Enhancing Decision Support Systems for Airport Ground Movement

Stefan Ravizza
MSc ETH Math. with distinction

School of Computer Science

University of Nottingham

Thesis submitted to the University of Nottingham for the degree of

Doctor of Philosophy (PhD)

July 2013

Abstract

With the expected continued increases in air transportation, the mitigation of the consequent delays and environmental effects is becoming more and more important, requiring increasingly sophisticated approaches for airside airport operations. The ground movement problem forms the link between other airside problems at an airport, such as arrival sequencing, departure sequencing, gate/stand allocation and stand holding. The purpose of this thesis is to contribute to airport ground movement research through obtaining a better understanding of the problem and producing new models and algorithms for three sub-problems. Firstly, many stakeholders at an airport can benefit from more accurate taxi time predictions. This thesis focuses upon this aim by analysing the important factors affecting taxi times for arrivals and departures and by comparing different regression models to analyse which one performs the best for this particular task. It was found that incorporating the information of the airport layout could significantly improve the accuracy and that a TSK fuzzy rule-based system outperformed other approaches. Secondly, a fast and flexible decision support system is introduced which can help ground controllers in an airport tower to make better routing and scheduling decisions and can also absorb as much of the waiting time as possible for departures at the gate/stand, to reduce the fuel burn and environmental impact. The results show potential maximum savings in total taxi time of about 30.3%, compared to the actual performance at the airport. Thirdly, a new research direction is explored which analyses the trade-off between taxi time and fuel consumption during taxiing. A sophisticated new model is presented to make such an analysis possible. Furthermore, this research provides the basis for integrating the ground movement problem with other airport operations. Datasets from Zurich Airport, Stockholm-Arlanda Airport, London Heathrow Airport and Hartsfield-Jackson Atlanta International Airport were utilised to test these sub-problems.

Dedicated to Nicoletta

Acknowledgements

First of all, I am deeply grateful to Dr. Jason Atkin for his ongoing support as my main supervisor. He has always been accessible and supported me with fruitful discussions and inputs in every project. In addition, I would like to thank my external supervisor Prof. Dr. Edmund Burke (University of Stirling) for his strategic guidance, leading the project where my PhD has been a part of it and securing the funding from EPSRC (EP/H004424/2) for my time in Nottingham. I also want to thank Dr. Andrew Parkes for agreeing to be my second internal supervisor.

There are many other academics and colleges I am grateful to. Prof. Dr. Marloes Maathuis from the ETH Zurich, Switzerland has added valuable expertise in the field of statistics to my PhD and has always been a pleasure to work with. Prof. Dr. John-Paul Clarke and his team at the Georgia Institute of Technology welcomed me for a scholar visit in Atlanta, USA and offered me interesting insights into research and non-research related differences between North America and Europe. Furthermore, I would like to thank the people from my project and in particular Dr. Jun Chen (University of Lincoln) with whom I have co-authored several papers. Equally importantly, I would like to thank the ASAP and ICOS research groups, and especially their PhD students, who made the time in England pleasant.

Moreover, my thanks go to professionals who gave me insights into the research topic and data to run experiments. My biggest thank you is directed to the team at Zurich Airport and especially to Giovanni Russo and Daniele Gullo. Most of the analysis within this thesis is based on their data and I valued their professional attitude during the entire time. I also thank Tim Walmsley for showing us Manchester Airport and Pelle Lövstrand and Simon Brown for providing datasets for Stockholm-Arlanda Airport and Heathrow Airport, respectively.

Last but not least, I am deeply grateful to Nicoletta and to my family for their unending support and patience through the entirety of my PhD and the time of my higher education in general.

Contents

List of Figures

List of Tables

1

Introduction

> Imagination is more important than knowledge.
>
> Albert Einstein

1.1 Background and Motivation

There has been a significant increase in air traffic over the past few years and this trend is predicted to continue. The SESAR (Single European Sky ATM Research) project predicts a doubling in the number of flights between 2005 and 2020 (SESAR 2006). SESAR aims to triple capacity by 2020 and to reduce delays on the ground and in the air (SESAR 2009). It is expected that the hub airports could form the bottlenecks for the overall air traffic management system within Europe (SESAR 2009; ACI EUROPE 2010). Hence, improvements in critical airport operations will be more and more important in the near future. The main operations which affect this bottleneck are arrival and departure management (sequencing and scheduling) at the runway, gate assignment, and ground movement.

Ground movement links the various other operations together and is the focus of this thesis. The main aim of this research is to better understand the ground movement problem, in such a way that important operations can be automated. Furthermore, this research provides the

basis for integrating the ground movement problem with other airport operations. The majority of previous research has focussed on the optimisation of a single airport operation at a time. Moreover, it has often either neglected real-world constraints or provided decision support systems which would take too long to find a beneficial solution to be used as real-time methods. The execution time of an approach can be especially important when solving integrated problems. From both an economic point of view (reducing delays and increasing throughput), and an environmental point of view (reducing noise, air pollution and carbon emissions), there are obvious benefits to be gained from treating the different airport operations as a whole.

This work is part of a research project called "Integrating and Automating Airport Operations" in which four universities are involved, namely the University of Lincoln, the University of Loughborough, the University of Nottingham and the University of Stirling. In addition, Manchester Airport and Zurich Airport have committed time and resource to it and their involvement has provided us with access to historic data from airports and to gain insights into their operations. During this research, I was able to start collaborations with three other airports to extend the analysis of some approaches to more airports.

1.2 Aims and Scopes

The purpose of this thesis is to contribute to airport ground movement research through obtaining a better understanding of the problem and producing new models and algorithms for various sub-problems. This thesis can broadly be split into three different research directions:

- Firstly, many stakeholders at an airport can benefit from more accurate taxi time predictions. Two chapters focus upon this aim by analysing the important factors affecting taxi times and by comparing different regression models to analyse which one performs best for this particular task.
- Secondly, a fast and flexible decision support system is introduced which can help ground controllers in an airport tower to make better routing and scheduling decisions.
- Lastly, a new research direction is explored which analyses the trade-off between minimis-

ing either taxi time or fuel consumption during the taxiing process. A sophisticated new model is presented to make such an analysis possible.

The focus is upon the ground movement problem, but some of the links to other airside airport operations are also discussed, to enable follow-up research related to the integration of various problems. In contrast to many other sources, this research analyses European hub airports, whereas much of the academic research in the literature is focused on North American hub airports, where the situation and the bottlenecks often differ.

1.3 Contributions of this Thesis

The contributions can be summarised as follows:

- Chapter 2
  - Overview and comparison of the various ground movement models and solution methods in the literature.
- Chapter 4
  - A linear function was developed to more accurately estimate taxi times at European hub airports.
  - Identification of the factors which are better correlated to taxi speeds, such as taxi distance and the traffic level.
  - First extensive analysis of taxi time prediction that covers both arrivals and departures at an airport.
  - A function which is useable by ground movement optimisers, allowing the elimination of individual factors.
- Chapter 5
  - Analysis compares different regression approaches, to determine the best approach for predicting taxi times of aircraft.

  - A TSK fuzzy rule-based system was found to outperform the other approaches.
  - New insights into the influence of distance and amount of traffic upon taxi-in times.

- Chapter 6
  - Introduction of a more realistic ground movement decision support system with average solution times of only a few milliseconds per aircraft, making it adequate for real-time use.
  - This work extends the basic ground movement problem of minimising the travel times to include the concept of absorbing possible waiting times for departures at the gate/stand.
  - The results show potential maximum savings in total taxi time of about 30.3%, compared to the actual performance at the airport.
  - An effective swap-operator which can further improve the quality of the solution with comparatively little additional computational time.
  - An approach which could be used to test different airport extension plans and to analyse the effects of different operational scenarios.

- Chapter 7
  - A sophisticated combination of two algorithms has enabled the development of a framework to analyse the trade-off between taxi time and fuel consumption for the ground movement problem.
  - The new model is able to tackle this hard problem in a comparatively efficient way.
  - Sensitivity analysis has highlighted that the potential trade-off depends very much upon the actual modelling of the fuel-based objective function.

1.4 Publications from this Thesis

A number of publications have been produced in the course of the development of this thesis underpinning the importance of the tackled problems and the achievements of our work. The key publications can be assigned to specific chapters as follows:

- Chapter 2
  - Jason A. D. Atkin, Edmund K. Burke, and Stefan Ravizza. "The airport ground movement problem: Past and current research and future directions". *Proceedings of the 4th International Conference on Research in Air Transportation (ICRAT)*, Budapest, Hungary, pp 131-138, 2010.
- Chapter 4
  - Jason A. D. Atkin, Edmund K. Burke, and Stefan Ravizza. "A statistical approach for taxi time estimation at London Heathrow Airport". *Proceedings of the 10th Workshop on Models and Algorithms for Planning and Scheduling Problems (MAPSP)*, Nymburk, Czech Republic, 2011.
  - Stefan Ravizza, Jason A. D. Atkin, Marloes H . Maathuis, and Edmund K. Burke. "A statistical approach for improving taxi time estimations at airports". *Journal of the Operational Research Society*, (advance online publication), doi:10.1057/jors.2012.123.
- Chapter 5
  - Jun Chen, Stefan Ravizza, Jason A. D. Atkin, and Paul Stewart. "On the utilisation of fuzzy rule-based systems for taxi time estimations at airports". *Proceedings of the 11th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS)*, Saarbrücken, Germany, pp 134-145, 2011.
  - Stefan Ravizza, Jun Chen, Jason A. D. Atkin, Paul Stewart, and Edmund K. Burke. "Aircraft taxi time prediction: Comparisons and insights". (submitted)
- Chapter 6
  - Jason A. D. Atkin, Edmund K. Burke, and Stefan Ravizza. "A more realistic approach for airport ground movement optimisation with stand holding". *Proceedings of the 5th Multidisciplinary International Scheduling Conference (MISTA)*, Phoenix, Arizona, USA, 2011.

  - Stefan Ravizza and Jason A. D. Atkin. "Exploration of the ordering for a sequential airport ground movement algorithm". Tech. Rep. 1543, University of Nottingham, 2011.
  - Stefan Ravizza, Jason A. D. Atkin, and Edmund K. Burke. "A more realistic approach for airport ground movement optimisation with stand holding". *Journal of Scheduling*, (advance online publication), doi:10.1007/s10951-013-0323-3.

- Chapter 7
  - Stefan Ravizza, Jun Chen, Jason A. D. Atkin, Edmund K. Burke, and Paul Stewart. "Trade-off analysis between taxi time and fuel consumption in airport ground movement". *Proceedings of the Conference on Advanced Systems for Public Transport (CASPT12)*, Santiago, Chile, 2012.
  - Stefan Ravizza, Jun Chen, Jason A. D. Atkin, Edmund K. Burke, and Paul Stewart. "Trade-off analysis between taxi time and fuel consumption in airport ground movement". *Public Transport*, (advance online publication), doi:10.1007/s12469-013-0060-1.

The order of the authors changed over time, from alphabetic order to order by contribution. I was main author on all publications, except the conference paper which was published at the ATMOS workshop in 2011. In addition to the aforementioned papers, I presented my work at the following events:

- "The airport ground movement problem: Past and current research and future directions", presentation at Student Conference on Operational Research (SCOR 2010), Nottingham, UK (2010-04-10)
- "The airport ground movement problem: Past and current research and future directions", presentation at International Conference on Research in Air Transportation (ICRAT 2010), Budapest, Hungary (2010-06-03)
- "Comparison of airport ground movement algorithms", presentation at EURO XXIV, Lisbon, Portugal (2010-07-12)

- "Ground movement at airports", presentation at "Airports and the Environment" seminar (PhD for PhD event), Loughborough, UK (2010-09-22)
- "Ground movement at airports", presentation at ASAP seminar, Nottingham, UK (2010-12-07)
- "Ground movement optimisation to increase stand holding", presentation at Aerodays 2011, Madrid, Spain (2011-03-30)
- "A statistical approach for taxi time estimation at London Heathrow Airport", presentation at Workshop on Models and Algorithms for Planning and Scheduling Problems (MAPSP 2011), Nymburk, Czech Republic (2011-06-20)
- "A more realistic approach for airport ground movement optimisation with stand holding", presentation at Multidisciplinary International Scheduling Conference (MISTA 2011), Phoenix, Arizona, USA (2011-08-10)
- "Exploration of the ordering for a sequential airport ground movement algorithm", presentation at OR 2011, Zurich, Switzerland (2011-09-02)
- "How OR can aid airport ground movement", presentation at OR 53, Nottingham, UK (2011-09-06)
- "Integrating and automating airport operations" (jointly with Geert De Maere), presentation at Air Transport Research Workshop, Lincoln, UK (2011-09-09)
- "How OR can aid airport ground movement", presentation at ASAP seminar, Nottingham, UK (2011-12-15)
- "How Operational Research can aid airport ground movement", presentation at University of Stirling, Stirling, UK (2012-03-06)
- "The trade-off between taxi time and fuel consumption in airport ground movement", presentation at Conference on Advanced Systems for Public Transport (CASPT 2012), Santiago, Chile (2012-07-23)

Finally, I presented two posters at two different competitions:

- "Ground movement optimisation to increase stand holding", 6th European Aeronautics Days (Aerodays 2011), Madrid, Spain, 2011 (shortlisted)
- "Less hassle for your next flight - A green approach to optimise airport operations", Research Showcase 2011 from the University of Nottingham, Nottingham, UK, 2011 (shortlisted)

1.5 Non-disclosure Agreement

Three non-disclosure agreements cover this thesis and prevent us from making the utilised datasets available. This work, as well as all related publications, were approved by our contacts at the particular airports, such as Giovanni Russo and Daniele Gullo at Zurich Airport, Pelle Lövstrand at Stockholm-Arlanda Airport and Simon Brown at Heathrow Airport.

