available trajectory, instead of requesting a trajectory that is fully tailored towards its current objectives.

The operational level component also interacts with the strategic level, but this link is of a different nature. This link first of all enables the possibility of taking strategic objectives into consideration at the operational level, which is eventually the level where the allocation decisions are made. An iterative procedure is used to ensure that the decisions that are about to be taken are still in line with the strategic objectives. It has been shown that this procedure works as intended and can actually further improve a solution by regarding the additional environmental impact of the planning period which is to be optimised. A necessary condition for this feedback mechanism to work is that the environmental impact of the planning

⁶Additional tests have shown that the solver actually prefers a uniform single runway solution, but this is not feasible due to capacity constraints.

period is not negligible compared to impact of the considered reference period. A second issue identified is that for the dose-response relationship that was selected as strategic objective function, there is a risk that the additional improvement in objective value is achieved at the cost of an unacceptable concentration of aircraft noise.

The multi-level optimisation system has been used to generate several numerical results. The first results presented were based on several single-objective optimisation problems. By comparison to a reference case, it was shown that the optimised solutions always outperform the reference solution on all regarded performance criteria. This means that improvements are feasible in all four regarded performance areas simultaneously. Additionally, the benefits of optimisation at multiple levels simultaneously have been demonstrated by comparing the multi-level results to results obtained using the trajectory level support system only.

To explore the trade-off between the reduction of fuel burn and the reduction of the noise response, a multi-objective optimisation has been performed as well. The presentation of the results focused on a particularly interesting part of the overall trade-off, where it was possible to obtain the majority of the potential fuel burn reduction, without simultaneously sacrificing most of the potential noise annoyance reduction.

Although the multi-level integrated environmental management system demonstrated in this chapter is in its current state already able to generate some interesting results, it should be recognised that the possibilities for further development of the system have certainly not been exhausted. Apart from the extension that would be required to also cover arrivals, there are many other options. Currently, aircraft noise is the performance area with the most extensive coverage in the system. Noise is regarded at all three levels of aggregation and it is possible to use either different or multiple performance indicators for aircraft noise simultaneously. It could also be argued that aircraft gaseous emissions are considered, albeit in a limited way, using the fuel burn as performance indicator. However, an extension to include full functionality for local emissions (see Section 3.2.3) or even air quality should be considered. Third party risk could also be added, as it is currently not regarded at all in the multi-level system, even though it was a performance area in the previous chapter.

Further work could be performed on one of the limitations identified above, related to the fact that the actual trajectory optimisation has been executed using trajectory level objectives instead of higher-level objectives. The results presented in Section 6.3.3 illustrate this issue and indeed show that the single event minimum awakenings objective is not a perfect proxy for the multi event minimum additional annoyance objective. It is recommended to explore the feasibility of extending the database of optimised trajectories with results based on the minimum additional annoyance objective and to analyse to what extent the multi-level optimisation results can be improved. In addition, it would also be recommended to investigate how often the trajectory level results can and should be recomputed, given the

fact the optimised trajectories will immediately be outdated because of the ever-changing reference noise exposure.

Further recommended extensions and explorations are related to the interaction between the strategic and the operational level. A first recommendation would be to deal with one of the limitations mentioned in Section 6.2.4. Here it was discussed that given the applied objective at the strategic level, this level does not apply optimisation techniques to derive the guidance for the operational level. Instead, the strategic level basically provides up-to-date sensitivity coefficients to be used by the operational allocation model. The alternative would be to first generate a high-level optimisation result at the strategic level and to adjust the objective of the operational level such that it will actively support the realisation of the strategic optimum.

To create the opportunity to study the long-term behaviour of the system, it is recommended to add the functionality required to simulate the use of the multi-level system over a longer period of time (e.g. a year), based on consecutively generating many solutions for short planning periods, such as a few hours, as used as planning period during for the examples in this chapter. As the simulation progresses, the noise load of the reference period evolves, due to the decisions made by the operational allocation model. This will affect the guidance provided by the strategic component and is therefore also expected to affect the decisions made at the operational level.

Conclusions and recommendations

This chapter presents the conclusions and recommendation. First, the main conclusions from the previous chapters are summarised, based on the work on the trajectory, operational, strategic and multi-level system (see figure 1.1) as presented in Chapter 3 through 6. These conclusions are complemented with the overall conclusions that can be drawn concerning the use of an integrated support system for the airport environmental management process, as proposed in the introduction of this thesis. The second part of this chapter provides an overview of the recommendations. These recommendations do not only address the potential options for further development of the presented support systems, but also propose future practices and research efforts.

7.1 Conclusions

Trajectory level

An aircraft trajectory optimisation tool has been adapted to investigate the benefits and consequences of a concept called custom optimised departure profiles. Under this concept, each departing aircraft would execute a departure profile that has been optimised for a specific flight, while still using predefined and published departure routes. This concept was presented as an alternative to full trajectory optimisation, involving optimisation of both the vertical profile and the route. It is expected that not optimising with respect to the routing will reduce the potential environmental benefits. However, the concept does not suffer from a considerable increase in airspace complexity. Especially this benefit could be of paramount importance for a potential implementation.

The profile optimisation results show that the use of optimised profiles leads to small reductions in fuel burn and impressive reductions with respect to the applied indicator for community noise when compared to two different standard-

ised departure procedures. The optimised trajectories feature a speed profile that differs from the current situation, and this observation led to a study into the effects of the optimised profiles on runway capacity. The results of this capacity study show that using the optimised departure profiles or mixing traditional and optimised departures procedures does not lead to a reduction in runway capacity when compared to the baseline situation. There may even be opportunities to increase throughput at the departure runway.

Regarding the tool itself, and when compared to the four-dimensional trajectory optimisation tool it was originally based upon, the profile-only version is several times faster and also more robust. These two characteristics facilitated the generation of a database containing more than 500 optimised departure profiles, based on different routes, aircraft types, take-off weights and optimisation criteria. Using this set of results, it was first of all shown that reductions in fuel burn and the noise indicator cannot only be obtained for a few example cases, but can be realised for the full set of optimised profiles. While the existence of this set of trajectories was a prerequisite in reaching this conclusion, the resulting trajectory database is also considered an important element because it enabled the runway capacity study and is also used for the multi-level optimisation study.

Operational level

An existing arrival management concept has been extended to include awareness of noise impact. Where an arrival manager is a typically a decision support tool for air traffic controllers used to assist in creating an efficient flow of aircraft towards the runway, the noise-aware version also assists in obtaining a particular noise allocation strategy. The desired noise allocation result is achieved through the selection of fixed arrival routing options.

The developed tool has been used to investigate the trade-off between noise and delay in different demand situations, from relatively busy to relatively quiet situations. An analysis of the results showed that in some cases the performance indicator for noise can be improved without any sacrifice in efficiency. This implies that it could even be advantageous to use a noise-aware arrival manager in cases where an increase in delay is deemed unacceptable.

The results also show that when the relative importance of the noise part of the objective function is increased further, delay also increases. For low demand situations, the resulting delay can even exceed the delay in a non-managed (i.e. first come, first serve) situation. For the busy, high demand situations the trade-off is more favourable.

The results also revealed that when using a noise indicator based on a single-event dose-response function and this indicator is dominant in the objective function, the model will tend to concentrate noise in the least noise-sensitive region. Assuming this may not always be considered acceptable, the obvious alternative would be to use an indicator based on cumulative noise exposure. An analysis into this alternative showed, however, that typical dose-response functions based

on cumulative exposure are no solution to this phenomenon.

Strategic level

A strategic runway allocation model has been developed, which can assign all annual flights to the runways of an airport, while respecting the feasibility of the applied runway configurations concerning required capacity and wind limits. The objective of this model is to minimise a particular combination of runway delay, community noise exposure and third party risk. The ultimate trade-off between the three objectives is guided by a graphical interactive optimisation procedure, and results in a balanced runway usage scheme.

Concerning the results, the optimiser shows significantly improved results with respect to the reference situation. The cumulative overall risk is reduced by about 30%. A similar reduction is observed for the performance indicator estimating annoyance due to aircraft noise. However, it should be acknowledged that the comparison between the optimised and reference situation may not be a completely fair comparison, because of some assumptions and simplifications in the design of the optimisation model. Still, runway allocation optimisation targeted at minimising environmental impact yields promising results at the strategic level.

A second important observation is that the results are high-level and strategic in nature. As such, the results cannot be used directly for the runway assignment decisions air traffic control would have to make during day-to-day operation. It is expected that some form of decision support is required to support implementation of strategic results in practice. Such a support system would operate at multiple levels of aggregation.

Multi-level

A fourth and final multi-level model was developed, by combining selected elements of the previously presented systems, as well as adding new components. The resulting implementation is in fact an operational decision support system that interacts with the adjacent levels. The operational component of the system first of all optimises environmental impact and the corresponding geographical allocation by selecting pre-processed optimal departure profiles for each specific flight. Simultaneously, it also influences aircraft routing and runway usage, by selecting airport operational modes.

Single-objective optimisation runs have been performed first, aimed at the minimisation of either fuel burn or one of the three noise-related performance criteria. The results show that all of the optimised solutions outperform the reference solution on any of the four regarded performance criteria. This observation means that improvements are feasible in all four performance areas simultaneously. Furthermore, the benefit of the multi-level optimisation has been shown by comparing the multi-level results in terms of fuel burn and noise response to results obtained using the trajectory level support system only. Finally, the res-

ults revealed that minimization of cumulative, multi-event noise exposure is not automatically achieved by combining different flight trajectories that have been optimised for single-event noise exposure.

Multi-objective optimisation has been performed as well, to explore the trade-off between the reduction of fuel burn and the reduction of the noise response. The results yielded the observation that it was possible to obtain the majority of the potential fuel burn and resulting cost reduction, without simultaneously sacrificing most of the potential noise annoyance reduction.

The integrated support system

The main goal of this research has been defined as to identify whether an integrated support system can improve the airport environmental management process, assuming such a system would feature three main capabilities: (i) the capability to use formal optimisation in order to select the best mitigation options, (ii) the capability to evaluate multiple performance areas simultaneously, and (iii) the capability to evaluate mitigation options at multiple levels of aggregation simultaneously.

The use of optimisation in order to maximise the efficiency of mitigation measures and minimise the corresponding environmental impact has been demonstrated numerous times before. Chapter 3, 4, and 5 provide references to examples that are available in the literature, related to the trajectory level, the operational level and the strategic level, respectively. These examples from literature, as well as the applications presented in this thesis, clearly illustrate the benefits of using optimisation through the impact reductions that can be achieved.

The second main capability is the use of an integrated approach in the sense of evaluating multiple environmental performance areas simultaneously, such as noise and fuel burn. This principle is also occasionally promoted in the literature and demonstrated multiple times in this thesis, in Chapter 3, 4, 5, and 6. It may prevent the situation that mitigation measures directed towards one specific area will result in negative side-effects in other areas, while the decision maker might not even be aware of this consequence. The examples presented show that when this principle is combined with optimisation, the trade-off between such conflicting objectives can be explored. The results show that in some cases it is possible to improve all areas simultaneously, compared to a non-optimised baseline situation. Where this cannot be achieved, the trade-offs may still be very advantageous, meaning that it might be possible to achieve fairly substantial reductions in one performance area, at the cost of only a marginal increase of impact in the conflicting performance area.