1.6 Collaborations with Other Universities

Prof. Dr. Marloes Maathuis from the ETH Zurich, Switzerland has helped in the work related to Chapter 4. In particular, she was involved in the validation of the statistical assumptions (Section 4.3.4.3) and the analysis with the general least squares models using autoregression models (Appendix A).

Dr. Jun Chen from the University of Lincoln, UK has contributed to the work in Chapters 5 and 7. Both fuzzy rule-based systems used in Chapter 5 have originally been implemented by him and an adopted version of the source code has been made available to us (Sections 5.3.5 and 5.3.6). Chapter 7 combines his research with the algorithm from Chapter 6. The part about the population adaptive immune algorithm (Section 7.3.4) has been developed by him and he made it available to run the combined model.

Prof. Dr. John-Paul Clarke from the Georgia Institute of Technology, USA has welcomed me as a scholar visitor for two months. He made it possible to work with data from Hartsfield-Jackson Atlanta International Airport to analyse a North American airport (Appendix D).

1.7 Structure of the Thesis

The main body of this thesis is organised as follows: Chapter 2 reviews existing contributions related to ground movement at airports, the integration with other airport operations and related research areas. In addition, the needs in this research areas are highlighted in this chapter. In Chapter 3, the different datasets are discussed, which are then utilised in the following chapters. A multiple linear regression approach is utilised in Chapter 4 to identify the significant factors when predicting taxi times for arrivals and departures. Chapter 5 continues with the prediction of taxi times, but instead of focusing on the relevant factors, different regression approaches are tested to find more accurate predictions. Chapter 6 introduces a graph-based algorithm to solve the routing and scheduling problem for aircraft on the surface of an airport which takes into account the findings of Chapters 4 and 5. In Chapter 7, an extension of the graph-based algorithm is presented where in addition to minimising the total taxi time a new objective was added, to taxi in a more environmentally friendly manner. Finally, a general concluding discussion is presented in Chapter 8.

2

Background and Related Work

> I suppose it is tempting, if the only tool you have is a hammer, to treat everything as if it were a nail.
>
> *The Psychology of Science: A Reconnaissance*
> Abraham Maslow

2.1 The Airport Ground Movement Problem

The airport ground movement problem is basically a routing and scheduling problem. It involves directing aircraft on the ground to their destinations in a timely manner, with the aim being to either reduce the overall travel time and/or to meet some target time windows, while simultaneously reducing environmental issues. Throughout the movement, it is crucial, for reasons of safety, that two aircraft never conflict with each other. The complexity of the problem can vary and should drive the choice of the solution approach. When an airport has only a few aircraft moving at once, with few potential conflicts between them, optimal routing can be achieved by simply applying a shortest path algorithm, such as Dijkstra's algorithm (Dijkstra 1959; Cormen et al. 2001), to each aircraft in turn. For larger airports, especially during peak hours, the interaction between the routes of different aircraft often requires the application of

a more complex routing algorithm.

The details of the problem descriptions and the constraints which have been utilised in previous work have varied according to the requirements of the airport which was being modelled. The various constraints upon the ground movement problem are considered in Section 2.1.1 and the different objectives in Section 2.1.2. Since it is important for improving the operations at an airport to integrate the related operations with the ground movement problem, this integration is discussed in Section 2.2. Section 2.3 presents the existing models and solution approaches of the ground movement problem and Section 2.4 relates this area to other research areas. Sections 2.5 and 2.6 review the areas of taxi time prediction and fuel efficient taxiing at airports, before Section 2.7 highlights some of the further needs related to ground movement at airports from an operational research point of view.

2.1.1 Constraints

The different constraints discussed in the ground movement research literature can be divided into the following categories:

2.1.1.1 Consideration of the route taken

It is important to ensure that aircraft follow a permitted route (see Figure 2.1 for an example). If the route for each aircraft is pre-determined, the ground movement problem is reduced to finding the best possible schedule (Smeltink et al. 2004; Rathinam et al. 2008). The other extreme occurs when no restrictions are set for the routing of each aircraft (Marín 2006; Marín and Codina 2008; Keith and Richards 2008; Clare et al. 2009; Clare and Richards 2009, 2011). The last possibility is for the restrictions to lie somewhere in between these extremes, where there is a predefined set of routes for each aircraft and the algorithm can choose amongst them (Pesic et al. 2001; Gotteland et al. 2001, 2003; Gotteland and Durand 2003; Herrero et al. 2005; García et al. 2005; Balakrishnan and Jung 2007; Roling and Visser 2008; Deau et al. 2008, 2009). How an airport is operated can differ and hence certain approaches can be more suitable for certain airports. The mentioned sources are further discussed in Section 2.3.

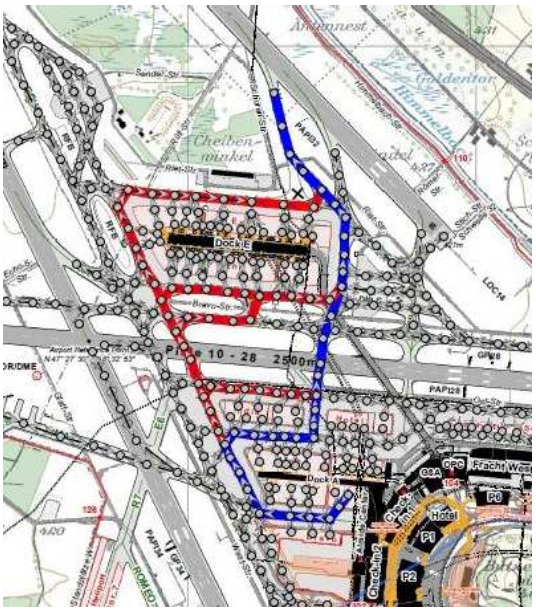

Figure 2.1: Different routes from the exit of runway 14 to pier A at Zurich Airport

2.1.1.2 Separation constraints between aircraft

As previously mentioned, it is crucial that aircraft moving around on the ground do not conflict with each other and have a separation based on jet blast to avoid affecting aircraft behind them. This is ensured during taxiing by applying separation constraints. The required minimum distances between aircraft appear to vary between authors. For example, Pesic et al. (2001) required it to be at least 60 metres, while Smeltink et al. (2004) required a value of 200 metres. Such constraints can also depend upon the aircraft type or size. If an aircraft is at a gate, no such restriction is usually used. At the point of take-off or landing, other restrictions are employed, which are presented in Section 2.2.

2.1.1.3 Aircraft movement speeds

Different aircraft require different lengths of time for taxiing. Recent research has taken this into account, modelling the speed depending either upon the type or size of an aircraft (Balakrishnan and Jung 2007; Roling and Visser 2008), or the kind of taxiway that is being followed (Gotteland et al. 2001). The time for making a turn can also be taken into account (Pesic et al. 2001),

where the maximal speed may need to be limited. Chapters 4 and 5 analyse how taxi speed can vary and this information is incorporated in the routing and scheduling approach in Chapter 6.

2.1.1.4 Timing constraints for arrivals

Arriving aircraft have to be routed from the runway to their stands (see Figure 2.1 for an example). From the point of view of the isolated ground movement problem, the arrival time for aircraft can be considered to either be fixed or to permit small deviations. The allocated gate is usually assumed to be vacant and, therefore, the aim is usually for the aircraft to reach the gate as soon as possible, since this is better from an environmental as well as an airline and passenger perspective.

2.1.1.5 Timing constraints for departures

Departing aircraft have to be routed and scheduled from their stands to the runway from which they will be departing. A pushback time (or earliest pushback time) is usually provided and is often seen as an earliest time for an aircraft to start taxiing. The aims for the ground movement of the departing aircraft can be more complicated than for arrivals. Assuming that the departure sequencing at the runway(s) has not been integrated into the problem, one of the following aims is usually adopted: 1) To reach the runway as early as possible. 2) To reach the runway in time to attain, or be as close as possible to, a pre-determined take-off time. 3) To reach the runway in time to take off within a specified time window, since many European aircraft have fifteen minute slots which are allocated by the Eurocontrol Central Flow Management Unit (CFMU) and have to be satisfied (Gotteland et al. 2003).

2.1.2 Objective Functions

The aim of the ground movement problem depends upon the scope of the problem. Much of the previous research has concentrated upon minimising the total taxi time including the waiting time for aircraft at the runway (Pesic et al. 2001; Marín 2006; Roling and Visser 2008; Rathinam et al. 2008), while other research has considered makespan (the duration from first

to last movement) minimisation (Herrero et al. 2005; García et al. 2005). Yet more research has treated this as a multi-objective problem. For example, penalising deviations from a scheduled time of departure/arrival (STD/STA) (Smeltink et al. 2004; Balakrishnan and Jung 2007; Deau et al. 2008, 2009), or from the CFMU slots (Gotteland et al. 2003), in addition to considering one of the total taxi time or makespan reduction objectives. In other research, longer taxi paths were penalised as well (Gotteland et al. 2001; Keith and Richards 2008; Clare et al. 2009; Clare and Richards 2009, 2011). Marín and Codina (2008) used a weighted linear objective function to simultaneously consider the total routing time, number of controller interventions, worst routing time, delays for arriving and departing aircraft and the number of arrivals and take-offs.

2.2 Integration of Other Airport Operations

The ground movement problem does not actually occur in isolation at an airport. The landing/arrival sequence will determine the times at which some aircraft enter the system; the gate/stand allocation problem will determine where arrivals leave the system and where departures enter the system. The departure sequencing problem determines the times at which departures leave the system. These systems can be seen to be intimately linked, so potential benefits from integrating all four problems are obvious. However, little research so far has considered this integration (see Sections 2.2.1 and 2.2.2). The complexity of these problems is such that it is currently impossible to simultaneously optimise all of these airport operations, but the real situation at the airport means that there has to be at least some coordination between the solutions of the sub-problems.

Before showing the relevant research areas which are linked to the ground movement problem, the interested reader is directed to review papers about air traffic management in general to better understand all of the relationships between the different areas (Wu and Caves 2002; Barnhart et al. 2003; Ball et al. 2006).

2.2.1 Integration of Departure Sequences

For departing aircraft, the ground movement can affect the departure sequencing, and vice versa. An optimal take-off sequence is of no use if it cannot be achieved by the taxiing aircraft, as discussed in Atkin et al. (2007). To maximise the throughput of a runway, two sequence-dependent separations are of major importance (Atkin 2008): wake vortex separations and en-route separations. The wake vortex separations depend upon the weight classes of the aircraft, so that larger separations are required whenever a lighter class of aircraft follows a heavier class. Separations also have to be increased when aircraft have similar departure routes (to ensure that en-route separations are met) or when the following aircraft is faster (to allow for convergence in the air).

Departure sequencing is sometimes considered within ground movement research (Gotteland et al. 2001), especially the newer research (Rathinam et al. 2008; Deau et al. 2008, 2009; Keith and Richards 2008; Clare et al. 2009; Clare and Richards 2009, 2011), in order to ensure that aircraft arrive at the departure runway at appropriate times, rather than merely reducing the overall taxi times. Only wake vortex separations are usually considered. However, the en-route separations are also sometimes taken into account (Keith and Richards 2008; Clare et al. 2009; Clare and Richards 2009, 2011).

Similarly, taxi times cannot be ignored in realistic departure sequencing systems. The movement near the runway can be especially important, for example, within flexible holding areas (Leese et al. 2001; Atkin et al. 2007), or the interleaving of runway queues (Bolender 2000). Even where the models for movement are not explicitly required, accurate taxi time predictions are often beneficial for improving sequencing (Atkin et al. 2006, 2008b), even when re-sequencing is performed at the runway, and would be more important if the re-sequencing was performed earlier.

Different techniques have been used to tackle the departure sequencing problem: Leese et al. (2001) and Balakrishnan and Chandran (2010) used dynamic programming algorithms, Cooper et al. (2002) and Gupta et al. (2010a) employed mixed integer programming formulations, Anagnostakis and Clarke (2003) recommended a two-stage approach and Atkin et al. presented work using metaheuristics (Atkin et al. 2007; Atkin 2008; Atkin et al. 2008b,a, 2009). A recent

review paper was written by Bennell et al. (2011). In addition, Li et al. (2009) presented an approach to plan runway configuration changes under stochastic wind conditions as the preferred direction for aircraft to depart is depended upon the wind.

The ultimate goal to support airside operations at an airport is to integrate ground movement with other operations. However, to be useful, such an approach needs to be able to find solutions in real-time, to deal with changes of the situation at an airport (such as delays, changes of gate allocations, etc.) and to be able to model the different operations in a realistic way by incorporating all the required real-world constraints.

2.2.2 Integration of Arrival Sequences

Aircraft enter the ground movement system by landing on a runway, or by leaving stands. The entry times into the system of landing aircraft will influence the ground movement operations. Better arrival time predictions can have a positive effect on the ground movement planning. There may be a choice of landing runway to be made. This choice can depend upon the current status of the ground movement and the assigned gate for the aircraft. After landing it will influence the later ground movement planning. Boysen and Fliedner (2011) analysed the problem to balance workload of ground staff by evenly distributing the number of arriving passengers, the arrivals per airline and the number of arrived passengers per airline over a certain time.

In some airport layouts, runway crossings may be necessary for taxiing aircraft. For realistic runway sequencing and taxiing optimisation, such crossings may need to be taken into account (Anagnostakis and Clarke 2003), requiring knowledge of the runway sequencing when planning the ground movement. Furthermore, runways are sometimes used in mixed mode, in which case departure and arrival sequences also have to be coordinated (Bianco et al. 2006; Böhme et al. 2007). Some of the approaches for the departure problem can also be used to solve the arrival problem and Bennell et al. (2011) presented a recent survey about both kinds of problem.

The problem of scheduling arrivals on a runway can be seen as a machine scheduling problem with sequence-dependent setup times (Bianco et al. 1999; Ernst et al. 1999; Bennell et al. 2011). Most of the literature either solves the problem with dynamic programming algorithms

(Psaraftis 1980; Chandran and Balakrishnan 2007; Balakrishnan and Chandran 2010) or with heuristics (Ernst et al. 1999; Bianco et al. 1999, 2006; Beasley et al. 2004; Soomer and Koole 2008; Soomer and Franx 2008; Salehipour et al. 2013). Artiouchine et al. (2008) presented a compact mixed integer programming formulation which was solved with a hybrid branch and cut framework. Very recently, Tavakkoli-Moghaddam et al. (2012) used fuzzy goal-programming to solve the problem on a single runway.

2.2.3 Integration of Gate Assignment

Gate assignment is another major problem which arises at congested airports (Ding et al. 2005; Dorndorf et al. 2008; Drexl and Nikulin 2008; Diepen et al. 2009; Jaehn 2010; Dorndorf et al. 2012). The aim is to find an assignment of aircraft to gates at terminals, or stands on the apron, so that some measure of quality (such as total passenger walking distance) is improved. This problem was fully discussed in a survey paper by Dorndorf et al. (2007), where the need for future work in multi-objective optimisation and robust assignments was also identified.

The ground movement problem could be integrated with the gate assignment problem, with the aim being to allocate gates/stands so that the total taxiing distance is reduced. This would have a beneficial impact upon the use of fuel, with consequent benefits for the environment as well as financial savings for airlines, delay benefits for passengers and a reduction in congestion on the apron. Kim et al. (2009, 2010) presented gate assignment research which considered minimising passenger flow in terminals and aircraft congestion on ramps. Other research areas for the gate assignment problems are how to reassign gates for flight delays (Tang et al. 2009; Maharjan and Matis 2011) or to have robust or stochastic gate assignments (Kim and Feron 2011; Azeker and Noyan 2012; Diepen et al. 2012).

A related field of research is to schedule baggage-handling facilities at airports (Abdelghany et al. 2006; Asco et al. 2011; Barth and Pisinger 2012) which is another assignment problem with different objective functions.

2.3 Existing Ground Movement Models and Solution Approaches

Table 2.1: Overview of approaches for the ground movement problem based on MILP formulations and GA models

Authors	Approach	Representation
Pesic et al. (2001)	GA	Times
Gotteland et al. (2001); Gotteland and Durand (2003)	GA	Ordering, Times
Gotteland et al. (2003)	GA	Ordering
Smeltink et al. (2004)	MILP	Ordering
Herrero et al. (2005); García et al. (2005)	GA	Times
Marín (2006)	MILP	Times
Balakrishnan and Jung (2007)	MILP	Times
Marín and Codina (2008)	MILP	Times
Roling and Visser (2008)	MILP	Times
Deau et al. (2008, 2009)	GA	Ordering
Keith and Richards (2008)	MILP	Ordering
Rathinam et al. (2008)	MILP	Ordering
Clare et al. (2009); Clare and Richards (2009, 2011)	MILP	Ordering
Yin et al. (2012)	MILP	Ordering

In this section, we present a comparison and categorisation of the existing research for the ground movement problem at airports, which has previously taken two main forms. The first form has involved the development of a Mixed Integer Linear Programming (MILP) formulation, to which a commercial solver was usually applied, yielding an optimal solution. Where models were formulated in a manner which would not be tractable to a MILP solver within a reasonable solution time, heuristic methods have been applied. This alternative approach has so far mainly involved the use of Genetic Algorithms (GAs). Of course, as heuristics, GAs give no guarantee of the optimality of the solutions found. However, their success over far shorter (and far more realistic in practice) execution times can often more than compensate for this.

We will first focus on the MILP formulations before discussing the GA-based approaches. For each approach, we will first discuss the various models which have been developed, before considering the previous research which has used these models in more depth. We will then compare the approaches, discussing the advantages and disadvantages of each. Finally, we

end this section by considering two important issues: firstly, how do the models handle the dynamic nature of the real problems at the airports, and secondly, how can speed uncertainty be handled to make the solution more robust in the real situation? An overview of the published ground movement optimisation research considered here can be found in Table 2.1, showing in chronological order both the solution approach which has been adopted and the defining characteristics of the model.

2.3.1 Mixed Integer Linear Programming (MILP) Formulations

MILP formulations are widely used by exact solution methods in operational research. In comparison to Linear Programming (LP) formulations where the objective function and constraints all have to be linear, MILP formulations introduce an additional restriction of integrality for some variables. Unfortunately, since this restriction changes the nature of the search space from continuous to discrete, it often leads to problems which are much harder to solve, so that solution times for large problems may no longer be practical.

Three different MILP modelling approaches are described below:

- Exact position approach: Here a time is allocated for each aircraft to traverse each individual part of its path. The approaches of Marín (2006), Balakrishnan and Jung (2007), Marín and Codina (2008) and Roling and Visser (2008) used a space-time network for this purpose. A spatial network representing the map of the airport is used as a starting point, then time is discretised and a copy of the underling spatial network is created for each time unit. These are then used to build a time expanded network. A good illustration of this can be found in Marín and Codina (2008).

- Ordering approach: In this case, rather than dealing directly with timings, the algorithm first aims to decide upon the sequencing, then uses this information to schedule times for each aircraft at each vertex or edge. This approach was adopted by Smeltink et al. (2004), Rathinam et al. (2008), Keith and Richards (2008), Clare et al. (2009), Clare and Richards (2009, 2011) and Yin et al. (2012). All of these only required a spatial network and modelled the sequencing constraints using binary variables, where the variables for

a pair (i,j) of aircraft at a vertex/edge are equal to one if and only if aircraft i passes this vertex/edge before aircraft j. With this approach, the times for each aircraft can be modelled as continuous variables, avoiding the disadvantages of time discretisation.

- Immediate predecessor/successor approach: It would also be possible to indicate only the immediate predecessor and successor for each aircraft at each vertex/edge rather than a full sequencing. As far as we can determine, this approach has not been used for solving the ground movement problem so far. Although the model in Smeltink et al. (2004) indicated the immediate predecessor aircraft, this was only to support the ordering model.

2.3.2 Review of Previous MILP-related Research

To our knowledge, Smeltink et al. (2004) was the first approach to handle the ground movement problem using the MILP formulation. This was performed for Amsterdam Schiphol Airport in 2004. Since this airport used standard, predefined taxi routes for aircraft, the problem was reduced to a scheduling problem. The approach worked on a spatial network where times were modelled as continuous variables and binary variables were used for the sequencing, as described above. The objective was to minimise the waiting time while taxiing and the deviation between the desired departure time and the scheduled departure time.

Marín (2006) presented a linear multi-commodity flow network model to simultaneously solve the aircraft routing and scheduling problem around airports. Two different methodologies were used to solve the MILP formulation: a branch and bound, and a fix and relax approach. In the latter case, the planning period was split into k smaller periods. Initially, only the variables within the first time period are taken as binary and a linear relaxation is applied to the variables for the other periods. The variables for the first period are then fixed, the variables for the second time period are made binary and the linear relaxation is maintained for the remaining variables. This is repeated for all k periods until all of the variables have been fixed. The objective of the MILP formulation was to minimise the total taxi time.

Marín and Codina (2008) later published further work where the model was multi-objective. The weighted linear objective function considered five other objectives, in addition to the previous goal of reducing the total routing time: 1) reducing the number of controller interventions,

2) reducing the worst routing time, 3) reducing the delays for arrivals, 4) reducing the delay for departures and 5) attempting to maximise the number of arrivals and take-offs. In contrast to other models, they allowed the aircraft to use the whole network and did not restrict them to a pre-determined set of paths. However, the presented algorithm was not able to deal with the separation constraints in an accurate way because the constraints were only modelled in the space-time network, which is independent of the type or size of aircraft.

Balakrishnan and Jung (2007) published another MILP formulation of the ground movement problem on a space-time network. In this approach, each aircraft could be allocated one of a limited set of routes. The relative benefits of different control approaches, such as controlled pushback and taxi path re-routing were also considered. Their aim was to minimise the total taxi time and to penalise situations where aircraft departed too late. It was pointed out that controlled pushback could reduce the average departure taxi time significantly, saving fuel.

An alternative MILP formulation for ground movement, which was also based on a space-time network, was provided by Roling and Visser (2008). A number of alternative routes were assigned to each aircraft beforehand, and only these were considered at the solution stage. It was possible for an aircraft to wait at the beginning of the journey, as well as on special vertices during the journey. The objective was to minimise a weighted combination of the total taxi time and total holding time at the gates. The objective function considered the entire route for each aircraft but the solution was only guaranteed to be conflict-free within the planning horizon, since these constraints were relaxed for later times.

Rathinam et al. (2008) used a MILP formulation which was based on the work of Smeltink et al. (2004) and primarily considered the ordering of the aircraft at vertices. Further separation constraints were added to the model, and it was simplified by reducing the number of binary variables. The algorithm used a spatial network and a predefined route for each aircraft, to minimise the total taxi time.

Keith and Richards (2008) introduced a new model for the coupled problem of airport ground movement and runway scheduling. Their MILP optimisation was influenced by the work of both Smeltink et al. (2004) and Marín (2006). The objective function was a weighted combination of minimising the makespan, the total taxi and waiting time and the total taxi distance. As

in Smeltink et al. (2004), a spatial network was used, with binary variables for handling the sequencing constraints and continuous variables for the timings. Although both wake vortex and en-route separations were considered for the take-off sequencing element, there were no route limitations applied. The work of Clare et al. (2009) extended their previous work. Their MILP formulation was changed to make it possible to introduce an iterative solution method for departing aircraft. In the first step, a relaxed MILP formulation was solved, and no guarantees were given for a conflict-free solution. An iterative procedure was then applied, where additional constraints were added where they were necessary to avoid any conflicts detected in the previous iteration. This was repeated until a conflict-free schedule was found. The papers by Clare and Richards (2009, 2011) expanded their work by incorporating arrival aircraft at Heathrow airport. The model simplifies the setting at an airport and cannot model pushbacks, variable taxi speeds and the detailed taxiway layout in the presented form. Moreover, it seems questionable if the approach is fast enough for real-time use and can deal with sudden changes.

Yin et al. (2012) published very recently a MILP formulation for George Bush Intercontinental Airport in Houston, Texas. The model aims to minimise weighted total taxi times by finding the taxi routes and the related schedules for each aircraft. However, the execution times for small simplified datasets were relatively long even when applying the concept of rolling horizon (see Section 2.3.7).

2.3.3 Genetic Algorithm (GA) Models

GAs are search methods inspired by evolutionary biology (Goldberg 1989; Sastry et al. 2005). They maintain a population of candidate solutions, have a method (called a fitness function) for evaluating solutions and apply a selection mechanism to guide the algorithm towards good solutions. The correct encoding of the problem can be key for the successful application of a GA (as we will consider in the next section), as can be the choice of appropriate mutation and crossover operators for the selected problem encoding.

The basic GA algorithm involves a repetition of the stages of evaluating the current population, selecting the population members to modify, applying the crossover and mutation operators and replacing old population members by new population members, as appropriate for the

replacement strategy (Sastry et al. 2005). These steps are repeated (iteratively improving the overall quality of the population of candidate solutions) until a given termination condition is met.

We now consider the important elements of the encodings which have been used for the ground movement problem over the last decade before considering, in Section 2.3.4, the specific encodings. As for the MILP approaches (discussed in Section 2.3.1), the GA approaches consider either the absolute timing or the relative sequencing of the ground movement.

All of the encodings which have been considered in the GA implementations, (Pesic et al. 2001; Gotteland et al. 2001, 2003; Gotteland and Durand 2003; Herrero et al. 2005; García et al. 2005; Deau et al. 2008, 2009), included the route allocation information, specifying the route r_i to allocate for each aircraft i. The additional information which was included differed between the approaches, but can be summarised into three categories:

- Applying an initial (aircraft-specific) delay/hold time, prior to pushback. The GA is responsible for determining this delay for each aircraft, as well as the route to allocate. This approach was adopted by Herrero et al. (2005) and García et al. (2005).

- Applying a delay at some point during the movement, but not restricting it to being applied at the start of the taxiing. This could be implemented either by specifying times for both initiating and terminating the delay (the approach which was adopted in Pesic et al. (2001) and Gotteland and Durand (2003)) or as a delay amount and (spatial) position at which to apply it to the aircraft, as in Gotteland et al. (2001). The GA is responsible for investigating when or where to apply the delay and the duration or end time of the delay as well as the route to allocate to the aircraft.

- Prioritising aircraft movement, where the GA is used to investigate the relative prioritisation of the aircraft rather than allocating holds directly. Here, the priority determines which aircraft take precedence when there are conflicts during the movement. This approach was adopted in Gotteland et al. (2001, 2003), Gotteland and Durand (2003) and Deau et al. (2009), where the GA explored the priorities to assign to aircraft as well as the routes.

2.3.4 Review of Previous GA-related Research

As far as we can determine, Pesic et al. (2001) published the first paper for optimising the ground movement problem at airports in 2001. They allowed a single delay per aircraft at a time determined by the GA. Their fitness function considered the number of time steps C, for which aircraft were in conflict during the movement, and the total travel time T for aircraft. The GA aimed to maximise the fitness value, which was $\frac{1}{2+C}$ in the presence of conflicts or $\frac{1}{2}+\frac{1}{T}$ in the absence of conflicts. Thus, all values larger than $\frac{1}{2}$ corresponded to solutions which were conflict-free and all values smaller than $\frac{1}{2}$ had at least one conflict and were therefore infeasible. Crossover and mutation operators were introduced, along with a diversification strategy and some simple termination criteria. For a random pair of parent solutions, the crossover operator chose for each aircraft the parent which had fewer conflicts with other aircraft, in order to increase the probability of producing an offspring population with better fitness values. This operator was appropriate because the problem was partially separable as defined and discussed in Durand and Alliot (1998). The mutation modified the details for the aircraft with the (potentially shared) worst local fitness value.