The third and final capability concerns the use of an integrated approach in the sense of considering multiple mitigation options at different levels of aggregation simultaneously. The evidence that this capability can contribute to improving the airport environmental management process is more limited, as this was demonstrated with the final, multi-level support system only. Synergy bene-

fits were demonstrated using a powerful combination of optimisation of departure profiles at the trajectory level, and runway and routing configurations at the operational level. This resulted in impact reductions that cannot be achieved by these measures in isolation, or even through combination of these measures when not coordinated properly.

It is important to realise, however, that all three of the desired capabilities play an important role in the example of the multi-level system, not just the evaluation of multiple measures at multiple levels. Optimisation techniques are used to find the optimal combination of measures already optimised, and the capability to evaluate multiple performance areas facilitates the trade-off and helps to identify potentially unwanted negative side-effects. It is therefore especially the combination of these three capabilities integrated into a support system that can contribute to improving the airport environmental management process.

7.2 Limitations and Recommendations

Concerning the custom optimised departure profile application at the trajectory level, it has already been mentioned in Chapter 3 that all noise-optimised profile results were based on a single dose-response relationship and that the model could be extended with other dose-response relationships. Additionally, the results of Chapter 6 showed that the single event minimum awakenings relation used was not a perfect proxy for the multi event minimum additional annoyance objective. It is therefore recommended to explore the use of alternative dose-response relationships, especially one based on cumulative noise exposure. Consequently, these new optimised departure profiles can be used to analyse to what extent the results for the multi-level optimisation can be improved as well.

For the optimised departure profiles, it was assumed that the concept will be supported by (near) future levels of flight deck automation. It is recommended to look into the validity of this assumption by defining the technical requirements and to compare these requirements with the capabilities of the flight management systems available today or in the next few years. A related assumption was that, with sufficient support from the flight deck systems, the concept would be safe, and acceptable in terms of pilot workload and passenger comfort. However, for the concept to reach the next state of maturity, this would need to be confirmed by further research.

Assuming a positive outcome of the safety, workload and passenger comfort analysis, a next step could be to perform a simulation experiment, involving an ANSP, a government representative and multiple airlines, the latter preferably each with their own profile optimisation capabilities and corresponding tools. There can be multiple reasons for conducting this experiment, such as to verify there are no incompatibilities between the profiles as planned by the different airlines, or to test the protocol used by the ANSP to communicate the profile requirements. From an environmental point of view, however, it would also be

very interesting to analyse the reactions of airlines in terms of selected departure profiles to changes in policy.

As described in Chapter 4, the noise-aware arrival management model has been used to generate results for high, medium and low arrival rates, but all in isolation. The current model could also be used to study its behaviour and resulting performance in case of alternating periods of high and low arrival rates, a situation not uncommon in practice. Given the results seen so far, it is expected that the model will prioritise noise impact during the busy periods and delay during the quiet periods, which may very well result in a highly favourable overall trade-off between the delay and noise objective, when considered over the entire period.

The results generated so far by the research model also showed that even when sacrificing efficiency is not unacceptable, there seem to be opportunities for noise allocation, especially during the busier periods. Given these results, it should at least be considered to extent existing or newly developed arrival managers with the functionality required to assess noise impact during the scheduling process. This recommendation only holds for situations where arrival managers make use of multiple but fixed alternative arrival routes, and only at locations with a policy in place that enables the model to determine what is considered optimal in terms of noise allocation.

Concerning the strategic noise allocation model outlined in Chapter 5, this model could be developed further by addressing some of the limitations mentioned earlier. More specifically, extending the model such that it could consider visibility conditions and runway operation in mixed mode can improve the accuracy and reliability of the results, as well as the applicability with respect to other airports.

What remains is that the results generated by the strategic noise allocation model are still high-level, limiting the applicability in the day-to-day operation. For that reason, it is recommended to also further explore the interaction between the strategic and the operational level. Instead of the multi-level approach tested in Chapter 6, where the strategic level only generates up-to-date sensitivity coefficients to be used by the operational allocation model, the recommendation is to first generate a high-level optimisation result at the strategic level. Consequently, the operational support systems should actively guide realisation of the strategic targets, for example by adding an optimisation objective that minimises the continuously developing deviation from the computed strategic optimum. At the same time it is recommended to add the functionality required to simulate the use of the multi-level system over a longer period of time, like a year, based on consecutively planning short periods.

Even with more research into the strategic-operational interaction, the possibilities for further development of the multi-level system have certainly not been exhausted. Related to the functionality of the system, two examples of interesting additions would be to add a user interface and to also apply optimisation

for arriving flights. With these additional optimisation possibilities for arriving flights, the model will probably be able to improve the overall results further. Also in terms of environmental impact analysis the model is not complete, as aircraft emission analysis is limited to fuel burn only, and third party risk is not considered at all. Adding these analysis capabilities increases the insight in the trade-offs between the different environmental performance areas and can prevent that decision makers might be uninformed about potential negative side-effects.

Even without all these potential extensions, it was concluded in the previous section that the integrated system can improve the airport environmental management process. Each of the three identified main system capabilities (the capability to use optimisation techniques, the capability to evaluate multiple performance areas, and the capability to evaluate mitigation options at multiple levels of aggregation simultaneously) contributes to this improvement. For this reason it is recommended to further develop these capabilities and to work on the implementation issues, and ultimately to make use of these capabilities, not necessarily at every airport, but especially for airports where the efforts towards environmental impact mitigation are considered significant and complex.

Aircraft navigation

A.1 Traditional aircraft navigation

The first aircraft that flew more than a century ago were navigated visually only. In flight, pilots would study their surroundings looking for landmarks such as rivers, cities and railways to determine their position and direction. Only a few basic instruments like an altimeter, airspeed indicator and a magnetic compass were available to assist the pilots. In the 1920s, the first navigational aids were installed in the US in the form of lighted beacons along important airmail routes^[76]. Although this enabled night-time flights along a few airways, poor visibility conditions could still seriously limit the effectiveness of this navigation system.

Airborne radio navigation was introduced in the 1930s. Using a manually rotated loop antenna and a radio receiver, pilots could determine the direction from which a particular radio signal was originating. By tuning in to commercial radio station with a known broadcasting location, these stations could be used as a homing beacon. Shortly thereafter, dedicated radio beacons (Non-directional Beacon (NDB)) were installed, replacing the use of commercial radio stations. The airborne system also evolved and was automated, resulting in the Automatic Direction Finder (ADF), a system still available in airplanes today^[99;113].

A more flexible and reliable radio navigation system was deployed after the Second World War. Instead of a non-directional beacon, a directional beacon or VHF Omni-range (VOR) was used. Based on phase comparison of two signals transmitted by the VOR station, the airborne equipment is able to determine the stations radial the aircraft is flying on. VOR stations are commonly co-located with Distance Measuring Equipment (DME) stations. This allows equipped aircraft to simultaneously also determine the distance to the VOR-DME station, based on the time delay between the interrogation sent by the aircraft and the ground station reply received by the aircraft. As this yield both range and azimuth relative to the position of the station, a single station is sufficient to obtain

a position fix^[99;113].

In 1960 ICAO selected the VOR as standard short-range navigation aid^[99]. Airways were typically defined as straight lines from one VOR station to the next and even today, many airways are still based on this principle. Figure A.1 shows an example of some VOR stations and how these stations can be used to navigate from A to B.




Figure A.1: A route from A to B based on navigation from VOR to VOR

This traditional navigation principle where aircraft would fly from one navigation aid to the next has some downsides. First of all, all traffic is concentrated on a limited number of airways and as air traffic grew, this could lead to capacity problems. Although additional airways could be defined through previously unused airspace, this would require building additional stations at considerable cost. A second disadvantage is that the system did not offer direct routes, as can also be seen in figure A.1. A direct route from A to B would be shorter, saving time and fuel. A solution to both limitations came from the domain of computer technology in the form of a system called RNAV.

A.2 Area Navigation (RNAV)

Area navigation is based on a different principle when compared to the previous beacon-based approach. RNAV systems simultaneously combine inputs from multiple navigation systems. Depending on the RNAV system capabilities, these inputs can be from ground-based navigation aids, such as (multiple) VOR and/or DME stations, but can also come from on-board inertial navigation systems, as well as from satellite navigation systems. RNAV systems use this to facilitate navigation along any desired route, no longer bound to the straight lines between the ground-based navigation aids¹^[113].

Waypoints are geographical positions that have been defined in RNAV systems and although these waypoints may be co-located with beacons, this is not required.

¹RNAV was originally also referred to as random navigation, explaining the acronym.

Waypoints can be defined by the pilots, but air navigation relies primarily on waypoints that have been published by ANSPs and these waypoints are typically available in a regularly updated navigation database on-board the aircraft. Pilots can define routes from their origin to their destination through the specification of a sequence of waypoints. The route segments between consecutive waypoints are called RNAV legs^[113].

The RNAV systems available on modern airliners can also compute navigation directions to remain on the planned route, both in the vertical plane (V-NAV) as well as in the horizontal or lateral plane (L-NAV). These directions can be presented to the pilots, but can also be provided to the autopilot systems^[99]. This setup allows modern aircraft to fly a pre-planned route automatically.

The standard departure routes out of an airport, as well as the standard arrivals routes towards an airport can also be flown with the V-NAV and L-NAV capabilities. Especially these functions are of interest in relation to airport environmental impact. RNAV routes offer more flexibility and can be flown with a higher precision. This yields opportunities for noise abatement, for example when overflying densely populated areas can be reduced.

Additional information on the custom optimised departure profiles

B.1 Optimal control

The aircraft trajectory optimisation problem presented in section 3.2 is an optimal control problem. Optimal control means optimally steering a dynamical system while satisfying boundary conditions, state constraints and control constraints. The general mathematical formulation is as follows:

$$\text{Minimise} \quad J = \varphi [\mathbf{x}(t_f), t_f] + \int_0^{t_f} L [\mathbf{x}(t), \mathbf{u}(t), t] dt \quad (\text{B.1})$$

subject to:

$$\dot{\mathbf{x}} = \mathbf{f} [\mathbf{x}(t), \mathbf{u}(t), t] \quad (\text{B.2})$$

and:

$$\begin{aligned} \mathbf{b}_{lower} &\leq \mathbf{g} [\mathbf{x}(t), \mathbf{u}(t), t] \leq \mathbf{b}_{upper} \\ \mathbf{b}_{0,lower} &\leq \mathbf{b} [\mathbf{x}(t_0)] \leq \mathbf{b}_{0,upper} \\ \mathbf{b}_{f,lower} &\leq \mathbf{b} [\mathbf{x}(t_f)] \leq \mathbf{b}_{f,upper} \end{aligned} \quad (\text{B.3})$$

where $\mathbf{x}(t)$ are the state variables, $\mathbf{u}(t)$ are the control variables and t is the time. The left part (or Mayer part) of cost function B.1 represent the end cost of the problem, where the right part (or Lagrange part) represents the running cost. Differential equation B.2 states that the system is subject to its dynamics, typically the equations of motion. The remaining equations set the constraints on the states, the controls and the time, and also enforce the initial and final conditions.