Gotteland et al. (2001) extended their previous work by considering how the GA could deal with speed uncertainty. We believe that this is an important consideration and will discuss it in Section 2.3.8. In addition to the encoding from their previous work (Pesic et al. 2001), they used a representation for prioritising aircraft movements, discussed in Section 2.3.3. The encoding included the route number and priority level for each aircraft. A fitness value was computed by applying an A* algorithm with the specified prioritisation of the aircraft. A space-time network was then generated and aircraft were routed in order of priority level. After an aircraft had been routed, the network was adjusted in such a way that the allocated route was removed, along with all potentially conflicting edges, so that the routing of the next aircraft avoided conflicts with previous aircraft.

The clustering of aircraft within these ground movement problems was considered in Gotteland et al. (2001). A two stage approach was adopted, where the clusters of aircraft with conflicts were solved independently in the first stage, before the different clusters were unified and solved in combination in the second stage.

Gotteland and Durand (2003) subsequently presented an alternative sequential algorithm: a branch and bound algorithm, with a first search strategy (Horst and Tuy 1990) replacing the A* algorithm to speed up the calculation of the fitness value, since there is always a preference to continue taxiing rather than to hold position.

Gotteland et al. (2003) explained the way in which their GA handles both take-off time prediction and CFMU slots. They modified their algorithms from Gotteland et al. (2001) with the aim of reducing the deviation from CFMU slots (rather than minimising the necessary taxiing time) by penalising (with a linear cost) deviations from the desired take-off times for each aircraft, with a steeper penalty when the scheduled take-off is outside the CFMU slot.

García et al. (2005) hybridised two earlier approaches which were previously detailed by the same authors in Herrero et al. (2005). A modified minimum cost maximum flow algorithm determined the initial population of a GA and was used to penalise the fitness function. The approach considered the application of an initial delay at the gate and the allocation of a route to each departing aircraft, with no possibility for waiting at intermediate points or slower taxiing during the ground movement. They used tournament selection, single-point crossover, a traditional mutation operator and an additional random variation of the delay time. Their fitness function penalised infeasible solutions and tried to minimise the makespan and the sum of the delays, while attempting to maximise the number of departing aircraft.

Two more recent papers from Deau et al. (2008, 2009), developed the ideas which had been discussed in Pesic et al. (2001), Gotteland et al. (2001, 2003) and Gotteland and Durand (2003). They proposed a two-phase approach which considered the runway sequencing in the first stage and the ground movement in the second stage. The separations to account for the wake vortices were the most important constraint for the runway sequencing element. A deterministic constraint satisfaction problem solution algorithm was used, which was based on a branch and bound methodology. They used an objective function which was similar to that used in Gotteland et al. (2003). Departing aircraft were moderately penalised if their scheduled time deviated from the desired time within the CFMU slot, but were much more heavily penalised if the scheduled time was outside this slot. Arriving aircraft had a fixed predicted time to land, so a solution was only feasible if these aircraft had, at most, a small delay (no more than one

minute) compared with the predicted landing time. In the second stage, their GA was modified to find a good solution for the ground movement problem given the runway sequencing from the first stage. The target runway sequence was considered as the ideal result of the routing stage but was not treated as a hard constraint. Thus, the fitness function for their GA penalised deviations from the target times. However, their publications do not present the execution times. Moreover, it seems unclear as to whether the approach is robust enough and can flexibly deal with changes.

2.3.5 Other Solution Approaches and Developments

Other solution approaches and developments for the ground movement problem at airports are discussed in the following section.

Andersson et al. (2000) used simple queueing models for the taxi-in and taxi-out processes to better understand the airport dynamics. These queuing models were later extended by Carr et al. (2002). The Ground-Operation Situation Awareness and Flow Efficiency (GO-SAFE) system was presented by Cheng and Foyle (2002) and Cheng (2003), but the technical details were omitted since the paper was from a commercial company and it was only stated that a dynamic programming approach using Dijkstra's algorithm was used for the ground movement problem. Experiments on a simple mock symmetric airport with six aircraft showed that an event-based A* algorithm outperformed a co-evolutionary strategy with respect to the cumulative time of completion (Brinton et al. 2002). Confessore et al. (2005) presented a discrete simulation-based architecture for a decision support system for taxiing on the apron. Different objectives where analysed on Rome Fiumicino Airport where significant improvements could be achieved compared to the approach used previously. Baik and Trani (2008) introduced a time-based simulation model for the analysis of airfield operations which incorporates a time-dependent shortest path algorithm published in Baik et al. (2002). A sequential A* algorithm was utilised in the papers by Lesire (2009, 2010) to deal with speed uncertainty on Toulouse-Blagnac airport (our approach in Chapter 6 is similar to this, but provides a better coverage of the solution space). The approach minimised the travel time of each aircraft and the execution time was very fast. Gupta et al. (2010b) presented results on a simplified model of the primary

international airport in Mumbai (BOM), with a network consisting of 15 nodes and 22 edges and a testset with 9 aircraft, where the approach was mainly based on the work by Smeltink et al. (2004). A different approach was presented by Mori (2010), who used cellular automata to simulate congestion for departing aircraft on predefined routes. The model divides the layout of the airport into small cells and the position and the speed of each aircraft is modelled in a process which was first developed to describe the traffic on highways (Nagel and Schreckenberg 1992; Esser and Schreckenberg 1997).

The benefits that can be gained by using the Airport Surface Detection Equipment Model X (ASDE-X) to have better awareness on the ground are highlighted in the work by Bhadra et al. (2011) and Srivastava (2011). ASDE-X obtains data from radars, sensors, aircraft transponders and ADS-B (Automatic Dependent Surveillance-Broadcast).

Finally, the paper by Cheng (2007) serves as a progress update on effort to introduce flight-deck automation systems for pilots, so that more detailed information can be delivered from the tower to the cockpit in an automatic manner.

2.3.6 Comparison of the Approaches

We now consider the major differences between the different models and solution approaches.

2.3.6.1 Differences in objectives and constraints

The optimisation of airport operations is a real-world problem, and as such it is important that the real objectives of the airport and real constraints upon the problem are considered. The majority of the published work has considered real airport settings, and it is apparent that both the objectives and the details of the constraints have differed between airports. Consequently, the models for the problems have also differed, resulting in the development of different solution approaches.

2.3.6.2 Optimality vs. execution time

The solution method which is adopted may also depend upon the load upon the airport (i.e. the number of aircraft which need to be simultaneously considered), since exact solution approaches become less practical as loads increase. With the expected increases in the density of air traffic meaning that airports have to be able to handle more aircraft in the near future, some solution techniques may potentially need to be adjusted over time.

GAs are heuristics rather than exact solution methods and can, therefore, often give neither any guarantee for the solution nor even an approximation ratio in many situations. However, a poor formulation of a MILP can also mean that an exact solution to the MILP can be a poor solution for the underlying real-world problem. For example, with time discretisation models, the way in which the time discretisation is handled can have a major effect upon the quality of the results: smaller intervals may give better results but will result in significantly larger problems to solve. Similarly, the way in which a model deals with the separation rules between aircraft can also affect the quality of the results. It should be noted that none of the papers which were discussed here measured the optimality gap for realistic scenarios, evaluating the effects of utilising only a heuristic (GA-based) solution approach or of the effects of time discretisation, perhaps due to the difficulty or impracticality of optimally solving these problems. In our opinion, it would be worthwhile to have some kind of comparison between the performance of the approaches, to be able to see the trade-off explicitly, but several publications do not provide all the needed information to reproduce their results and enable a rigorous comparison.

Due to the fact that airports are usually interested in real-time decisions, the execution time of an algorithm is a crucial measure. From this point of view, heuristics such as GAs often outperform MILP formulations. For example, in Roling and Visser (2008) it was shown that the execution time increased dramatically as the number of aircraft increased for their MILP formulation.

Different researchers have also used different objective or fitness functions, due to having slightly different aims. We believe that the generation of some generic benchmark scenarios to allow such an analysis to be performed, comparing exact and heuristic solution approaches and the effects of different objective functions, would be of huge benefit.

As far as we are aware, there has been no investigation using other metaheuristics such as simulated annealing (Aarts et al. 2005), or tabu search (Gendreau and Potvin 2005). Furthermore, there seems to be an unexploited potential for hybrid approaches which can make use of the advantages of different models. More discussion about heuristics can be found in Section 6.5.

2.3.7 Dealing with the Dynamics

One major characteristic of the problem of ground movement at airports is the dynamic nature of the problem. Predictions become less accurate the further they are in the future: predicted positions for current aircraft may be wrong as may be predictions of when new aircraft will be ready to pushback from the gates or to land. Predictions, therefore, have to be regularly updated and, since some approaches need a significant execution time, attempts have been made to decompose the problems into smaller sub-problems. In this section, we summarise the approaches which have been used to cope with the dynamic nature of the routing problem.

- A simple modelling approach, called *shifted windows*, was introduced by Pesic et al. (2001) for their GA. Every Δ minutes, the situation was resolved for a fixed time window. Only arriving or departing aircraft within the time window were considered but the time window was enlarged for these aircraft to avoid horizon effect problems.

- Smeltink et al. (2004) evaluated three different variants of a *rolling horizon* approach, not only for handling the dynamics of the problem, but also to reduce the size of the problem to be solved. In each case, the planning period was split into disjoint, equal length time intervals. In the first variant, the routes which had been allocated in previous intervals were considered to be fixed, whilst in the second variant they could be modified. In the third variant, the aircraft were sorted according to their pushback or landing time, respectively, and a *sliding window* was applied to consider m aircraft in each iteration. The first iteration considered aircraft 1 to m, then aircraft 1 was fixed and aircraft 2 to $m+1$ were considered, then aircraft 2 was fixed, and so on. Unfortunately, this variant had a significantly higher execution time without increasing the solution quality significantly.

- The fix and relax approach (discussed in Section 2.3.2) which was used by Marín (2006) for

solving his MILP formulation, worked in a similar way to the sliding window approach. He also used an alternative time-interval-based approach, where only aircraft in a particular interval were used for planning but the interval was not enlarged to guarantee a conflict-free solution. Instead, a shortest path algorithm was used to estimate the remaining time for the aircraft which do not reach their destination within the interval.

- Lesire (2009, 2010) suggested a routing and scheduling approach which considers the aircraft sequentially. Even though a simultaneous approach may find better solutions to this problem, this method has several benefits. A decision support system can better react if an aircraft has a delay and less or no changes to the other aircraft's routes have to be applied, since the other aircraft's routes are considered to be fixed in most cases. Chapter 6 introduces a sequential routing and scheduling approach and highlights that the solution quality is very good. Moreover, both the approach by Lesire (2009, 2010) and Chapter 6 report very fast execution times making them more suitable for the real-time use at airports.

2.3.8 Robustness and Speed Uncertainty

Almost all published approaches were based on deterministic data. However, the real world situation at airports is less predictable. Therefore, we think it is important to take solution robustness into consideration. Uncertainty in the data for the ground movement problem can appear in different areas, one of which is speed predictions (see Chapters 4 and 5). An approach to cope with this was presented and illustrated in Gotteland et al. (2001). They modelled the speed uncertainty as a fixed percentage of the predefined speed. Hence, an aircraft was assumed to occupy not only a single position in the network but multiple possible positions at the same time. While an aircraft was taxiing, the number of occupied positions grew and when an aircraft was waiting at a holding point, the speed uncertainty and number of occupied positions decreased. More discussion about buffer times can be found in Section 6.4.7.

2.4 Related Research Areas

Similar application problems have been considered in other areas of research, such as job-shop scheduling with blocking (Hall and Sriskandarajah 1996), the control of Automated Guided Vehicles (AGVs) (Vis 2006; Nishi et al. 2011; Schüpbach and Zenklusen 2012) or robots (Nishi et al. 2005) and train routing and scheduling (Cordeau et al. 1998; Caimi et al. 2011). Of course, the details of the constraints and objectives differ, so there are limits to the applicability of the research.

2.4.1 Job-shop Scheduling with Blocking

The job-shop scheduling problem is a standard problem in Operations Research in which jobs are allocated to resources in a sequential manner. In general, a finite set of jobs is given, together with a chain of operations which need to be executed on a finite set of machines or resources. Each operation has an execution time and a resource can normally only handle one job at a time. The goal is to find an allocation of the job's operations to a time interval of the resources, such that the makespan is minimal. Many different variations have been studied and applied to different areas (Pinedo 2012). The version which is closest to the ground movement problem at airports seems to be the job-shop scheduling with blocking problem (Hall and Sriskandarajah 1996; Mascis and Pacciarelli 2002; Brucker 2007). Aircraft can be treated as jobs, resources represent the different taxiway parts and a fixed route of an aircraft defines the chain of operations to which this aircraft have to be allocated. The constraint of fixing the aircraft's route may affect the solution quality and may even make a certain ground movement scenario infeasible which could be solved with more flexibility of the possible routes. In the variation of the job-shop scheduling with blocking problem, a job remains on the resource after its processing until the downstream resource becomes available. Gröflin and Klinkert (2009) presented an approach to solve a generalised version of the job-shop scheduling with blocking problem with tabu search on a generalised disjunctive graph.

2.4.2 Controlling AGVs or Robots

In the case of controlling AGVs or robots, it is less important to avoid permanent changes to the routes and schedules of vehicles which are already moving. Vis (2006) presented a survey paper about the design and control of AGV systems. Nishi et al. (2005) first presented a decentralised approach using Lagrangian decomposition and coordination techniques. Afterwards, another approach was introduced with the decomposition of Petri nets and Lagrangian relaxation (Nishi et al. 2009). Results showed execution times of the approach were less than 1 second on average for scenarios with up to 9 AGVs, but scenarios with 15 AGVs had an average execution time of 13 seconds, with one particular scenario using 36 seconds. Nishi et al. (2011) later presented a bilevel decomposition approach using a mixed integer formulation for the simultaneous conflict-free routing of AGVs. Their results with up to 4 AGVs on a relatively simple layout suggest that the approach is too slow to be used in an online setting. Richli (2009) in his Master's dissertation introduced different approaches to solve the simultaneous AGV routing problem, but again the reported execution times of the algorithm suggest that sequential approaches, as presented by Gawrilow et al. (2008), are more appropriate for real-time cases. A different approach again by Tanaka et al. (2010) tackled the simultaneous routing problem using Petri nets, but reported execution times which were too long for the use at airports. Very recently, Schüpbach and Zenklusen (2012) published work on adaptive routing for personal rapid transit. The idea behind this work is to reoptimise the routes of all currently used vehicles, when a new vehicle requests a new route. The approach was based on solving a minimum cost multi-commodity flow problem on a time-expanded graph with column generation and randomised rounding.

2.4.3 Train Routing and Scheduling

In the field of conflict-free routing and scheduling for trains, Caimi et al. (2011) presented an integer linear programming formulation which utilises information from conflict cliques as strong cutting planes. This new approach was able to massively reduce the execution time compared with other known algorithms. Earlier, D'Ariano et al. (2007) published work for scheduling trains in a railway network based on a branch and bound algorithm. Often such

approaches fix the speed and prohibit any waiting outside of the stations, either since they are focusing on the routing and scheduling within station areas only or to explicitly simplify the problem.

2.4.4 Complexity of the Problem

As far as we are aware, there is no published proof that the simultaneous ground movement problem is NP-hard. However, there are many similarities with the job-shop scheduling with blocking problem (Mascis and Pacciarelli 2002; Brucker 2007) and with the multi-commodity integral flow problem. The multi-commodity integral flow problem with only two commodities has been shown to be NP-complete (Even et al. 1976). Basically a space-time network of the airport layout can be used as a directed graph where each edge has a capacity of 1 and the aircraft can be considered as commodities with source, target and demand which are set to 1 showing the close relation between the two problems. Schüpbach and Zenklusen (2011) showed very recently that a simplified version of the conflict-free vehicle routing problem is NP-hard even on paths, using a reduction from the 3-partitioning problem. The ground movement problem varies depending on the problem description and also on the objective function. It is unclear whether special cases exist which can be solved in polynomial time, but that would only be of particular interest if such special cases were relevant for the support of airport operations.

2.5 Taxi Time Prediction

In visually analysing the average taxi speed of different airports, it was obvious that major differences appeared depending on various factors. Major differences are apparent between arriving and departing aircraft as well as from whether the amount of traffic at the airport is low, medium or high at the time. Since the effect of the taxi times do not appear to have been sufficiently incorporated into the current state-of-the-art research, we are also interested in predicting taxi times accurately and use this information later in our approach to have a more realistic decision support system. Chapters 4 and 5 provide more background about this problem domain.

2.6 Fuel Efficient Taxiing

There is not much coverage of environmental considerations in taxiing. Previous research mainly focused upon stand holding in order to reduce fuel burn (Burgain et al. 2009; Atkin et al. 2010a, 2011a; Simaiakis et al. 2011). The assumption made was that by reducing the total taxi time one can simultaneously improve the efficiency of airport operations and reduce the fuel consumption, but Chen and Stewart (2011) indicated that this may not always be the case. We suggested the value of considering speed profiles when routing aircraft to avoid unnecessary fuel burn due to acceleration and deceleration (Atkin et al. 2010b).

Chapter 7 shows an analysis of the trade-off between the total taxi time and the fuel consumption for the conflict-free routing problem for aircraft on the airport's surface, which is also published in Ravizza et al. (2012b). In contrast to the approach of Chen and Stewart (2011), the interactions between aircraft are considered, instead of analysing a single trajectory of an unimpeded aircraft. Interactions affect speed profiles of the aircraft involved and massively increase the solution space of the routing approach.

2.7 Spotting the Needs in the Area

In this section, we describe several important open research directions for the airport ground movement problem.

2.7.1 Consistency and Comparability

As discussed in Section 2.3.6, the constraints and objectives vary widely within the published research. No comparison has so far been performed between different approaches, so it is difficult to estimate the gap between the exact optimisation methods (e.g. MILP formulations) and the heuristic approaches (e.g. GA) for either the quality of the solution or the execution time of the algorithms. More consistency is desirable. Results about our routing and scheduling approach are discussed in Section 6.6.

2.7.2 Integration of Other Airport Operations

The integration of other airport operations, such as departure and arrival sequencing and gate assignment, is highly desirable and, ultimately, optimisation across multiple airports would be even better. Of course, the complexity of the integrated problem would grow and, since the computation is time-critical, there seems to be more potential for heuristic and hybrid methods than exact approaches. With the integration of different airport operations, the problem may also have to be treated as a multi-objective optimisation problem.

2.7.3 Robustness and Uncertainty

Uncertainty in the input data is common at airports. Pushback time uncertainty and taxi speed/duration uncertainty are known to be major limiting factors upon the accuracy of models. We see the need for more investigation into models of the airport ground movement problem which are more robust against such uncertainty.

2.7.4 Environmental Considerations in Taxiing

Consideration of the environmental effects of airports has become increasingly important and could be taken into account for ground movement. For example, where possible, delays for an aircraft should be scheduled prior to starting the engines, i.e. as initial delays at the gate/stand.

Perhaps more interestingly from the point of view of the problem modelling, aircraft engines are more efficient when a constant taxi speed can be maintained rather than having a lot of acceleration and deceleration. Speed changes and multiple stops should, therefore, be avoided or reduced. It may be advisable to consider some kind of post-processing to calculate speeds for link traversals, so that the pilots could be given appropriate information to allow them to replace higher speed taxi operations plus waits by a lower speed operation. Chapter 7 takes this into account.

2.7.5 Limiting Changes

When the real-world dynamic case is considered, it is possible that routes or sequencing can change over time. This may be highly undesirable if information has been transmitted to pilots. Thus, the effects of avoiding changes should at least be considered. These issues are considered in Chapter 6 and this has not been widely discussed in other publications.

2.8 Conclusions

A good proportion of this chapter was published as a conference paper (Atkin et al. 2010b) and provided the first overview and comparison of the various ground movement models and solution methods in the literature. It is apparent that there are significant differences between both the objectives and the constraints which were utilised in previous research. To some degree this is inevitable due to the differences between airports and different stakeholder aims. However, there is obvious benefit to be gained from a formalisation of these. The state-of-the-art approaches use mostly either a MILP formulation or a genetic algorithm approach and a categorisation of the representations has been provided for both.

In addition to highlighting the state-of-the-art in this research area, a number of interesting and important future research directions have also been identified. Of particular importance is the integration of other (highly-related) airport operation problems. Runway sequencing (for both departures and arrivals) and gate assignment are highly connected to the problem of airport ground movement and we suggest that there would be benefits from handling them simultaneously. More consistency within airport operations would also be helpful and generic benchmark scenarios would be useful for both quantifying algorithms and encouraging further research by those who may not have direct contact with an airport. Finally, we have identified the importance of handling uncertainty in taxi speeds and generating robust solutions and of considering the operational limitations of communicating instructions to pilots and the environmental effects of decisions.

3

Datasets Used in Experiments

> We have no idea about the 'real' nature of things ... The function of modeling is to arrive at descriptions which are useful.
>
> Richard Bandler and John Grinder

3.1 Introduction

This chapter reviews the datasets which are utilised in the experiments presented in the upcoming chapters. It is important in such a project to have access to real world data, because one of the main aims of this project was to bridge the gap between theory and practice and enhance decision support systems at airports. Datasets from three different airports in Europe were used for the different experiments. It is important to clarify that non-disclosure agreements were signed for all of the different datasets which restricted us from openly publishing the datasets, stating very detailed information about the data in this thesis and from running all of the experiments with all datasets.

The setting and experiments with the world's busiest airport (Hartsfield-Jackson Atlanta International Airport) are shown in Appendix D as a case study of how to apply some of the findings introduced in this thesis to a new airport.

The rest of the chapter is structured as follows: An overview is first given of the available datasets and how the airport layouts are modelled as a graph. Afterwards, each of the three airports is explained and visualised, before ending the chapter with an aircraft categorisation.

The main elements of all of the supplied data consisted of information about each aircraft, detailing the terminal gate or remote stand, the runway, the start and end time of taxiing, the aircraft type and whether the aircraft was an arrival or a departure. The considered data also included information about the airport layout, the positions of stands and the runway entrance and exit points and the layout of all of the taxiways. This information was used to represent an entire airport layout as a directed graph, where the edges represent the taxiways and the vertices represent the junctions or intermediate points. Aircraft are considered to occupy edges in our routing and scheduling approach (see chapter 6), and conflicts are avoided by preventing any two aircraft from using the same edge simultaneously.

3.2 Airports

3.2.1 Zurich Airport

This thesis utilised data from Zurich Airport (ZRH), which is the largest airport in Switzerland and a hub airport for Swiss International Air Lines. It was reported that the airport had around 24.3 million passengers and 279000 movements in the year 2011.

The airport has three runways, named 10/28, 14/32 and 16/34, according to their direction of operation, with the first and the last runways intersecting each other (see Figure 3.1). It was confirmed by the field staff that, as long as no heavy winds occur, ZRH operates with three operational modes: a) before 7am, runway 34 is used for arrivals and 32 and 34 for departures; b) during the day, runways 14 and 16 are used for arrivals and 28 and 16 for departures c) after 9pm, only runway 28 is used for arrivals and runways 32 and 34 are used for departures. The mentioned rules only apply on weekdays and outside the holiday times of Baden-Württemberg.

The considered data included information about the airport layout, the positions of stands and runway entrance and exit points and the layouts of all of the taxiways. Figure 3.2 visualises the

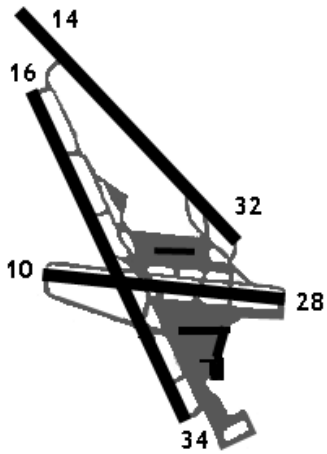

Figure 3.1: Sketch of Zurich Airport (ZRH)

graph which was created to represent the airport's layout in Zurich, with 465 vertices, 553 edges and 119 gates. It also included the real timings for the aircraft using the airport during each day. This information was used to develop a taxi time prediction function, as discussed later, to improve the accuracy of the taxi time predictions which are used in the ground movement model. We had access to data for an entire week's operations between the 27th of June and the 3rd of July 2011. No extraordinary occurrences took place and there were 5613 movements in total (2806 arrivals and 2807 departures). This dataset is referred to as "ZRH 2011" within this thesis. Additionally, an older dataset contains 679 movements in total (337 arrivals and 342 departures) from an entire day's operations for the 19th of October 2007. This dataset is referred to as "ZRH 2007". Figure 3.3 shows the total amount of traffic on the surface at different times of the day for "ZRH 2007", with different colours for arriving and departing aircraft. Taxi times at Zurich Airport varied from around 1 to 12 minutes (with a mean of 4.43 minutes) for arrivals and 4 to 24 minutes (with a mean of 8.88 minutes) for departures. Departures often need longer for the taxiing process due to waiting in a runway queue before take-off.

Figure 3.4 visualises the amount of traffic for the dataset "ZRH 2011". Instead of providing the information of whether an aircraft is arriving or departing, the different days are indicated by different colours. During the working days the patterns look very similar, with considerably fewer movements on Wednesday. The amount of traffic at the weekend (indicated with dashed

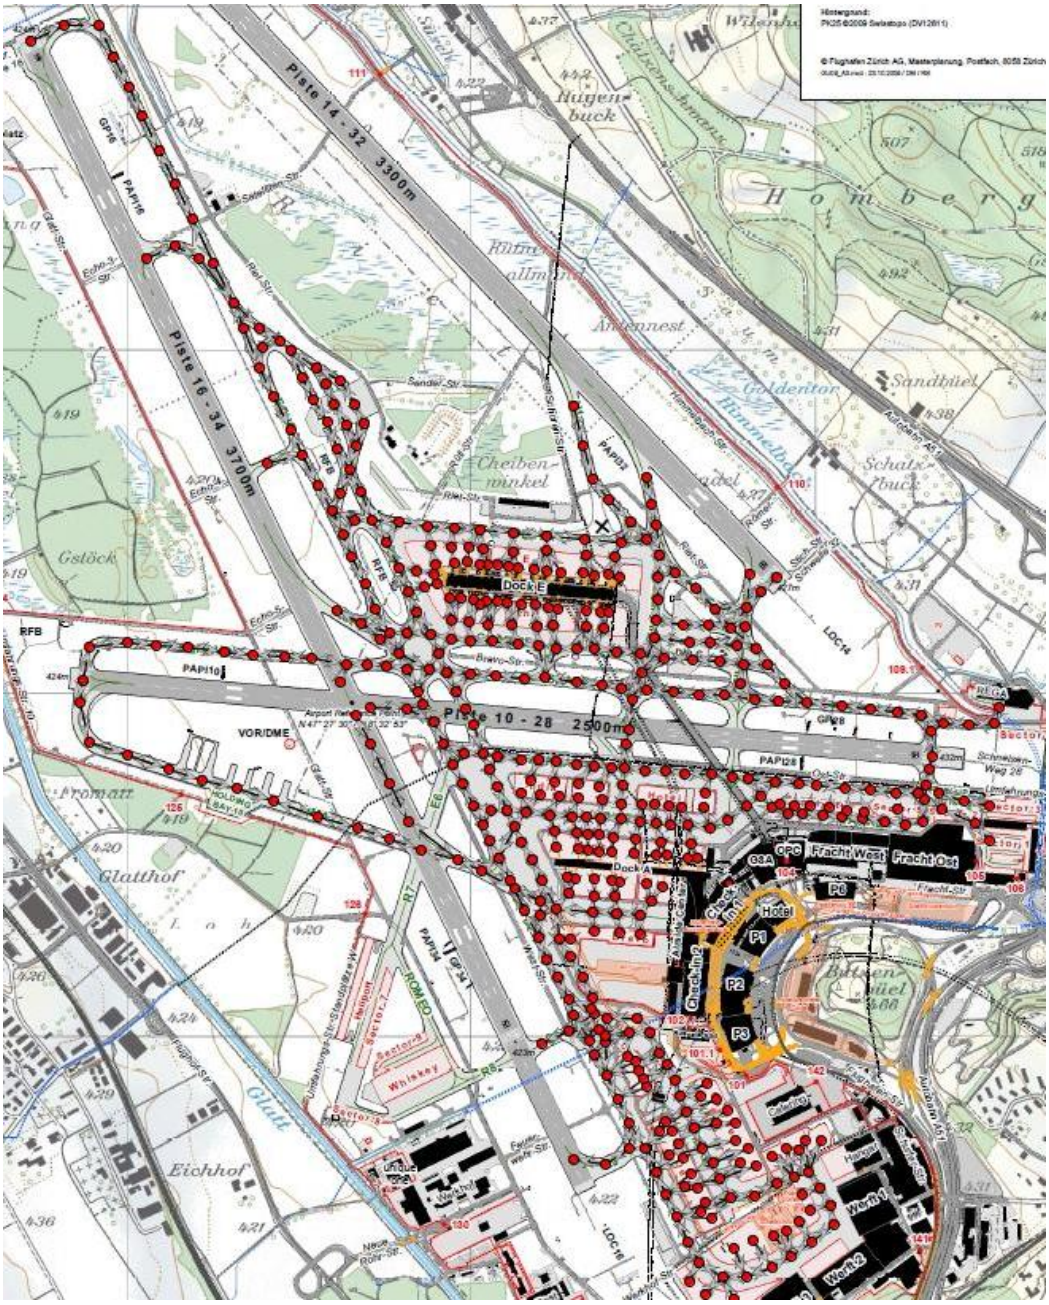

Figure 3.2: Layout from Zurich Airport modelled as a graph with vertices and edges