B.2 Collocation and nonlinear programming

Multiple methods are available for solving optimal control problems. The method that is employed by EZopt, the software package that has been used, is a direct approach based on using Nonlinear Programming (NLP) and collocation^[59]. This is a numerical approach that provides an approximation to the optimal solution. The main purpose of this method is to transform the optimal control problem into a NLP problem and to solve the latter.

The method first discretises the trajectory, as already discussed in section 3.2.4. This results in a representation based on piecewise constant controls and piecewise linear states¹. The collocation method also defines collocation points, exactly halfway between the nodes. Figure B.1 shows a part of the discretised trajectory, with two nodes and one collocation point halfway between these two nodes.

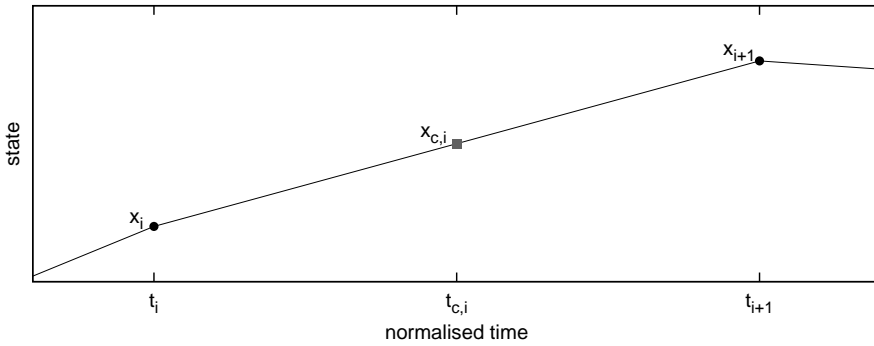


Figure B.1: *The collocation point halfway between nodes*

The time, states and time derivatives of the states at the collocation points are defined as:

$$\begin{aligned}
 t_{c,i} &= \frac{t_i + t_{i+1}}{2} \\
 \mathbf{x}_{c,i} &= \frac{\mathbf{x}_i + \mathbf{x}_{i+1}}{2} \\
 \dot{\mathbf{x}}_{c,i} &= \frac{\mathbf{x}_{i+1} - \mathbf{x}_i}{t_{i+1} - t_i}
 \end{aligned} \tag{B.4}$$

and evaluating \mathbf{f} from equation B.2 at a collocation point i yields $\mathbf{f}_{c,i}$. Now define

¹According to the method as used by EZopt. The original method^[59] is based on linear controls and cubic polynomials to represent the states.

the defects at the collocation points as:

$$\begin{aligned}\Delta_i &= \mathbf{f}_{c,i} - \dot{\mathbf{x}}_{c,i} \\ &= \mathbf{f}_{c,i} - \frac{\mathbf{x}_{i+1} - \mathbf{x}_i}{t_{i+1} - t_i}\end{aligned}\tag{B.5}$$

When the defect Δ is set to be zero, the resulting nonlinear algebraic equality constraint is an approximation to the solution of equation B.2. When using this approximation, the differential equation has been eliminated from the problem formulation. This means that the discretised version of the problem has effectively been reduced to a NLP problem.

B.3 Solving NLP Problems

EZopt is supplied with an NLP solver that is called NPSOL^[55]. NPSOL itself uses a Sequential Quadratic Programming (SQP) algorithm to solve the NLP problem. As the name suggests, this is an iterative procedure that involves solving multiple Quadratic Programming (QP) problems. The process starts with an initial solution as provided by the user. Based on this initial solution, NPSOL creates a QP subproblem that is basically a local approximation of the NLP problem. This QP problem is solved using a QP solver. The solution of the QP problem is interpreted as a search direction. NPSOL improves the initial solution by taking a step in this search direction and creates the next QP subproblem. This iteration process continues with the intention to converge to a point that satisfies the optimality conditions.

For the EZopt user, familiarity with NPSOL and the employed algorithm is not required. However, during the optimisation run, NPSOL will provide details of the iteration process, such as the iteration number, the step size, the current value of the (augmented) objective function and the status of the convergence tests. These parameters can be used to monitor the optimisation process if desired.

B.4 Aerodynamic modelling

The drag of the aircraft must be computed to solve the equations of motion (equation 3.1 in section 3.2.1). The computation of the drag is performed using lift-drag polars of the different aircraft types involved. First, the lift coefficient C_L is determined using:

$$C_L = \frac{W}{S} \frac{2}{\rho V^2 \cos(\mu)} \tag{B.6}$$

where S is the wing area of the current aircraft and ρ is the density of air for the current altitude in the International Standard Atmosphere.

Next, the lift coefficient is used to compute the drag coefficient from a lift-drag polar, as shown in figure B.2(a). Finally, the drag coefficient is used to compute


the drag D :

$$D = \frac{1}{2}\rho V^2 C_D S \tag{B.7}$$

which completes the drag computation itself.

Lift-drag polars are typically valid for a particular aircraft configuration. This means that the selection of different flap setting or the extension or retraction of the undercarriage would require switching drag-polars. Several configurations of NOISHHH indeed feature this functionality. The EZopt software allows the definition of multi-phase problems. For each phase, users can define different equations, including the equations that model the drag characteristics. If desired, constraints can be added to ensure that a particular flap settings matches a particular range of airspeeds. During the optimisation run, the solver can also optimise the time to switch from one phase to the next. As such, NOISHHH generates the flap schedule for each optimised trajectory.

This approach has not been used for the concept of custom optimised departure profiles. Instead, a single lift-drag polar (figure B.2(a)) was used for each aircraft type. This curve has been constructed from the drag-polars of the different aircraft configurations, as shown in figure B.2(b). The overall drag-polar in figure B.2(a) represents the minimum drag envelope of the aircraft type.



(a) Example of a lift-drag polar

(b) Lift-drag polars for multiple configurations

Figure B.2

Advantages and disadvantages can be identified for both approaches. Ulti-

mately, the decision to use the continuous drag polar for this application was made to reserve the use of different phases for the implementation of the conditional operational constraints, see section 3.3.1. If both the flaps selection and the conditional operational constraints would be modelled using phases, the problem would become far more complex, as the modeller needs to specify the sequence of the different phases beforehand.

B.5 Influence of wind fields

The equations of motion and most of the examples presented in chapter 3 are based on the assumption of no wind. The results in section 3.3.6 on the other hand have been generated with a slightly modified version of the profile optimisation tool. This modification allows the tool to account for effects of wind on the motion of the aircraft.

The implementation of this modification is still based on the pseudo three-dimensional situation and related equations of motion, as described in section 3.2.1. The aircraft is also still assumed not to deviate from the specified ground track. Figure B.3 shows the wind triangle for this situation.




Figure B.3: *The wind triangle showing the ground speed, wind speed and air speed in the horizontal plane*

In this situation, the course is always known for any given position. Together with the wind speed, wind direction and the true airspeed in the horizontal plane, the resulting ground speed V_{ground} can be solved using the law of cosines. The resulting ground speed is used to update state s :

$$\dot{s} = V_{ground} \tag{B.8}$$

replacing equation 3.3.

The wind model itself is at this point very uncomplicated. It is based on an input file listing altitudes with corresponding wind directions and velocities. For the actual altitude, the wind speed and direction are obtained through linear interpolation, if required.

B.6 Overview of optimised profiles

Table B.1 shows the flights that have been optimised for this study, using the default conditions of a final altitude of 6000 ft and no wind. All of these 21 flights have been computed for five different sets of constraints and optimisation criteria (ICAO-A, ICAO-B, Fuel, Noise (0.02), and Noise (0.05)).

runway	route	aircraft type
18L	ANDIK2E	B737-300
	ANDIK2E	B747-400
	ARNEM2E	B737-300
	ARNEM2E	B737-800
	ARNEM2E	B747-400
	LOPIK2E	B737-300
	LOPIK2E	B737-800
24	ANDIK1S	B737-300*
	ANDIK1S	B747-400*
	ARNEM1S	B737-300*
	ARNEM1S	B737-800*
	ARNEM1S	B747-400*
	BERGI1S	B737-700*
	BERGI1S	B747-400*
	LOPIK1S	B737-300*
	LOPIK1S	B737-800
	SPYKER1S	B737-300
	SPYKER1S	B737-700
	SPYKER1S	B747-400
	VALKO1S	B737-700
	VALKO1S	B737-800*

Table B.1: *Overview of all optimised flights. Combinations marked with * are part of the subset used for the capacity simulations.*

Additionally, all of these combinations have been computed for six different take-off weights, depending on the aircraft type, as specified in table B.2. In total this leads to 630 optimised profiles.

B.7 Model verification

The model that was developed for generating the custom optimised departure profiles is based on NOISHHH (see Section 3.2) and most of its components have been applied and verified before. Within the scope of this project, verification activities focused on new components of the tool. For example, two different types of checks have been performed to confirm that the actual location of the

	B737-300	B737-700	B737-800	B747-400
1	49500	56500	60500	274000
2	51500	59000	64000	298000
3	53500	61500	67500	322000
4	55500	64000	71000	346000
5	57500	66500	74500	370000
6	59500	69000	78000	394000

Table B.2: *Modelled take-off masses in kg for different aircraft types*

aircraft, which is required for the noise calculation, is correctly resolved from the new state parameter that records the along-track distance.

No such detailed checks have been performed for the tool as a whole. Still, it can be stated that the overall results in terms of fuel burn, awakenings and transit time are comparable to results obtained with older configurations of NOISHHH for cases based on the same aircraft type and similar routing.

The second model, used for the runway capacity simulations, was developed from scratch and verified by visualising several of the generated solutions. For each solution, all aircraft positions were plotted as on a radar screen, while moving on a fast time scale. This animation was used to check: (i) that the separation requirements between the different pairs of aircraft are not violated, and (ii) that each aircraft is at minimum separation with another aircraft at least once during the run.

Linear Programming

C.1 General introduction

Linear Programming or LP is a mathematical optimisation method for problems with a single linear objective function and linear constraints. The mathematical formulation of the standard form of an Linear Programming (LP) problem is^[69]:

$$\text{Maximise} \quad Z = c_1x_1 + c_2x_2 + \dots + c_nx_n \quad (\text{C.1})$$

subject to the restrictions:

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 + \dots + a_{1n}x_n &\leq b_1 \\ a_{21}x_1 + a_{22}x_2 + \dots + a_{2n}x_n &\leq b_2 \\ &\vdots \\ a_{m1}x_1 + a_{m2}x_2 + \dots + a_{mn}x_n &\leq b_m \end{aligned} \quad (\text{C.2})$$

and:

$$x_1 \geq 0, \quad x_2 \geq 0, \quad \dots, \quad x_n \geq 0 \quad (\text{C.3})$$

Problems in other forms can still be a legitimate LP problem, if they have one or more of the following characteristics:

- Minimise the objective function
- Greater-or-equal constraints
- Equality constraints
- Negative lower bounds on the variables

C.2 Integer Programming

There are several variants on pure LP problems as introduced in the previous section^[69]. If all of the variables of the problem or not real but integer, the problem becomes an Integer Programming or IP, sometimes also called integer linear programming. A special variant of the integer variable is the binary variable, which can only take the values of zero and one. If a problem only consist of binary variables, it becomes a Binary Integer Programming or BIP problem.

Mixed forms of which only a part of the variables is integer exist as well. In this case the problem is classified as Mixed Integer Programming or MIP. Mixed Integer Linear Programming is an alternative name to the same type of problem. The majority of the mathematical programming problems in this thesis is based on MIP.

C.3 Solving LP Problems

Solving by hand for small simple problem can be done using the simplex method, which is an algebraic procedure^[69]. First, any inequality constraints are converted into equality constraints by adding variables, called slack variables. Next, this augmented form of the problem is written in a matrix-like format: the simplex tableau. Starting from a feasible initial solution, this solution is improved iteratively by performing algebraic operations on the simplex tableau, resulting in a new tableau after each iteration. The process ends when the solution cannot be improved any further.

For problems with more than just a few variables, solving it by hand would be very laborious, and this task is normally done using software. CPLEX by IBM's ILOG is a commercial mathematical programming solver^[14]. It can handle pure LP problems using several variants of the simplex method, but CPLEX can also solve (mixed) integer problems. It uses algorithms to detect network structures in LP problems and solve these more efficiently than standard LP problems. Finally, this solver can also minimise problems having a convex quadratic objective function and maximise problems that have a concave quadratic objective function.

The solver is provided in two different forms, first as an interactive executable and second as set of callable libraries. The interactive solver reads a problem from user input or from files in several formats. After solving, the solution can be presented to screen, or written to file. The callable libraries can be used by programmers to embed the solvers in applications written in C, C++, C# or Java.

As an alternative, there is also LP_SOLVE, a free, open source Linear Programming solver^[33]. This solver also handles mixed-integer problems, but has no capabilities for quadratic problems. As with CPLEX, LP_SOLVE is provided as a command line executable and as a set of callable libraries that supports even more programming languages, including high-level languages as MATLAB, Py-

thon and AMPL. On top of that, there is also a graphical Windows interface to the program.

All mathematical programming problems in this thesis have been solved by either CPLEX or LP_SOLVE. Since LP_SOLVE can also read a particular CPLEX input format, some of the problems have been constructed using this input format and have then been solved using both solvers. Apart from the possibility to compare performance, this also allows to compare both solutions, which can provide a good indication on whether the reported optimum is indeed correct.

Additional information on the environment aware arrival manager

This appendix describes the environment aware arrival manager in terms of the linear programming model that has been used. It provides more detail with respects to the transit times for the different arrival routes, the constraints that have been used to maintain the cps sequence, apply the selection of the route, maintain separation between the aircraft and to realise that aircraft are not scheduled to land earlier than feasible, based on their current location.

D.1 Transit times

The model uses the transit times from the metering fixes to the runway threshold when generating the constraints. These transit times vary for each route and aircraft type, and have been determined using the NLR ATC simulator (NARSIM). Table D.1 presents the values used.

D.2 CPS sequence

The feasible sequences according to the constraint position shifting principle have already been shown in figure 4.5. However, to explain the actual modelling in terms of MIP, a part of a more detailed version of the same decision tree is depicted in figure D.1. For this version, nodes labels have been added (A,B,...,I), together with the weight classes of the different aircraft (M,H). The names of the arcs between the nodes are not shown in the figure, but are based on the concatenation of the label of the two nodes that are connected, e.g. arc AB connects node A and B.

Fix:	Southern			Northern		
Route:	A	B	C	A	B	C
CRJ700	1000	1113	1235	1107	1113	1128
Dash 8-400	1055	1182	1314	1178	1182	1191
A320-200	1019	1158	1282	1126	1158	1175
B737-800	1021	1142	1258	1128	1142	1151
A330-300	963	1081	1208	1070	1081	1101
B767-400	1014	1136	1246	1121	1136	1139
B777-300	919	1043	1190	1026	1043	1083
B747-400	959	1106	1233	1066	1106	1126

Table D.1: *Transit times in seconds for all route, fix and aircraft type combinations*




Figure D.1: *Detailed version of the decision tree for Constrained Position Shifting*

For the MIP model, the constraint position shifting sequence is based on finding a path between the first and final node of the network. This is achieved by specifying equality constraints based on binary variables. These constraints enforce that each node has the same number of active (i.e. has value 1) incoming arcs as active outgoing arcs. For example, for nodes B and C, the following equations are added to the model:

$$\begin{aligned} \sum B : AB - BD - BE &= 0 \\ \sum C : AC - CF &= 0 \\ AB, AC, BD, BE, CF &\in \{0, 1\} \end{aligned} \quad (D.1)$$

This is repeated for all nodes in the network, except for the first and final node. At the first node, the network is initialised by forcing the solver to select one of the two initial arcs.

$$AB + AC = 1 \quad (D.2)$$

Due to the form of the constraints in combination with the initialisation, each node will either have one active incoming and one active outgoing arc or none at all. This means that together, these equations will generate a feasible path through the CPS decision tree.

D.3 Route selection

Although the landing sequence has now been set, no approach routes have been selected. The model needs to select exactly one of the three available routes (A,B,C) for each arc in the network, but only if that arc is active. This is also accomplished using equality constraints. For example for link AB:

$$\begin{aligned} AB - AB_{rA} - AB_{rB} - AB_{rC} &= 0 \\ AB_{rA}, AB_{rB}, AB_{rC} &\in \{0, 1\} \end{aligned} \quad (D.3)$$

Because each link corresponds to a particular type of aircraft, the noise exposure per link-route combination can be computed a-priori and the results can be added to the model, again using equality constraints. For example, for link AB, the noise exposure (NE) is specified using this constraint:

$$NE_{AB} - NE_{AB,rA} \cdot AB_{rA} - NE_{AB,rB} \cdot AB_{rB} - NE_{AB,rC} \cdot AB_{rC} = 0 \quad (D.4)$$

The total noise exposure of all aircraft combined is computed by summing the noise contributions of all links.

D.4 Earliest landing opportunity

Each aircraft should be allowed sufficient time to complete the flight towards one of metering fixes and complete the assigned approach route, see table D.1.

This means that there is an earliest earliest landing opportunity for each aircraft-route combination, and the scheduled landing time should be at this earliest opportunity or later. For each landing slot, this requirement is modelled using equality constraints. For example, for landing position two:

$$\begin{aligned}
<2 - 1000AB_{rA} - 1113AB_{rB} - 1235AB_{rC} - \dots \\
&1208AC_{rA} - 1222AC_{rB} - 1231AC_{rC} - D2 = 0 \\
<2, D2 \geq 0
\end{aligned} \tag{D.5}$$

where $LT2$ is the scheduled landing time for the second aircraft to land and $D2$ is any additional delay (not including delay due to an assignment of a longer route) that should be absorbed to meet the scheduled landing time¹.

D.5 Separation

All aircraft are required to maintain a predefined minimum separation with respect to their predecessor in the landing sequence. This separation, expressed in time, applies to the landing interval and is dependent on the wake vortex categories of the aircraft pair under consideration. As mentioned in section 4.2.5, the separation values that have been used are 95 seconds or more following a medium aircraft, 125 seconds or more for a heavy aircraft following a heavy one and 155 seconds or more for a medium following a heavy. For each landing slot, this requirement is modelled using inequality constraints. For example, for landing position number two and three:

$$\begin{aligned}
<2 - LT1 - 95AB - 95AC > 0 \\
<3 - LT2 - 95BD - 95BE - 155CF > 0 \\
<1, LT3 \geq 0
\end{aligned} \tag{D.6}$$

It is possible that when two consecutive aircraft would approach over the same route and the leading aircraft would be faster than, although separation has been assured at the runway threshold, separation is insufficient at the start of route. To prevent this situation, additional constraints are added to the model, but only for the pairs of aircraft where the leading aircraft would be faster and approaching over the same route. For example, for landing position number two and three:

$$LT3 - LT2 - M \cdot AB_{rA} - (M + 111) \cdot BD_{rA} > -2M \tag{D.7}$$

where M is a sufficiently large number, e.g. one thousand seconds. For this specific example, this would mean that if both aircraft would approach over route A, then the required separation (still in terms of landing interval) would be at least 111 seconds instead of 95, as previously required by the second constraints in equation D.6.

¹For this specific example, link AB involves a CRJ700 approaching over the southern fix at $t \geq 0$ and link AC means a B737-800 approaching over the northern fix at $t \geq 80$

It could be argued that more or at least adjusted separation constraints may be required. For example, time-based constraints result in a particular spatial separation, depending on the ground speed of the aircraft. In high head wind situations, when ground speeds are lower for similar air speeds, the intervals may need to be increased. Furthermore, although separation is ensured at the runway threshold, separation is not explicitly guaranteed at the merging points of the different routes or any other remaining points along the common path. However, considering the relatively short length of the common path when consecutive aircraft approach over different routes, and based on the decelerating behaviour of the aircraft on approach, it is assumed that separation assurance at the runway threshold is sufficient here.

D.6 Model verification

As described in Section 4.2.6, the solution as returned by the LP solver is post-processed for ease of interpretation. During this post-processing step, the solution can also be converted to an input file for an air traffic control simulator. With this file, the simulator can be instructed to playback the solution on a radar screen. A few of the solutions have been reviewed this way for verification purposes. As a result of this step, the constraint that provides separation between aircraft that approach over the same route was added to the final version of the model.

As a second verification activity, one of the minimum delay solutions has been checked manually to verify that: (i) the solution does not violate any of the imposed constraints, (ii) the solution is indeed the most efficient option (i.e. minimum delay), and (iii) that the reported landing times and noise exposure values are correct.

Appendix E

Additional information on strategic environmental impact allocation

This appendix describes the strategic level in terms of the linear programming model that has been developed. It provides more detail with respects to the three preparatory steps that are executed in order to generate the model equations. The first step is based on an analysis of historical weather and the annual flight schedule. Based on this analysis a number of unique situations is identified. During the second step, the feasible runway allocation options are determined for each of these situations. In the third step, the corresponding impact of each option is quantified. The final part of this appendix discusses the verification of the model.

E.1 Identification of situations

The flight schedule is analysed to identify which hours of the year show similar traffic patterns, where similar is defined here as the same number of departures, the same number of arrivals and the same period of day (i.e. day, evening or night). Each traffic situation is marked by an identifier. As an example, traffic pattern *T1* may look like:

T1
- period of day: evening
- arrivals per hour: 15
- departures per hour: 22
- occurrence: 10 hours per year
- total departures: 20x B744, 200x B738
- total arrivals: 50x B744, 100x B738

The flight schedule analysis is followed by a weather pattern analysis. The goal

of this analysis is to obtain the statistical availability of runway configurations in relation to the weather conditions. For the current implementation, this analysis is limited to the parameters wind speed and wind direction. Other weather conditions that may limit the use of particular runway configurations, like visibility and cloud base, have not been regarded at this stage.