lines) follows a different pattern and fewer movements are recorded. The mean taxi times for this dataset were 3.60 and 11.13 minutes for arrivals and departures, respectively, with maximal taxi times of around 15 and 34 minutes, respectively,

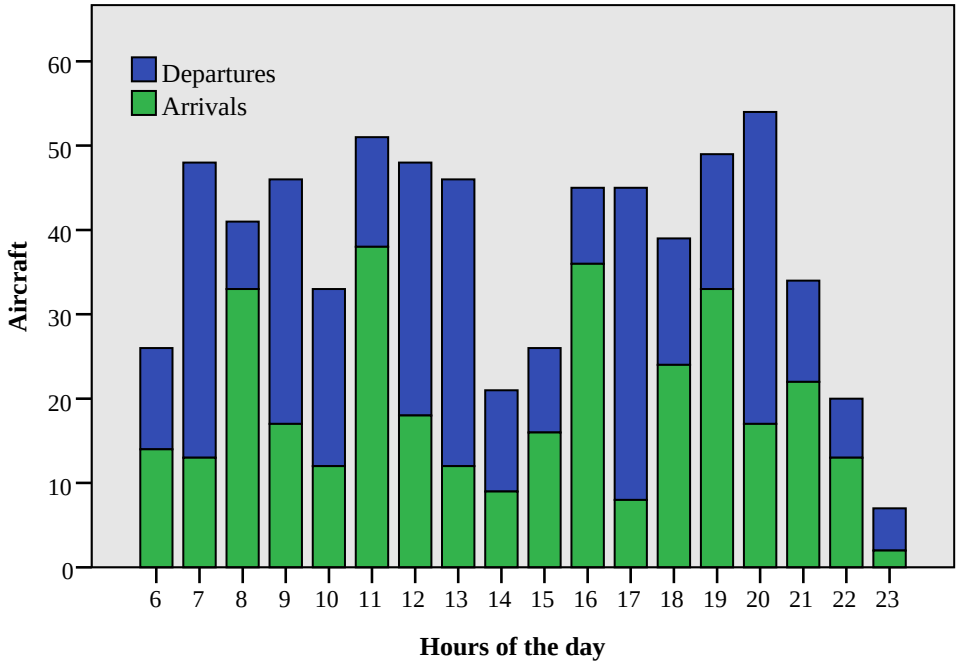

Figure 3.3: Hours of the day at Zurich Airport for the datatset from the year 2007

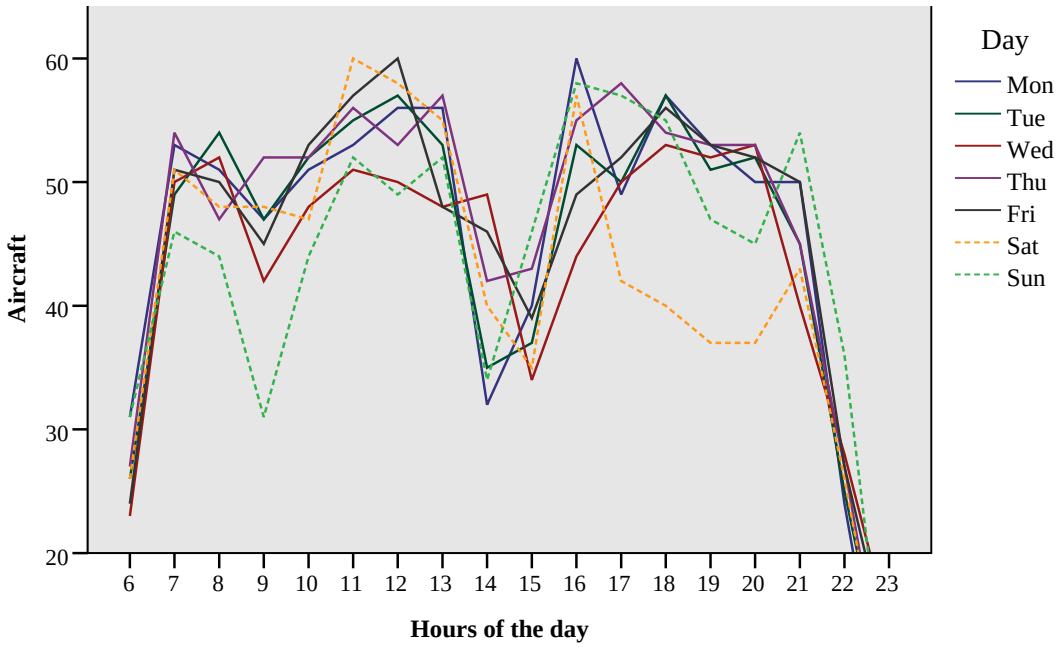

Figure 3.4: Hours of the day at Zurich Airport for the datatset from the year 2011

3.2.2 Stockholm-Arlanda Airport

Another dataset which was used for some of the analysis in this thesis was from Stockholm-Arlanda Airport (ARN), the largest airport in Sweden. Their main hub carrier is Scandinavian Airlines. Around 19 million passengers and 105400 landings were reported during the year 2011.

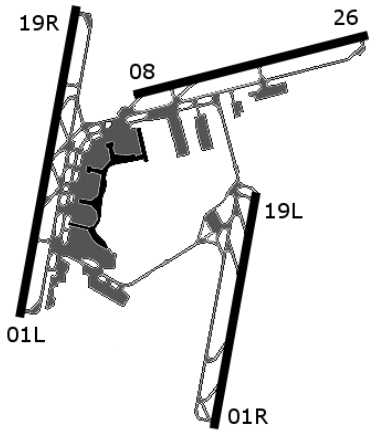

Figure 3.5: Sketch of Stockholm-Arlanda Airport (ARN)

The airport has three runways, named 08/26, 01L/19R and 01R/19L. The latter two runways are parallel and named with an "R" and an "L" depending upon whether they are on the right or left side according to the facing of the aircraft. A sketch of the airport layout is provided in Figure 3.5. Figure 3.6 shows the graph representing taxiways on the surface, which has 317 vertices, 349 edges and 91 gates. We had access to the data for an entire day's operations at Stockholm-Arlanda Airport for the 7th of September 2010 (661 movements, with 326 arrivals and 335 departures) with no extraordinary occurrences. This is referred to as "ARN 1". Aircraft within this dataset were most often landing on either runway 19L or 19R and departing from either runway 19R or 08. The average taxi time was around 7.4 minutes, with a maximal taxi time of 20 minutes. It is clear from Figure 3.7 that airside airport operations have a peak in the morning between 7am and 9am and also more movements during the late afternoon and early evening. At a later date, a second datasets was made available for the 14th of October 2010. It consists of 656 movements and it will be referred to as "ARN 2".

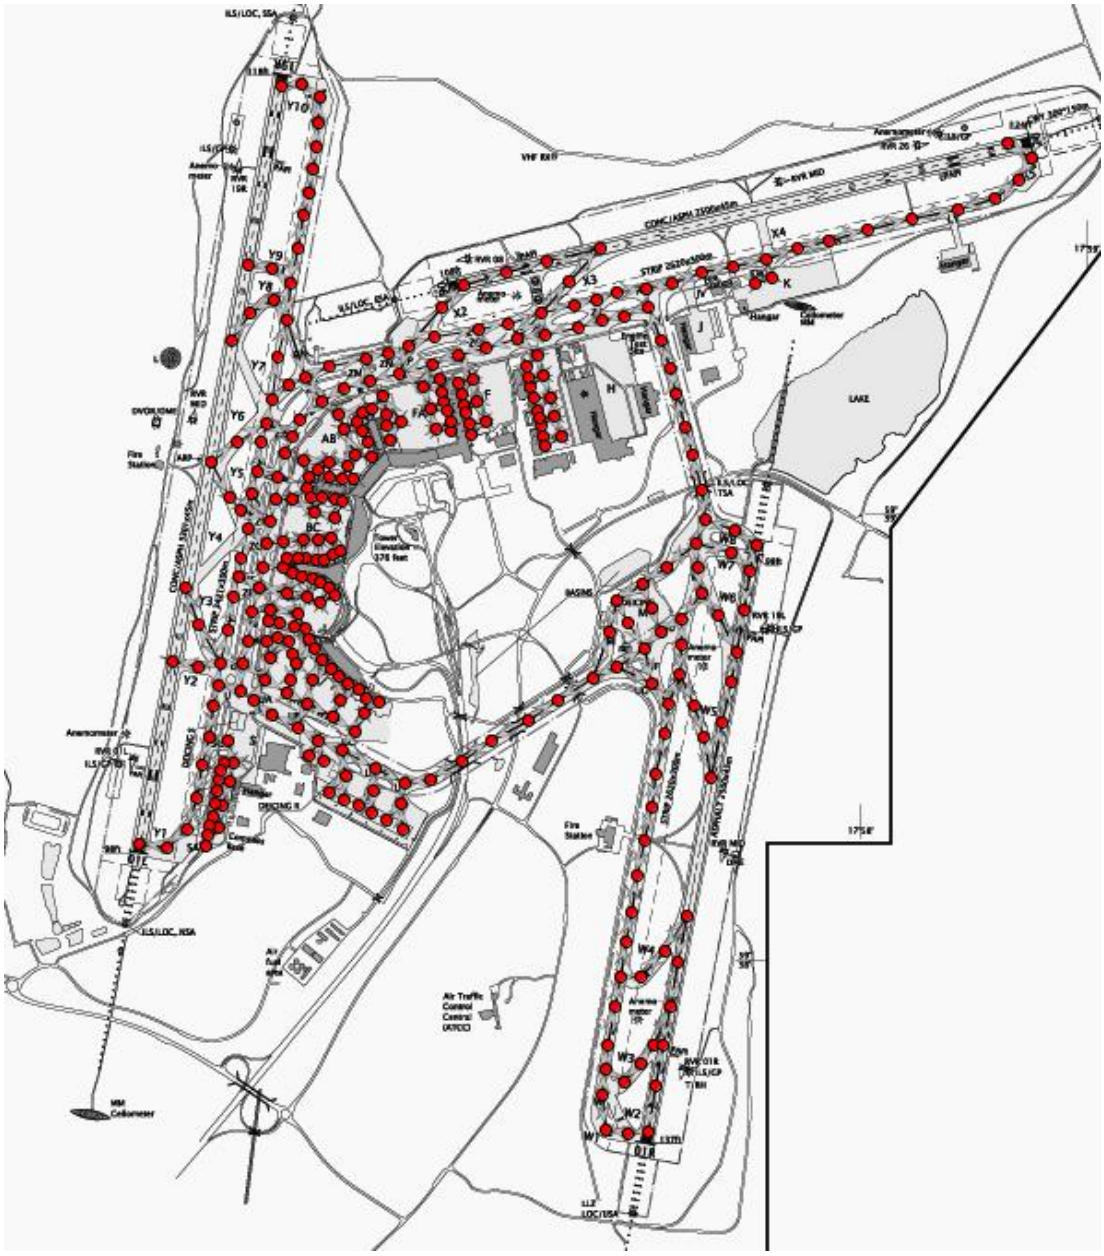

Figure 3.6: Layout from Stockholm-Arlanda Airport modelled as a graph with vertices and edges

3.2.3 London Heathrow Airport

The last dataset utilises real recorded data from London Heathrow Airport (LHR), supplied by NATS Ltd. London Heathrow Airport is one of the busiest international airports in the world with around 65 million passengers and 455000 movements a year, despite the fact that

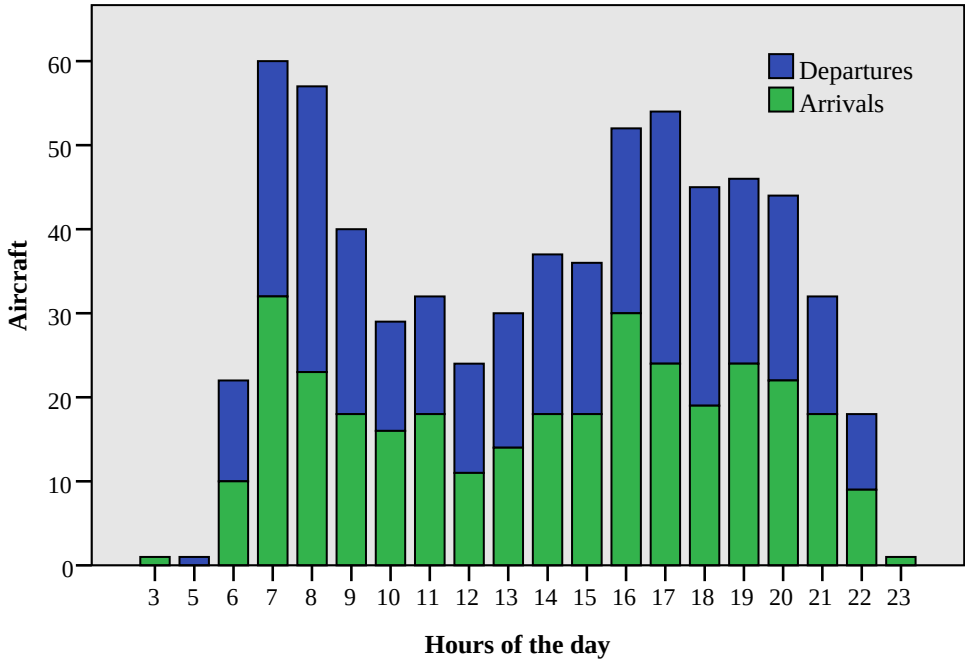

Figure 3.7: Hours of the day at Stockholm-Arlanda Airport

it operates with only two runways and (for noise control reasons) is restricted to using only a single runway at a time for departures. Heathrow is a primary hub airport for British Airways and British Midland International (BMI) and also a base for Virgin Atlantic Airways Limited.