The weather pattern analysis starts with a wind rose analysis using long-term historical weather. The outcome of this analysis is a list containing for each cell of the wind rose the historic occurrence percentage of the associated wind condition. Next, for each cell in the list it is determined which runway combinations are available, based on the prevailing cross and tail wind limitations. All cells that share the same list of available runway configurations are clustered and these clusters are called wind situations. Each wind situation is again given an identifier. Table E.1 provides an example of how a list of wind situations and corresponding available runway configurations may look like.

Pattern	Occurance	Available configurations
W1	12.25%	S1, S2, S3, S4, S5, S6
W2	11.49%	S1, S3, S5, S6
W3	9.26%	S2, S4, S6
...
W22	0.09%	S6

Table E.1: *Example list of wind situations and corresponding available runway configurations*

The process continues by combining each traffic situation with each wind situation, resulting in a list of unique combined situations (as shown in figure 5.4). This list basically represents every situation that may occur during a year, combined with the probability that it occurs. The identifiers of the wind situations and traffic situations are merged to obtain a new identifier for these combined situations. For the example, when merging traffic situation $T1$ with all wind situations, this would yield the identifiers $T1W1$ through $T1W22$.

E.2 Feasible runway allocation options

For each of these unique combined situations, a check is performed to determine which of the available runway configurations have sufficient capacity to handle the traffic. Runway configurations that, although available from a wind limit perspective, do not provide sufficient capacity are discarded. This analysis yields the main decision variables for the MIP model.

As an example, assuming that for the combination of traffic situation $T1$ and wind situation $W1$ only three viable runway combinations remain, the main decision variables could be $T1W1S1$, $T1W1S2$, and $T1W1S4$. Selection of decision variable $T1W1S1$ implies that for situation $T1W1$, the traffic would be handled

using runway configuration $S1$. If this configuration would consist of no more than one runway for arrivals and one runway for departures, all arriving traffic of this situation would be allocated to the arrival runway, and all departing traffic would be allocated to the departing runway of this configuration.

On the other hand, many of the available runway configurations consist of multiple runways for either arrivals, departures, or both. In these cases, additional decision variables are defined, in order to allow modelling the distribution of the traffic over the different runways available. In line with the example above, main decision variable $T1W1S1$ would in this case be supplemented with a set of seven variables¹, named $T1W1S1c1$ through $T1W1S1c7$. This set represents a spectrum from a maximum use of the first runway (within the capacity limits of that runway) at $T1W1S1c1$, through an even distribution at $T1W1S1c4$, to a maximum use of the second runway at $T1W1S1c7$.

As the analysis of the feasible runway allocation options is complete, the result is used to generate the majority of the constraints of the LP-problem. These constraints are primarily used to force a selection of a runway configuration for each situation. For example, if for situation $T1W1$ all feasible runway configurations consist of no more than one runway for arrivals and one runway for departures, the following equality constraint would be used:

$$T1W1S1 + T1W1S2 + T1W1S4 = 1 \quad (\text{E.1})$$

In this case, the variables are not declared integer. This means that the solver can select multiple configurations for this situation, as long as their total use is one.

If at least one of the feasible configurations offers more than a single runway for arrivals or departures, the same equality constraint is used, but the variables are declared integer:

$$\begin{aligned} T1W1S1 + T1W1S2 + T1W1S4 &= 1 \\ T1W1S1, T1W1S2, T1W1S4 &\in \{0, 1\} \end{aligned} \quad (\text{E.2})$$

This results in the choice for exactly one of the configurations. Simultaneously, additional constraints are added to manage the distribution of traffic over the runways, for all configurations. For example, now assuming that configuration $S1$ features two departure runways, the following constraint would be added:

$$T1W1S1c1 + \dots + T1W1S1c7 - T1W1S1 = 1 \quad (\text{E.3})$$

This equality equation manages that if configuration $S1$ is selected for situation $T1W1$, the variables that are used to model the traffic distribution over the two runways also become non-zero.

¹A single additional variable would suffice to model a distribution over two runways. These seven variables are used to model the non-linear runway delay function using a piece-wise linear representation. The set is declared as a Special Ordered Set of order 2. As such, at most two variables can be non-zero, and these non-zero variables must be consecutive in their ordering^[33].

E.3 Consequences of the options

The next step is to determine the consequences of each allocation option in terms of noise, risk and delay. Each allocation option corresponds to a traffic situation, for which the exact number of operations, the aircraft types and the period of day are known. At the same time, each allocation option also corresponds to the use of a particular runway configuration, and potentially also to a particular distribution of the traffic over the different runways of the configuration. Combining these data yields the specification of flight operations per runway, and these data can be used to compute the noise, risk, and delay contribution of each allocation option to the annual total. These contributions will be used as coefficients for equality constraints defining w_1 , w_2 , and w_3 in the objective function, see equation 5.3 in section 5.2.6. The noise, risk, and delay model are described in section 5.2.3.

E.4 Model verification

The model developed for the strategic level is extensive and consists of many decision variables. Therefore, it is not possible to manually inspect and verify an entire solution, as was a viable approach for the operational level model. As an alternative, a small sub-problem was defined, based on only two wind situations in combination with a few traffic situations. This small scale sub-problem was checked manually in terms of feasibility and reported values of the noise, risk, and delay contribution.

Appendix **F**

Additional information on multi-level integrated environmental management

This appendix provides additional information related to the multi-level integrated environmental management system described in chapter 6. It first presents the details and related equations of the operational allocation model. Furthermore, it specifies the traffic scenario used to generate the numerical examples and finally describes an additional test that was performed in relation to the hypothesis that the model has a preference towards a uniform solution when specifying a very short reference period.

F.1 Operational allocation model

The operational allocation model used for the multi-level integrated environmental management system is a MIP model that consists of three sets of equality constraints that manage the selection of the operational mode, the selection of a departure procedure and finally a set that is used to assist the specification of the objective function.

Operational mode selection

The model should select exactly one (eligible) operational mode for each block within the planning period. For example, for the seventh block of the planning period, the model should choose between mode C, D and E (see section F.2) and this requirement is met using the following equality constraint:

$$\begin{aligned} P7C3 + P7C4 + P7C5 &= 0 \\ P7C3, P7C4, P7C5 &\in \{0, 1\} \end{aligned} \tag{F.1}$$

where the variable $P7C3$ is used to model selection of configuration 3 (i.e. operational mode C) for block number 7 to be either true or false.

Flight procedure selection

When an operational mode has been selected for a particular block, each flight within the block should be assigned a departure procedure and route combination that is compatible with the current operational mode. This is accomplished using equality constraints like:

$$F85C3P1 + F85C3P2 + F85C3P3 + F85C3P4 + F85C3P5 - P7C3 = 0$$

$$F85C3P1, F85C3P2, \dots, F85C3P5 \in \{0, 1\}$$
(F.2)

This equation forces the selection of exactly one of the 5 available departure procedures for flight number 85, if and only if the third configuration is selected for block number seven.

Objective function variables

The objective function for this problem is shown in chapter 6, equation 6.1 and consist of four variables: the sum of the fuel burn, sum of the number of awakenings $Awak$, sum of the additional annoyance ΔA and sum of the the additional high annoyance ΔHA of all departures under consideration. The value of these four variables is computed using four additional equality constraints. For example, the total fuel burn is determined using the equation:

$$F_{F1C1P1}F1C1P1 + F_{F1C1P2}F1C1P2... + \dots + F_{F120C6P5}F120C6P5 - \sum fuel = 0$$
(F.3)

where F_{F1C1P1} is the fuel burn for flight number one, in case of selection of operational mode number one and flight procedure number one. The values of these coefficient are determined by the operational component of the multi-level system.

F.2 Traffic scenario

The traffic scenario used to demonstrate the multi-level integrated environmental management system is based on a planning period of three hours, subdivided into nine blocks of 20 minutes. In total, 120 flights are scheduled for departure during these nine blocks. Table F.1 shows the number of departures per block. This table also indicates which of the operational modes (as depicted in figure 6.2) are available for selection by the model. Modes are made unavailable if the number of departures per block per runway would exceed the maximum runway capacity, which is assumed to be twelve per block per runway. For the dual runway configuration, the distribution of the flights over the two runways is determined by

Block	N	Suitable configurations					
1	9	A	B	C	D	E	F
2	6	A	B	C	D	E	F
3	10	A	B	C	D	E	F
4	21				D	E	
5	16			C	D	E	F
6	22				D	E	
7	18			C	D	E	
8	8	A	B	C	D	E	F
9	10	A	B	C	D	E	F

Table F.1: *Number of operations per block and corresponding suitable configurations, based on capacity considerations*

	B737-300	B737-700	B737-800	B747-400	Total
ANDIK/SPYKER	13	-	-	11	24
BERGI	-	16	-	8	24
VALKO	-	8	4	-	12
LOPIK	7	-	17	-	24
ARNEM	12	-	19	5	36
Total	32	24	40	24	120

Table F.2: *Number of departures per type and directions*

the departure direction or TMA exit point of the flights considered. This explains the unavailability of mode C and F for some of the blocks, as these modes would lead to more than twelve departures on one of the runways, even though the total number of departures is below 24.

The composition of the fleet is based on the aircraft types that are available in the database of optimised departure trajectories. As shown in table F.2, 20% of the operations are executed using the Boeing 747-400. The remain 80% consist of a mix of different types of the Boeing 737. The table also shows the distribution of the aircraft types over the different departure directions.

F.3 Tests with alternative configurations

This section describes the additional model runs that have been performed to test the hypothesis that the model tends towards a uniform and concentrated solution in case of a very short reference period. These runs are based on adapting the set of eligible runway configurations, contrary to what is shown table F.1.

For the first test, all runway and route configurations are made eligible for all blocks within the planning period. This effectively means that during this test, the runway capacity constraints are ignored. The results show that the solver now applies configuration B throughout the entire planning period instead

of configuration E, as seen earlier in section 6.3.5. This first of all indicates that configuration B, and not E, is in fact the preferred configuration for this optimisation problem. However, with configuration B not available during a part of the planning period, configuration E is selected, not only during the four busy blocks (4-7), but for all blocks instead. This seems to confirm the tendency of applying a uniform and therefore concentrated solution.

A second test is based on the previous one and allows the selection of all configurations for all blocks, now except for configuration B. In this case, the solver reports an optimum based on configuration A for all nine blocks. Apparently, with configuration B not available, configuration A becomes the optimum. This means that configuration E is not even the second-best option in absence of the capacity constraints. This further strengthens the notion that configuration E is selected for the whole planning period, only because it is the optimum configuration during the busy blocks.

F.4 Model verification

The verification activities related to the multi-level model are primarily based on a careful review of the results, both in terms of decision variables, as well as objective value. Several of these checks are also described in Chapter 6. For example, a confirmation of a correct selection of profile types can be found in Section 6.3.3. As a second example, this section also looks into the details of the minimum fuel burn solution regarding the selected operational modes. Based on the results available in the database of optimised trajectories, it is confirmed that the selected operational mode is indeed the most fuel efficient one.

Bibliography

- [1] B.J.M. Ale and M. Piers. The assessment of management of third party risk around a major airport. *Journal of Hazardous Materials*, Vol. 71, No. 1-3:1–16, 2000.
- [2] G.S. Anderson and N.P. Miller. Alternative analysis of sleep-awakening data. *Noise Control Engineering Journal*, Vol. 55, No. 2:224–245, 2007.
- [3] A. Andreeva-Mori and H. Ishii. Multiple aircraft approach path optimization for noise abatement considering the influence of meteorological conditions on sound propagation. In *Proceedings 53rd AIAA Aerospace Sciences Meeting*, Reston, VA, USA, January 2015. The American Institute of Aeronautics and Astronautics (AIAA).
- [4] Anon. Procedure for the calculation of airplane noise in the vicinity of airports. Technical Report SEA-AIR-1845, Society of Automotive Engineers (SAE), January 1986.
- [5] Anon. Procedures for emission inventory preparation: Volume IV: Mobile sources. Technical Report EPA-450/4-81-026d, U.S. Environmental Protection Agency (EPA), 1992.
- [6] Anon. International standards and recommended practices, environmental protection. Annex 16 to the Convention on International Civil Aviation, Volume II, Aircraft engine emissions. Technical report, International Civil Aviation Organisation (ICAO), July 1993.
- [7] Anon. Noise abatement departure procedures. Technical Report AC 91-53A, Federal Aviation Administration (FAA), 1993.
- [8] Anon. ICAO engine exhaust emissions databank, First edition. Technical Report Doc 9646- AN/943, International Civil Aviation Organisation (ICAO), Montreal, Canada, 1995.

-
- [9] Anon. Effects of aviation noise on awakenings from sleep. Technical report, Federal Interagency Committee on Aviation Noise (FICAN), 1997.
- [10] Anon. Final report of the investigation into the probable cause of the accident with Transavia Boeing 757, registration PH-TKC, on December 24th, 1997 at Amsterdam Airport Schiphol. Technical Report 97-75/A-26, Dutch Transport Safety Board, The Hague, The Netherlands, November 1999.
- [11] Anon. Addendum to final report 97-75 / A-26. Technical report, Dutch Transport Safety Board, The Hague, The Netherlands, September 2000.
- [12] Anon. ACI airports economic survey. Technical report, Airports Council International (ACI), Geneva, Switzerland, 2001.
- [13] Anon. Directive of the European Parliament and of the Council of 25 June 2002 relating to the assessment and management of environmental noise. Technical Report 2002/49/EC, European Parliament, Council, 2002.
- [14] Anon. *ILOG CPLEX 9.0 user's manual*. ILOG S.A., Gentilly, France, October 2003.
- [15] Anon. Review of the quota count (QC) system used for administering the night noise quotas at Heathrow, Gatwick and Stansted Airports. Technical report, U.K. Department for Transport, London, UK, 2003.
- [16] Anon. ExternE - Externalities of energy - Methodology 2005 update. Technical report, European Commission, 2005.
- [17] Anon. Report on standard method of computing noise contours around civil airports. Technical Report Doc 29, 3rd edition, European Civil Aviation Conference (ECAC), 2005.
- [18] Anon. Procedures for air navigation services - aircraft operations (PANS-OPS) - volume I, flight procedures. Technical Report Doc. 8168, International Civil Aviation Organisation (ICAO), Montreal, Canada, 2006.
- [19] Anon. SOURDINE II - Deliverable D9-1: Final report. Technical report, SOURDINE Consortium, September 2006.
- [20] Anon. Procedures for air navigation services - Air traffic management (PANS-ATM). Technical Report Doc 4444, International Civil Aviation Organization (ICAO), June 2007.
- [21] Anon. SESAR definition phase - Deliverable 3: The ATM target concept. Technical report, SESAR Consortium, September 2007.
- [22] Anon. Luchtvaartnota: Concurrerende en duurzame luchtvaart voor een sterke economie. Technical report, Ministry of Transport, Public Works and Water Management, April 2009.

-
- [23] Anon. Control of development in airport public safety zones. Technical Report 01/2010, Department for Transport, March 2010.
- [24] Anon. Flying on demand - Airbus Global Market Forecast 2014-2033. Technical report, AIRBUS S.A.S., Blagnac, France, September 2014.
- [25] N. Antoine, I. Kroo, K. Willcox, and G. Barter. A framework for aircraft conceptual design and environmental performance studies. In *Proceedings 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Reston, VA, USA, August 2004. The American Institute of Aeronautics and Astronautics (AIAA).
- [26] M. Arntzen, S.J. Heblj, and D.G. Simons. Weather dependent airport noise contour prediction concept based on ray-tracing. *Journal of Aircraft*, Vol. 51, No. 5:1351–1359, 2014.
- [27] H. Balakrishnan and B. Chandran. Scheduling aircraft landings under constrained position shifting. In *Proceedings AIAA Guidance, Navigation, and Control Conference and Exhibit*, Reston, VA, USA, August 2006. The American Institute of Aeronautics and Astronautics (AIAA).
- [28] H. Balakrishnan and B. Chandran. Efficient and equitable departure scheduling in real-time: new approaches to old problems. In *Proceedings 7th USA/Europe ATM R&D Seminar*, Brussels, Belgium, July 2007. EUROCONTROL.
- [29] H. Balakrishnan and B.G. Chandran. Algorithms for scheduling runway operations under constrained position shifting. *Operations Research*, Vol. 58, No. 6:1650–1665, 2010.
- [30] D.R. Ballal and J. Zelina. Progress in aeroengine technology (1939-2003). *Journal of Aircraft*, Vol. 41, No. 1:43–50, 2004.
- [31] S.L. Baughcum, T.G. Tritz, S.C. Henderson, and D.C. Pickett. Scheduled civil aircraft emission inventories for 1992: Database development and analysis. Technical Report CR 4700, National Aeronautics and Space Administration (NASA), 1996.
- [32] B. Berglund, T. Lindvall, and D.H. Schwela. Guidelines for community noise. Technical report, World Health Organization (WHO), London, UK, 1999.
- [33] M. Berkelaar, K. Eikland, and P. Notebaert. *LP_SOLVE reference guide*, 5.5.0.15 edition, September 2009.
- [34] D.A. Black, J.A. Black, T. Issarayangyunc, and S.E. Samuels. Aircraft noise exposure and resident’s stress and hypertension: A public health perspective for airport environmental management. *Journal of Air Transport Management*, Vol. 13, No. 5:264–276, 2007.
-

-
- [35] E. Boeker and A. Rapoza. Verification of aircraft acoustic source data for the Aviation Environmental Design Tool. In *Proceedings United States National Noise Control Conference (NOISE-CON)*, Springfield, IL, USA, August 2013. Institute of Noise Control Engineering (INCE).
- [36] E.R. Boeker, E. Dinges, B. He, G. Fleming, C.J. Roof, P.J. Gerbi, A.S. Rapoza, and J. Hemann. *Integrated Noise Model version 7.0 technical manual*, January 2008.
- [37] A.M. Bolender and G.L. Slater. Evaluation of scheduling methods for multiple runways. *Journal of Aircraft*, Vol. 37, No. 3:410–416, 2000.
- [38] M. Brunet, T. Chaboud, N. Huynh, P. Malbéqui, and W. Ghedhaifi. Environmental impact evaluation of air transport systems through physical modeling and simulation. In *Proceedings 9th AIAA Aviation Technology, Integration, and Operations Conference*, Reston, VA, USA, September 2009. The American Institute of Aeronautics and Astronautics (AIAA).
- [39] R. Bütikofer. Concepts of aircraft noise calculation. *ACTA Acustica united with Acustica*, Vol. 93, No. 2:253–262, 2007.
- [40] R.E. Cadoux and J.A. Kelly. Departure Noise Limits and Monitoring Arrangements at Heathrow, Gatwick and Stansted Airports. Technical Report ERCD report 0207, Civil Aviation Authority (CAA), London, UK, March 2003.
- [41] B. Capozzi, S. Augustine, and T.R. Thompson. Dynamic noise avoidance routing through terminal airspace. In *Proceedings AIAA Guidance, Navigation and Control Conference and Exhibit*, Reston, VA, USA, August 2003. The American Institute of Aeronautics and Astronautics (AIAA).
- [42] B. Capozzi, S. Augustine, T.R. Thompson, and J.E. Robinson III. An initial assessment of benefits for noise-aware decision-support tools. In *Proceedings AIAA Aircraft Technology, Integration, and Operations Technical Forum*, Reston, VA, USA, October 2002. The American Institute of Aeronautics and Astronautics (AIAA).
- [43] J-P. Clairbois, Ph. Debroux, D. Meganck, and D. van der Poorten. Liege/Bierset airport, Brussels South/Charleroi airport: An integrated approach of the noise impact of airports. In *Proceedings 34th International Congress and Exposition on Noise Control Engineering*, Springfield, IL, USA, August 2005. Institute of Noise Control Engineering (INCE).
- [44] J-P.B. Clarke, N.T. Ho, L. Ren, J.A. Brown, K.R. Elmer, K-O. Tongand, and J.K. Wat. Continuous descent approach: Design and flight test for Louisville International Airport. *Journal of Aircraft*, Vol. 41, No. 5:1054–1066, 2004.
-

-
- [45] P.G. Cowell, P.B. Foot, R.J. Gerrard, D. Kent, S.M. Manson, and A. Rivoire. A methodology for calculation individual risk due to aircraft. Technical Report R&D Report 0007, National Air Traffic Services (NATS), 2000.
- [46] R. de Neufville and A. Odoni. *Airport systems planning, design, and management*. McGraw-Hill, Columbus, OH, USA, 2002.
- [47] K. Deb. *Multi-objective optimization using evolutionary algorithms*. Wiley, Hoboken, NJ, USA, 2001.
- [48] R. Doganis. *Flying off course: The economics of international airlines*. Routededge, New York City, NY, USA, 2002.
- [49] P. Ekins and S. Speck. Proposals of environmental fiscal reforms and the obstacles to their implementation. *Journal of Environmental Policy and Planning*, Vol. 2, No. 2:93–114, 2000.
- [50] L.J.J. Erkelens. Research into new noise abatement procedures for the 21st century. In *Proceedings AIAA Guidance, Navigation and Control Conference*, Reston, VA, USA, August 2000. The American Institute of Aeronautics and Astronautics (AIAA).
- [51] N.J. Eyre, E. Ozdemiroglu, D.W. Pearce, and P. Steele. Fuel and location effects on the damage costs of transport emissions. *Journal of Transport Economics and Policy*, Vol. 31, No. 1:5–24, 1997.
- [52] S. Fidell, D.S. Barber, and Th.J. Schultz. Updating a dosage-effect relationship for the prevalence of annoyance due to general transportation noise. *Journal of the Acoustical Society of America*, Vol. 89, No. 1:221–233, 1991.
- [53] I.H. Flindell and I.J. Witter. Non-acoustical factors in noise management at heathrow airport. *Noise & Health*, Vol. 1, No. 3:27–44, 1999.
- [54] S.P. Galis, M.A. Brouwer, and T. Joustra. Optimization of yearly airport capacity within noise limits at Schiphol Airport. In *Proceedings 33rd International Congress and Exposition on Noise Control Engineering*, Prague, Czech Republic, August 2004. Czech Acoustical Society (CAS).
- [55] P.E. Gill, W. Murray, M.A. Saunders, and M.H. Wright. *User's guide for NPSOL 5.0: A Fortran package for nonlinear programming*. Stanford University, Systems Optimization Laboratory, Stanford, CA, USA, June 2001.
- [56] R. Girvin. Aircraft noise-abatement and mitigation strategies. *Journal of Air Transport Management*, Vol. 15, No. 1:14–22, 2009.
-