The dataset was recorded during the entire week from the 5th to the 11th of July 2010. All of the 4727 arrivals and 4728 departures (9455 in total) were landing or departing from either runway 27R or 27L (direction west) and none from the east on runway 09L or 09R. The graph representing the surface of Heathrow is shown in Figure 3.8 which consists of 559 vertices, 642 edges and 197 gates.

The taxi times for incoming aircraft were on average 8.25 minutes, with a maximum of 63.12 minutes. The outgoing taxi times were considerably longer with an average of 22.09 minutes and a maximum of 171.53 minutes.

Figure 3.8: Layout of London Heathrow Airport, modelled as a graph with vertices and edges

3.3 Aircraft Categorisation

The following aircraft categorisation was mainly used for the experiments in Chapter 7, with data from Zurich Airport. Aircraft were classified into different groups and for each group the settings for a representative aircraft type were used for calculations. This procedure was necessary due to the lack of detailed data in the provided datasets. Aircraft were distinguished by their wake vortex separation group. The group 'light' was represented by the settings for a Cessna 172 Skyhawk. The settings for an Airbus A320 were used for the wake vortex group 'medium'. Finally, the group 'heavy' was represented by the settings for an Airbus A333. All of these aircraft were the most common aircraft type in their category at Zurich Airport. The technical details of the aircraft and their engines can be found in Table 3.1.

Table 3.1: Specifications of aircraft and engines

	Cessna 172 Skyhawk	Airbus A320	Airbus A333
Maximum take-off weight	1100 kg	78000 kg	230000 kg
Rolling resistance	162 N	11.48 kN	33.85 kN
Engine	O-320	CFM56-5A1	PW4168
Number of engines	1	2	2
Maximal fuel flow	1×0.0112 kg/s	2×1.051 kg/s	2×2.884 kg/s
Rated output	unknown	2×112 kN	2×302.5 kN

The same approach as in Chen and Stewart (2011) was used to calculate the total rolling resistance and is defined as follows:

$$F_r = \mu \cdot m \cdot g, \tag{3.1}$$

where $\mu = 0.015$ is the rolling resistance coefficient on a concrete surface, m is the maximum take-off weight of the aircraft and $g = 9.81$ m/s^2 is the acceleration due to gravity. The maximal fuel flow and the rated output values are based on the International Civil Aviation Organization engine emissions database (ICAO 2008) and the research by Stettler et al. (2011). They have to be multiplied by the number of engines specified in each setting.

4

A Statistical Approach for Taxi Time Estimation

> Essentially, all models are wrong, but some are useful.
>
> *Empirical Model-Building and Response Surfaces*
> George E. P. Box

4.1 Introduction

Airlines and airports face several key challenges in the near future. Firstly, the number of flights is predicted to increase in the next few years (SESAR 2006). Secondly, there is an increasing focus upon environmental considerations, and this is likely to increase in importance. Thirdly, the use of computerised tools is enabling increased aircraft utilisation, reduced idle times, and increased passenger connection options, leading to ever more complex and interlinked flight schedules. The on-time performance of flights at each airport and the earlier visibility of any delays (allowing corrective measures to be put into place) is becoming increasingly important, since many downstream flights can be affected by delays to each single aircraft. Consequently, the operations at busy hub airports are experiencing an increased focus of attention, and this

is likely to increase in the face of future challenges.

Total taxi times from stand/gate to runway are needed if advanced predictions of take-off times are required, for use by en-route controllers (or decision support systems to help them) or for improving arrival time predictions for the destination airports, allowing the effects of any predicted delays to be mitigated. Taxi times are already needed by several existing search algorithms for take-off time prediction and take-off sequencing (Atkin et al. 2007; Eurocontrol 2012) and for allocating appropriate stand holds to aircraft to absorb ground delay at the gate/stand, decreasing the fuel burn and environmental effects (Burgain et al. 2009; Atkin et al. 2010a). Although the effects have been less well studied and perhaps being less sensitive to small prediction errors, taxi times are also useful for arrivals, being necessary for predicting stand/gate arrival times, to ensure that adequate resources are available at the correct time (Eurocontrol 2012). Taxi time predictions will become even more important if the efficiency of stand resource utilisation is to be improved in the future. The current common practice is to use standard mean taxi times for each taxi source/destination pairs. A better understanding of the influencing factors, and a model to estimate such taxi times to a higher level of accuracy, would have positive effects for both the published approaches and the systems which are currently in use.

The importance of the ground movement problem was explained in Chapter 2, highlighting how it links several other airport operations such as runway sequencing and gate assignment. Improved ground movement can increase on-time performance at airports, so ground movement simulations and optimisers are extremely useful. These usually explicitly model the interactions between aircraft (modelling delays due to other aircraft and any necessary re-routing on longer paths to avoid conflicts) and, thus, require predictions for taxi times which do not already include these elements (Gotteland and Durand 2003; Smeltink et al. 2004; Balakrishnan and Jung 2007; Roling and Visser 2008; Lesire 2010). The use of historic data would be preferable for calibrating models. However, such recorded data usually includes significant delays due to the interactions between aircraft. There are obvious benefits from being able to quantify the effects of this interaction and the model which is considered in this chapter aims to provide this facility. Although average speeds have often had to be used in the past due to the lack of

reliable predictions, it is important to understand aircraft speed in more detail if more realistic ground movement decision support systems are desired.

The causes and effects of taxi time variability are both often neglected in academic literature. However, some elements have been considered in the past. Rappaport et al. (2009) analysed the effect on taxi times of having to reduce speed for turns and it was shown that aircraft travelling straight forward reached higher average speeds than those with upcoming turns. Atkins et al. (2008) also suggested remarkable variabilities of taxi times around a corner during taxiing. In addition, Idris et al. (2002) performed a statistical analysis of departing aircraft at Boston Logan International Airport (BOS) with the conclusion that the taxi-out time for each airline/runway configuration combination was highly dependent upon the take-off queue size. However, the analysis by Idris et al. (2002) only covered taxi times for departing aircraft. The problem also seems to differ between North American and European airports, with much shorter take-off queues usually being observed at European hubs. More recently, two further estimation approaches were published for North American airports. Simaiakis and Balakrishnan (2009) presented a queuing model and considered the potential impact on emissions reduction. The statistical analysis exclusively used the size of the take-off queue to estimate the taxi-out time. Balakrishna et al. (2009) presented a model for taxi-out time prediction based on reinforcement learning algorithms. Recently, Srivastava (2011) published work on departure taxi time prediction using ASDE-X surveillance data in a linear regression model. The explanatory variables were aircraft queue position, distance to the runway, arrival rates, departure rates and weather and were evaluated at John F. Kennedy International Airport (JFK).

In other work, Tu et al. (2008) analysed push-back delays at Denver International Airport with seasonal trends and daily propagation patterns. Gate-waiting is defined as the phenomenon when an arrival has to wait until a gate becomes available, e.g. when the gate was blocked by another aircraft (Idris 2001; Wang et al. 2009a; Wang 2011). Wang et al. (2009b) showed that 10 major US airports were affected regularly. However, this seems to be less of an issue for the European airports which were analysed in this study. A recent study by Carpenter and Stroiney (2012) showed the potential of managing ramp congestions. Ramps are areas around terminals, especially at US airports, which are mostly managed by an airline, whereas the rest

of the taxiways are usually controlled by another authority. Such a division is not the normal case at airports in Europe.

The aim of this chapter is to extensively study the variation of taxi times not only for departing aircraft, but also for arriving aircraft. In contrast to earlier studies, we focus on European hub airports in this chapter where the taxi process is less dominated by queuing at the runway and hence other factors have a proportionately greater effect upon taxi times. The consideration of the airport layout is essential for this research and was not considered in the past. The outcomes will enable researchers to make increasingly accurate taxi time predictions and to develop more realistic decision support systems for ground movement, potentially resulting in smoother airport operations, emission reductions for the taxi process and better on-time performance at airports.

The remainder of this chapter is structured as follows: Section 4.2 provides a description of the problem and the available data. The statistical taxi time prediction method is then detailed in Section 4.3, where the influence of the ground movement model will be observed. The results and their applications are discussed in Sections 4.4, 4.5 and 4.6. The chapter ends by drawing important conclusions from this work in Section 4.7.

4.2 Problem Description

The problem considered in this chapter involves the identification of a function to estimate taxi times for both arriving and departing aircraft, which can then be used in an airport decision support system. The problem description in this section has two parts. Firstly, we summarise the airport ground movement problem, explaining why accurate taxi times are very important. Secondly, we discuss the data which we can expect to be available from an airport for use in calibrating ground speed models.

4.2.1 The Airport Ground Movement Problem

This research was motivated by our work on the airport ground movement problem (Atkin et al. 2011b; Ravizza and Atkin 2011), which is basically a routing and scheduling problem (see

Chapter 6). It involves directing aircraft on the surface of an airport to their destinations in a timely manner, with the aim usually being to reduce the overall travel time, to meet some target time windows and/or to absorb the delay at a preferred time, such as when the engines are not running. It is crucial, for reasons of safety that two aircraft never conflict with each other throughout the ground movement.

For larger airports, especially during peak hours, decision support systems are advantageous to deal with the complexity of the problem (Gotteland and Durand 2003; Smeltink et al. 2004; Balakrishnan and Jung 2007; Roling and Visser 2008; Lesire 2010). Sophisticated algorithms are needed to route and schedule all the aircraft simultaneously on the surface. In doing so, some aircraft might be allocated to a longer route and/or waiting times might need to be added to some schedules to handle conflicts, aiming for a globally better solution.

For the purpose of this chapter, the important feature of this problem is that decision support systems need taxi time predictions for aircraft in isolation, ignoring the presence of other aircraft, but historic data is rarely able to provide this information. However, it is clear that the use of historic data is vital in order to ensure that results are realistic and can be compared to the status quo at an airport, in order to quantify any potential improvements from new airport ground movement decision support systems, without running expensive trials.

4.2.2 Utilised Airport Data

This analysis utilised data from two hub airports in Europe: Stockholm-Arlanda Airport and Zurich Airport. Sketches of the two airport layouts are provided in Figure 3.5 and Figure 3.1 in Chapter 3. We utilised data for an entire day's operation at each airport and used the datasets "ARN 1" (661 movements) and "ZRH 2007" (679 movements). Both datasets represent days with no extraordinary occurrences to be taken into account.

In visually analysing the average taxi speeds, it was obvious that there were major differences between different groups of aircraft. A boxplot is presented in Figure 4.1, showing the general variability in the average speed of the aircraft for two stand groups at Stockholm-Arlanda Airport (the average speed was calculated based upon the taxi time and the shortest path).

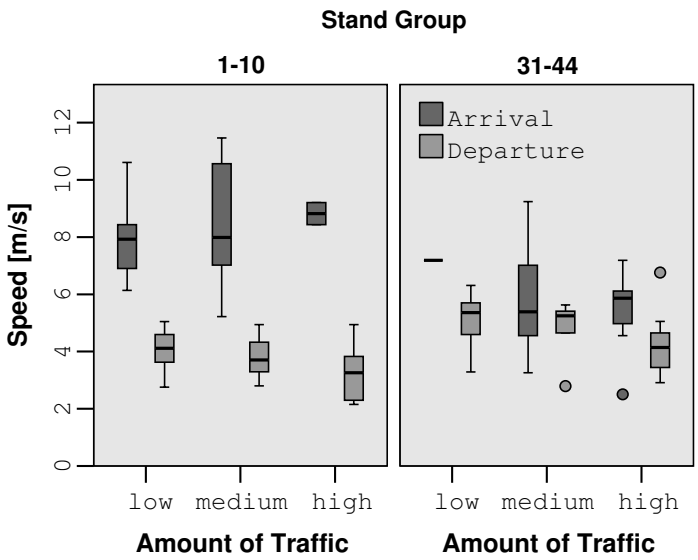

Figure 4.1: Average speed at Stockholm-Arlanda Airport from two different stand groups to the runway 19R

Major differences are apparent between arriving and departing aircraft as well as between low, medium and high traffic situations at the airport.