-
- [57] J.M. Gulding, J.R. Olmstead, R. Bryan, L. Mirsky, G.G. Fleming, J. D'April, and P.J. Gerbi. *Integrated Noise Model users guide version 6.0*. Federal Aviation Administration (FAA), September 1999.
- [58] A. Hale. Risk contours and risk management criteria for safety at major airports, with particular reference to the case of Schiphol. *Safety Science*, Vol. 40, No. 1-4:299–323, 2002.
- [59] C.R. Hargraves and S.W. Paris. Direct trajectory optimization using non-linear programming and collocation. *Journal of Guidance, Control, and Dynamics*, Vol. 10, No. 4:338–342, 1987.
- [60] S. Hartjes, J. Dons, and H.G. Visser. Optimization of area navigation arrival routes for cumulative noise exposure. *Journal of Aircraft*, Vol. 51, No. 5:1432–1438, 2014.
- [61] S. Hartjes, H.G. Visser, and S.J. Hebli. Optimisation of RNAV noise and emission abatement standard instrument departures. *The Aeronautical Journal*, Vol. 114, No. 1162:757–767, 2010.
- [62] B. He, E. Dinges, J. Hemann, D. Rickel, L. Mirsky, C.J. Roof, E.R. Boeker, P.J. Gerbi, and D. Senzig. *Integrated Noise Model user's guide version 7.0*. Federal Aviation Administration (FAA), April 2007.
- [63] S.J. Hebli, V.A. Hanenburg, R.A.A. Wijnen, and H.G. Visser. Development of a noise allocation planning tool. In *Proceedings 45th AIAA Aerospace Sciences Meeting and Exhibit*, Reston, VA, USA, January 2007. The American Institute of Aeronautics and Astronautics (AIAA).
- [64] S.J. Hebli and H.G. Visser. Air traffic control decision support for integrated community noise management. In *Proceedings 7th AIAA Aviation Technology, Integration, and Operations Conference*, Reston, VA, USA, September 2007. The American Institute of Aeronautics and Astronautics (AIAA).
- [65] S.J. Hebli and H.G. Visser. Impact of environmentally optimized departures on runway capacity. In *Proceedings 9th AIAA Aviation Technology, Integration, and Operations Conference*, Reston, VA, USA, September 2009. The American Institute of Aeronautics and Astronautics (AIAA).
- [66] S.J. Hebli and H.G. Visser. Advanced noise abatement departure procedures: Custom-optimized departure profiles. *The Aeronautical Journal*, Vol. 119, No. 1215:647–661, May 2015.
- [67] S.J. Hebli and R.A.A. Wijnen. Development of a runway allocation optimisation model for airport strategic planning. *Transportation Planning and Technology*, Vol. 31, No. 2:201–214, 2008.
-

-
- [68] S.C. Henderson and U.K. Wickrama. *Aircraft emissions: current inventories and future scenarios*, chapter 9, pages 291–331. Cambridge university press, Cambridge, UK, 1999.
- [69] F.S. Hillier and G.J. Lieberman. *Introduction to operations research*. McGraw-Hill, Columbus, OH, USA, 2001.
- [70] R.H. Hogenhuis, S.J. Heblj, and H.G. Visser. Optimization of area navigation noise abatement approach trajectories. *Proceedings of the Institution of Mechanical Engineers - Part G - Journal of Aerospace Engineering*, Vol. 225, No. 5:513–521, May 2011.
- [71] R. Horonjeff and F.X. McKelvey. *Planning and design of airports*. McGraw-Hill, Columbus, OH, USA, 1993.
- [72] J. Huber, J-P.B. Clarke, and S. Maloney. Aircraft noise impact under diverse weather conditions. In *Proceedings 9th AIAA/CEAS Aeroacoustics Conference and Exhibit*, Reston, VA, USA, May 2003. The American Institute of Aeronautics and Astronautics (AIAA).
- [73] D.R. Isaacson, T.J. Davis, and J.E. Robinson III. Knowledge-based runway assignment for arrival aircraft in the terminal area. In *Proceedings AIAA Guidance, Navigation and Control Conference*, Reston, VA, USA, August 1997. The American Institute of Aeronautics and Astronautics (AIAA).
- [74] D.R. Isaacson and J.E. Robinson III. A Knowledge-based conflict resolution algorithm for terminal area air traffic control advisory generation. In *Proceedings AIAA Guidance, Navigation, and Control Conference and Exhibit*, Reston, VA, USA, August 2001. The American Institute of Aeronautics and Astronautics (AIAA).
- [75] Y.C. Jung and D.G. Isaacson. Development of conflict-free, unrestricted climbs for a terminal area departure tool. In *Proceedings 3rd Annual Aviation Technology, Integration and Operations Forum*, Reston, VA, USA, November 2003. The American Institute of Aeronautics and Astronautics (AIAA).
- [76] M. Kayton and W.R. Fried. *Avionics navigation systems*. Wiley, Hoboken, NJ, USA, 1997.
- [77] A.D. Kershaw, D.P. Rhodes, and N.A. Smith. The influence of ATC in approach noise abatement. In *Proceedings 3rd USA/Europe ATM R&D Seminar*, Brussels, Belgium, June 2000. EUROCONTROL.
- [78] M. Koning, E. Verkade, and J. Hakfoort. Gevolgen van uitbreiding Schiphol: Een kengetallen kosten-batenanalyse. Technical report, Netherlands Bureau for Economic Policy Analysis (CPB), The Hague, The Netherlands, 2002.
-

-
- [79] M. Krzyzanowski, B. Kuna-Dibbert, and J. Scheinder. Health effects of transport-related air pollution. Technical report, World Health Organization (WHO), 2005.
- [80] R.B.T. Lieshout, J. Veldhuis, and P. Balke. Klachtenanalyse 2005 (in Dutch). Technical Report 915, SEO Economic Research, Amsterdam, The Netherlands, 2006.
- [81] E. Maris, P.J. Stallen, R. Vermunt, and H. Steensma. Noise within the social context: Annoyance reduction through fair procedures. *Journal of the Acoustical Society of America*, Vol. 121, No. 4:2000–2010, 2007.
- [82] T.R. Meerburg, R.J. Boucherie, and M.J.A.L. van Kraaij. Noise load management at Amsterdam Airport Schiphol. Technical Report Memorandum 1859, Department of Applied Mathematics, University of Twente, Enschede, The Netherlands, December 2007.
- [83] W. Michalowski and A. Piotrowski. Solving a multi-objective production planning problem. In *Proceedings Multi-Objective Decision Making*, London, UK, April 1983. Academic Press.
- [84] H.M.E. Miedema and C.G.M. Oudshoorn. Annoyance from transportation noise: Relationships with exposure metrics. *Environmental Health Perspectives*, Vol. 109, No. 4:409–416, 2001.
- [85] G. Noel, D. Allaire, S. Jacobson, K. Willcox, and R. Cointin. Assessment of the Aviation Environmental Design Tool. In *Proceedings 8th USA/Europe ATM R&D Seminar*, Brussels, Belgium, July 2009. EUROCONTROL.
- [86] E.O. Olson and D.N. Mavris. Aircraft conceptual design and risk analysis using physics-based noise prediction. In *Proceedings 12th AIAA/CEAS Aeroacoustics Conference*, Reston, VA, USA, May 2006. The American Institute of Aeronautics and Astronautics (AIAA).
- [87] J.E Penner, D.H. Lister, D.J. Griggs, D.J. Dokken, and M. McFarland (Eds.). Aviation and the global atmosphere. Technical report, Intergovernmental Panel on Climate Change (IPCC), 1999.
- [88] S. Pietrzko and R. Bütikofer. FLULA - Swiss aircraft noise prediction program. In *Proceedings Annual Conference of the Australian Acoustical Society*, Darlinghurst, NSW, Australia, November 2002. Australian Acoustical Society (AAS).
- [89] X. Prats. *Contributions to the optimisation of aircraft noise abatement procedures*. PhD thesis, Technical University of Catalonia, Castelldefels, Spain, January 2010.
-

-
- [90] X. Prats, V. Puig, and J. Quevedo. Equitable aircraft noise-abatement departure procedures. *Journal of Guidance, Control, and Dynamics*, Vol. 34, No. 1:192–203, 2010.
- [91] H.J. Davison Reynolds, T.G. Reynolds, and R.J. Hansman. Human factors implications of continuous descent approach procedures for noise abatement. *Air Traffic Control Quarterly*, Vol. 14, No. 1:25–45, 2006.
- [92] G.J.J. Ruijgrok. *Elements of aviation acoustics*. Het Spinhuis, Amsterdam, The Netherlands, 2000.
- [93] G.J.J. Ruijgrok and D.M. van Paassen. *Elements of aircraft pollution*. VSSD, Delft, The Netherlands, 2007.
- [94] D. Rutishauser, G.L. Donohue, and R.C. Haynie. Measurements of aircraft wake vortex separation at high arrival rates and a proposed new wake vortex separation philosophy. In *Proceedings 5th USA/Europe ATM R&D Seminar*, Brussels, Belgium, June 2003. EUROCONTROL.
- [95] G. Schäfer, M. Baryn, M. Fritz, A. Johansson Augier, and U. Wieland. Europe in figures - Eurostat yearbook 2008. Technical report, Statistical office of the European Union (Eurostat), Luxembourg, 2008.
- [96] Y. Schipper. Environmental costs in European aviation. *Transport Policy*, Vol. 11, No. 2:141–154, 2004.
- [97] P.J.M. Stallen. A theoretical framework for environmental noise annoyance. *Noise & Health*, Vol. 1, No. 3:69–79, 1999.
- [98] J. E. Tate and M.C. Bell. Evaluation of a traffic demand management strategy to improve air quality in urban areas. In *Proceedings 10th International Conference on Road Transport Information and Control*, New York City, NY, USA, April 2000. Institute of Electrical and Electronics Engineers (IEEE).
- [99] R. B. Underdown and D. Cockburn. *Ground studies for pilots: Radio aids*. Wiley, Hoboken, NJ, USA, 2001.
- [100] P. Upham, C. Thomas, D. Gillingwater, and D. Rapera. Environmental capacity and airport operations: Current issues and future prospects. *Journal of Air Transport Management*, Vol. 6. No. 3:145–151, 2003.
- [101] H.M.M. van der Wal, P. Vogel, and F.J.M. Wubben. Voorschrift voor de berekening van de Lden en Lnight geluidbelasting in dB(A) ten gevolge van vliegverkeer van en naar de luchthaven Schiphol - Part 1: Berekeningsvoorschrift (in Dutch). Technical Report NLR-CR-2001-372-PT-1, National Aerospace Laboratory (NLR), Amsterdam, The Netherlands, July 2001.
-