4.3 Approach for Estimating Taxi Speed

The aim of this research is to estimate a function which can more accurately predict taxi times for aircraft or, equivalently, better predict their average speeds. It is not obvious which factors are important for calculating such taxi times and which factors can be ignored. Discussions with practitioners can help in understanding the problem and identifying potential factors but this has its limits for mathematically determining the importance of factors. Multiple linear regression was able not only to answer this question, but also to estimate a function which could predict the taxi speed and was easy to interpret. Of course, the accuracy of the estimation has to be verified, but given such a function, the aim is to eliminate the effects of factors which represent the actual amount of traffic at the airport, by setting the respective variables to 0. Our aim is to be able to predict the taxi times for independent aircraft, for use in a more advanced ground movement decision support system. This would provide the opportunity to

compare scenarios with the way in which an airport is currently operating.

4.3.1 Summary of Multiple Linear Regression

A brief summary of multiple linear regression is given here for reasons of completeness, before providing the details of how it has been applied to the problem of estimating taxi times by incorporating details of the airport layout. The interested reader is directed to the book by Montgomery et al. (2001) for a more in-depth presentation of multiple linear regression.

Multiple linear regression is a statistical approach which attempts to model a dependent variable y as a function of other explanatory variables $x_1, \ldots, x_p$, by a function of the following form:

$$y_i = \beta_0 + \beta_1 x_{i1} + \cdots + \beta_p x_{ip} + \epsilon_i, \quad i = 1, \ldots, n, \tag{4.1}$$

where $\beta_1, \ldots, \beta_p$ are the true (but unknown) coefficients of the regression, x_{ij} is the ith value of the jth explanatory variable, y_i is the ith value of the response variable, and ϵ_i is the ith value of a random error term. The random error terms $\epsilon_1, \ldots, \epsilon_n$ are assumed to be uncorrelated and to have a normal distribution with mean zero and constant variance σ^2. The regression coefficients can be estimated using least squares regression, yielding estimated coefficients $\hat{\beta}_1, \ldots, \hat{\beta}_p$. The predicted y value for the ith observation is then given by

$$\hat{y}_i = \hat{\beta}_0 + \hat{\beta}_1 x_{i1} + \cdots + \hat{\beta}_p x_{ip}. \tag{4.2}$$

The difference between y_i and $\hat{y}_i$ is called the residual, e_i.

A measure is needed to analyse the accuracy of the model. We used the coefficient of determination R^2 to perform this function, which is defined as follows:

$$R^2 = 1 - \frac{SS_{Res}}{SS_T}, \tag{4.3}$$

where SS_{Res} is the residual sum of squares, i.e. $SS_{Res} = \sum_i (y_i - \hat{y}_i)^2$, and SS_T is the total sum of squares, i.e. $SS_T = \sum_i (y_i - \bar{y})^2$, where $\bar{y}$ is the mean of y. Since models with fewer explanatory variables should be preferred over models using many explanatory variables that

fit equally well, an adjusted R^2 value is often used, including an additional factor to count the number of observations n and the number of explanatory variables p:

$$R^2_{Adj} = 1 - \frac{SS_{Res}/(n-p-1)}{SS_T/(n-1)}. \tag{4.4}$$

Both R^2 and R^2_{Adj} are values between 0 and 1 for regression models, with values closer to 1 representing better models.

Multiple linear regression is widely used to describe data, to understand the correlation between variables, to forecast similar observations and to estimate parameters. Regression models work well within their range of observed data, but they can be very poor for forecasting events outside of this range.

4.3.2 Analysis of the Dependent Variable

We discovered that estimations of taxi speeds (in m/s) better fit the requirements of the linear regression models than direct estimates of the taxi times of aircraft. In particular, the assumptions that the statistical errors are normally distributed were not fulfilled with the direct estimates. Furthermore, we also discovered that a logarithmic transformation of the dependent variable (Equation (4.5)) was required in order to fulfil the stated assumptions of multiple linear regression. Therefore this transformation is used throughout the following sections:

$$y := \log_{10}(Speed). \tag{4.5}$$

A good estimate for $\log_{10}(Speed)$ can then be used for the calculation of a good estimate of the taxi time.

4.3.3 Analysis with only one Explanatory Factor

Different individual factors are analysed in this section. The analysed factors were derived from a combination of previously published work in this area, discussions with practitioners and data-driven transformations. The factors which appeared to be statistically relevant were

then included together in a combined model. For reasons of simplicity, we focus within this section only on the settings for Stockholm-Arlanda Airport, although many results are similar for both airports, as can be observed in Section 4.3.4.

4.3.3.1 Distances

The first factor which was analysed considered the distance (in metres) that an aircraft was taxiing. To determine such distances, it was useful to model the airport ground layout as a graph, where the edges represented the taxiways and the vertices represented the junctions or intermediate points (see Figure 3.6). Based on this underlying graph, it was then assumed that aircraft were travelling on their shortest path and Dijkstra's algorithm (see Cormen et al. (2001) for more details) was used to determine, for each aircraft, the taxi distance from the stand to the runway or back again. The incorporation of the actual airport layout was essential for the approach as will be seen later. We note that further improvements may be possible from using the actual route taken, but that information was not available at the time.

Regressing $\log_{10}(Speed)$ on 'Distance' yielded an adjusted coefficient of determination $R^2_{Adj} = 0.473$, with a p-value smaller than 2.2e-16 (the p-value comes from the F-test that compares the given model to a model with only an intercept). Figure 4.2(a) shows a plot of the observed values, y, against the explanatory variables, x.

Analysis of the results in Figure 4.2(a) encouraged the application of a logarithmic transformation to the distance. The visualisation of the fit can be seen in Figure 4.2(b), where the fit has a better linear shape. Regressing $\log_{10}(Speed)$ on $\log_{10}(Distance)$ yielded an R^2_{Adj} value of 0.479 (p-value $<$ 2.2e-16), which is only marginally better, but it will be observed later that it leads to significant improvements in the final model for both airports.

The R^2_{Adj} value indicates that almost half of the variance can be explained by this factor, showing the importance of this indicator. Therefore, additional time was invested in analysing it. Instead of only using the entire distance of an aircraft as a variable, it was divided into three different components based upon the known behaviour of aircraft as they taxi around the airport. 'Distance0' represented the length of the path directly around the gates, 'Distance2' represented the length of the path which was comprised of long sub paths without any junctions and 'Dis-

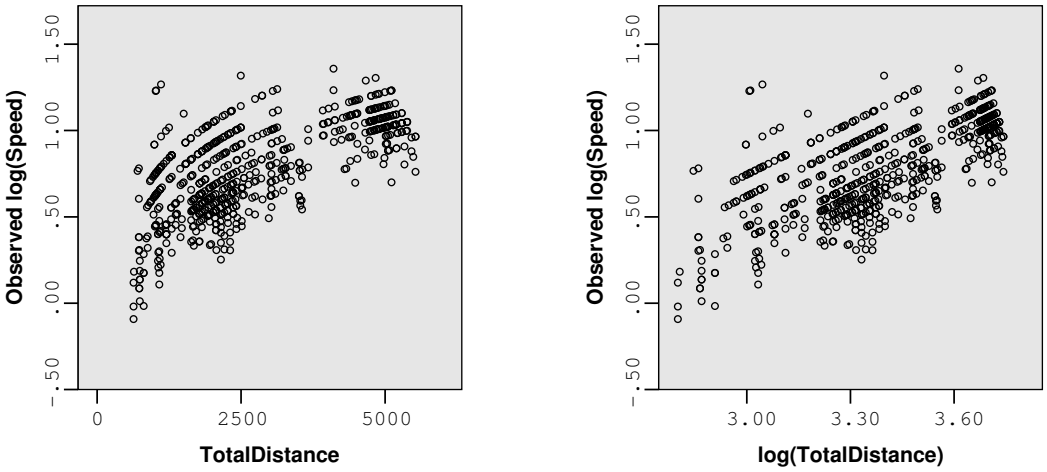

(a) Distance

(b) $\log_{10}$(Distance)

Figure 4.2: Scatterplots showing the logarithmic transformation

tance1' represented the remaining distance (where all values were in metres). These distances were determined using the directed graph model of Stockholm-Arlanda Airport, by assigning each edge in the graph to one of the three distances. The 'Distance', 'Distance0', 'Distance1', 'Distance2', $\log_{10}(Distance)$, $\log_{10}(Distance0)$, $\log_{10}(Distance1)$, and $\log_{10}(Distance2)$ values were all included in the analysis. The resulting regression model yielded an improved R^2_{Adj} value of 0.604 (p-value $<$ 2.2e-16).

4.3.3.2 Angle

The total amount of turning which an aircraft had to achieve was another promising predictor of taxi speed, since aircraft have to slow down to make turns (Gong 2009). A factor was introduced to measure the total turning angle (in degrees), calculated as the total angular deviations between adjacent edges on the shortest path for the aircraft. Again, the graph model of the airport layout was used for this, as shown in Figure 4.3. This turned out to be another major factor ($R^2_{Adj} = 0.470$, p-value $<$ 2.2e-16) and the importance was improved further when $\log_{10}(Angle)$ was considered ($R^2_{Adj} = 0.482$, p-value $<$ 2.2e-16).

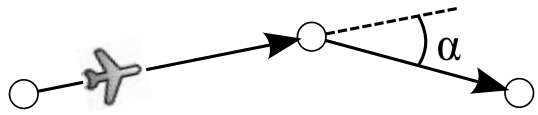

Figure 4.3: Measuring turning angle of aircraft on one vertex

4.3.3.3 Departures vs. arrivals

As shown in Figure 4.1, the speed for departures can differ significantly from the speed for arrivals and factors were introduced for this. In contrast to the factors which have been introduced so far, this information is nominal rather than being a continuous variable. A dummy variable called ARR was introduced, defined to be 1 for arrival aircraft and 0 for departure aircraft. The regression showed an R^2_{Adj} value of 0.380 for this single factor, demonstrating its importance (p-value $<$ 2.2e-16).

4.3.3.4 Amount of traffic

Another important factor affecting the taxi speed of aircraft is the amount of traffic on the airport surface while the aircraft is taxiing. As a first attempt for an indicator of surface load, we divided the operational hours into three different categories. The indicator 'Traffic_high' was set to 1 for hours where more than 50 aircraft were moving and to 0 otherwise. 'Traffic_medium' was set to be 1 for hours with between 36 and 50 moving aircraft and 0 otherwise. Both indicators were set to zero for the last category representing low surface load (the same categorisation is visible in Figure 4.1). This approach with these variables resulted in an R^2_{Adj} value of only 0.007 and a p-value of 0.036.

A more advanced measure was introduced based upon the paper by Idris et al. (2002). The value N_i counts the number of other aircraft which are taxiing on the airport surface at the time that the particular aircraft i started to taxi, as shown in Equation (4.6), where the Iverson bracket denotes the value 1 if the condition in square brackets is satisfied and is 0 otherwise. The parameters t^i_{start} and t^i_{end} represent the time at which aircraft i starts and ends its taxi operation.

$$N_i = \sum_{j \in Aircraft \setminus \{i\}} \left[t^i_{start} \in (t^j_{start}, t^j_{end}) \right] \tag{4.6}$$

The value Q_i was also adopted to count the number of other aircraft which cease taxiing during the time aircraft i is taxiing, as shown in Equation (4.7), again using the Iverson bracket.

$$Q_i = \sum_{j \in Aircraft \setminus \{i\}} \left[t^j_{end} \in (t^i_{start}, t^i_{end}) \right] \tag{4.7}$$

Since the paper by Idris et al. (2002) was restricted to taxi-out times, this approach was further developed to cope with separate departures and arrivals. Eight integer variables were used to allow consideration of the effects of the counts of arrivals and departures depending upon whether the current aircraft was an arrival or departure. These were named $N_{DEP,\#DEP}$, $N_{DEP,\#ARR}$, $N_{ARR,\#DEP}$, $N_{ARR,\#ARR}$, $Q_{DEP,\#DEP}$, $Q_{DEP,\#ARR}$, $Q_{ARR,\#DEP}$ and $Q_{ARR,\#ARR}$. In this notation, the N or Q indicated whether it was the count of already moving aircraft or of aircraft which ceased their movement. The first index for each value represented the type of aircraft under consideration (ARRival or DEParture). The second index indicated whether it was the count of arrivals or departures ($\#ARR$ or $\#DEP$) which was to be considered for counting, i.e. for a departing aircraft, all of the variables with a first index ARR are treated as if they are 0 and for arriving aircraft all of the variables with the first index of DEP are treated as if they are 0.

A highly significant regression model considering only these eight factors led to an R^2_{Adj} value of 0.422 (p-value $<$ 2.2e-16). Further investigation was performed to determine whether the model could be further improved by considering only aircraft destined for, or originating from, the same runway as the aircraft under consideration. In that case, the fit was worse ($R^2_{Adj} =$ 0.382, p-value $<$ 2.2e-16). One possible explanation for this is that often one runway is used for departures and another one for arrivals, in which case half of the factors have the same value as in the unrestricted case and the other half have the value 0, resulting in less information being considered by the model than in the unrestricted case.

4.3.3.5 Less important factors

A number of other elements were taken into consideration, for example whether the model could be improved by using the square of some of the values or by including some interaction terms but no improvement was found. Another approach was to consider the number of engines of the aircraft (R^2_{Adj} = 0.007, p-value = 0.039) or by using the wake vortex categorisation of the aircraft (R^2_{Adj} = 0.032, p-value = 4.4e-05). These results for the European airports which we studied fit the findings of Idris et al. (2002) (for a North American airport), where a poor correlation was observed between taxi time and aircraft type, and the type determines both the number of engines and the wake vortex categorisation.

Further analysis studied the effect of the different runways and stand groups. Although nothing relevant was found for Stockholm-Arlanda Airport, some effects were found at Zurich Airport by analysing different operational modes (which runway(s) is/are being used for take-offs/landings). The details are reported later, in the analysis of the whole model for Zurich Airport.

4.3.4 Multiple Regression with Several Factors

This section presents multiple regression models for Stockholm-Arlanda