-
- [102] R.P. van Loo. Integration of NOISHHH, an optimization tool for noise abatement, into a decision support system for strategic airport planning. Master's thesis, Delft University of Technology, Delft, The Netherlands, February 2005.
- [103] B.M.S. van Praag and B.E. Baarsma. Using happiness surveys to value intangibles: The case of airport noise. *The Economic Journal*, Vol. 115:224–246, 2005.
- [104] M.L.S. Vercammen. The design of sound insulation measures for dwellings around Amsterdam Airport Schiphol. In *Proceedings of the 33rd International Congress and Exposition on Noise Control Engineering*, Prague, Czech Republic, August 2004. Czech Acoustical Society (CAS).
- [105] H.G. Visser. Generic and site-specific criteria in the optimization of noise abatement trajectories. *Transportation Research Part D*, Vol. 10, No. 5:405–419, 2005.
- [106] H.G. Visser and R.A.A. Wijnen. Optimization of noise abatement departure trajectories. *Journal of Aircraft*, Vol. 38, No. 4:620–627, 2001.
- [107] H.G. Visser and R.A.A. Wijnen. Optimisation of noise abatement arrival trajectories. *The Aeronautical Journal*, Vol. 107, No. 1076:607–615, 2003.
- [108] F.J. Vormer, M. Mulder, M. M. van Paassen, and J.A. Mulder. Optimization of flexible approach trajectories using a genetic algorithm. *Journal of Aircraft*, Vol. 43, No. 4:941–952, 2006.
- [109] F.-J.P. Vormer. *Theoretical and operational aspects of optimal airport arrival trajectories*. PhD thesis, Delft University of Technology, Delft, The Netherlands, December 2005.
- [110] T.-C. Wang and Y.-J. Li. Optimal scheduling and speed adjustment in en route sector for arriving airplanes. *Journal of Aircraft*, Vol. 48, No.2:673–682, 2011.
- [111] R.A.A. Wijnen and H.G. Visser. Optimal departure trajectories with respect to sleep disturbance. *Aerospace Science and Technology*, Vol. 7, No. 1:81–91, 2003.
- [112] R.A.A. Wijnen, W.E. Walker, and J.H. Kwakkel. Decision support for airport strategic planning. *Transportation Planning and Technology*, Vol. 31, No. 1:11–34, 2008.
- [113] D. Wyatt and M. Tooley. *Aircraft Communications and Navigation Systems*. Taylor & Francis, London, UK, 2013.

Acknowledgements

When reflecting on my PhD time, I should distinguish between the first phase as a full time PhD-student and the time after that, when I combined my project with a full time job. During the first part, I especially enjoyed the freedom. With freedom, I mean having the opportunity to dive into interesting subjects that, sometimes, only seemed slightly related, without having to worry about budgets and deadlines. Perhaps I should have worried more about deadlines, especially the one marking the end of this first phase and the beginning of the second phase of my PhD period. I think it is fair to say that I underestimated my plan to finish this thesis during weekends and holidays. With the research about 90% complete and half of the text available in draft as well, I really thought I was almost there. I never could have imagined that it would take me years before I would be writing these final words.

I would like to thank a few people and start with Dries, my co-promotor. I still remember our meeting where this all started, a few weeks before my graduation. I wouldn't have considered a PhD if you hadn't asked me back then. You were a great support from start to end, even sometimes -or actually surprisingly often- straight from the beaches of Haifa, Lanzarote, Tenerife, and all the other warm and sunny destinations I can't remember.

I would also like to thank Ricky, my promotor. When we first met in 2007, in that hotel bar in Belfast, I think neither of us would have thought you would become my promotor. But when Michel left the university, you picked up perfectly.

The third person without whom this wouldn't have been possible is Roland. It started when you arranged the project funding at NLR and placed me in your department. But perhaps I should thank you even more for your numerous pushes towards the finish line.

I have a few more persons at NLR who I would like to acknowledge. First of all Lennard for your supervision during the first part of the project. Furthermore, thanks to Christiaan and Ed for getting me started with NARSIM, even though I used it far less than anticipated. Similar thanks to Henk, who did the noise computations for Chapter 4, saving me the need of familiarizing myself with the

Dutch noise model, at least for a while. Jürgen, thank you for your company and for showing me around at NLR. Annette, thanks to you for distracting me with other interesting projects and keeping me connected with ATEP while my desk was still upstairs. Finally a word of thanks for all the other ATEP & ATAP colleagues for the good times during breaks and social events.

Not only at NLR, but also at the university there are several other people I would like to thank. First of all Roland for your inspiration and Sander for your work on NOISHHH. Hans and Paul, I can both recommend you as roommates and I thank the other colleges of the department for your good company during our lunches and social events. I still have great memories of our pre-conference stay in that house on the banks of Lake Tahoe, a magnificent winter wonderland all around us. At the university, several MSc students also contributed to the work in this thesis, and I would like to thank Robin, Roel, Viktor, Brian and Menno.

I am also thankful to my parents, Henk and Gina. You laid the foundation of course, but also supported me when I moved to Delft. Together with my other family and friends you kept asking when my thesis would be ready, and yes, that also helped.

Finally I thank Béate, my wife, first of all for your comments and views on the parts that you read. But more importantly, for the countless times you somehow convinced me that it was really a good day for writing, regardless of my other plans. Without your dedication, my thesis would never have been completed. About time, as our next endeavour has already started!

Sander Heblj

List of Publications

P.C. Roling, S.J. Heblj, R.A.A. Wijnen, and H.G. Visser. Strategic planning for airport systems: An educational perspective. In *Proceedings 45th AIAA Aerospace Sciences Meeting and Exhibit*, Reston, VA, USA, January 2007. The American Institute of Aeronautics and Astronautics (AIAA).

S.J. Heblj, V.A. Hanenburg, R.A.A. Wijnen, and H.G. Visser. Development of a noise allocation planning tool. In *Proceedings 45th AIAA Aerospace Sciences Meeting and Exhibit*, Reston, VA, USA, January 2007. The American Institute of Aeronautics and Astronautics (AIAA).

S.J. Heblj and H.G. Visser. Air traffic control decision support for integrated community noise management. In *Proceedings 7th AIAA Aviation Technology, Integration, and Operations Conference*, Reston, VA, USA, September 2007. The American Institute of Aeronautics and Astronautics (AIAA).

S.J. Heblj and R.A.A. Wijnen. Development of a runway allocation optimisation model for airport strategic planning. *Transportation Planning and Technology*, Vol. 31, No. 2:201-214, 2008

H.G. Visser, S.J. Heblj, and R.A.A. Wijnen, Improving the Management of the Environmental Impact of Airport Operations. In: Filip N. Gustavsson (ed.), *New Transportation Research Progress*, pp.1-65, Hauppauge, NY, USA, 2008. Nova Science Publishers.

S.J. Heblj and H.G. Visser. Impact of environmentally optimized departures on runway capacity. In *Proceedings 9th AIAA Aviation Technology, Integration, and Operations Conference*, Reston, VA, USA, September 2009. The American Institute of Aeronautics and Astronautics (AIAA).

S. Hartjes, H.G. Visser, and S.J. Heblj. Optimisation of RNAV noise and emission abatement standard instrument departures. *The Aeronautical Journal*, Vol. 114, No. 1162:757-767, 2010.

R.H Hogenhuis, S.J. Heblj, and H.G Visser. Optimization of area navigation noise

abatement approach trajectories. *Proceedings of the Institution of Mechanical Engineers - Part G - Journal of Aerospace Engineering*, Vol. 225, No. 5:513-521, May 2011

M.L. Braakenburg, S. Hartjes, H.G. Visser, and S.J. Heblj. Development of a multi-event trajectory optimization tool for noise-optimized approach route design. In *Proceedings 11th AIAA Aviation Technology, Integration, and Operations Conference*, Reston, VA, USA, September 2011. The American Institute of Aeronautics and Astronautics (AIAA).

S.J. Heblj, M. Arntzen, and D.G. Simons. A weather dependent noise contour prediction concept: Combining a standard method with ray tracing. In *Proceedings 9th European Conference on Noise Control*, Madrid, Spain, June 2012. The European Acoustics Association (EAA).

B.R. Kuiper, H.G. Visser, and S.J. Heblj. Efficient use of an allotted airport annual noise budget through minimax optimization of runway allocations. *Proceedings of the Institution of Mechanical Engineers - Part G - Journal of Aerospace Engineering*. Vol. 227, No. 6:1021-1035, June 2012

M. Arntzen, S.J. Heblj, and D.G. Simons. A weather dependent noise contour prediction concept: Calculating multi-event noise contours with ray-tracing. In *Proceedings 12th AIAA Aviation Technology, Integration, and Operations Conference*, Reston, VA, USA, September 2012. The American Institute of Aeronautics and Astronautics (AIAA).

S.J. Heblj, V. Sindhamani, M. Arntzen, D.H.T. Bergmans, and D.G. Simons. Noise attenuation directly under the flight path in varying atmospheric conditions. In *Proceedings 42nd International Congress and Exposition on Noise Control Engineering*. Vienna, Austria, September 2013. Austrian Noise Abatement Association (ÖAL).

M. Arntzen, S.J. Heblj and D.G. Simons. Weather dependent airport noise contour prediction concept based on ray-tracing. *Journal of Aircraft*, Vol. 51, No. 5:1351-1359, September 2014.

S.J. Heblj and H.G. Visser. Advanced noise abatement departure procedures: Custom-optimized departure profiles. *The Aeronautical Journal*, Vol. 119, No. 1215:647-661, May 2015.

About the author

Sander Heblj was born on 24 May 1980 in Apeldoorn, the Netherlands. After completing his primary education in 1992, he started on the Atheneum programme at De Heemgaard, a comprehensive school in Apeldoorn. There he obtained his VWO diploma in June 1998.

In September of the same year, he enrolled as Aerospace Engineering student at Delft University of Technology. During his studies in Delft he accepted a one year position as board treasurer for EUROAVIA Delft, the local branch of the European Association of Aerospace Students. Furthermore, he also did a three months internship at the Netherlands Aerospace Centre (NLR) in Amsterdam. In 2003 Sander obtained his B.Sc. degree, followed by his M.Sc. degree in 2004. His final thesis was on a strategic noise allocation model.

Directly after his graduation, he started as associate researcher at the Flight Mechanics and Propulsion Group of Delft University of Technology, Faculty of Aerospace Engineering. He worked on several research projects related to aircraft noise, and Schiphol Airport in particular. About one and a half years later, he started as PhD student and worked on his research for four years.

As of 2010, Sander works as an R&D Engineer for the Environment and Policy Support department at the Netherlands Aerospace Centre. In this position, he works primarily on noise modelling, including noise model development and verification studies, noise optimisation studies and noise assessment studies.

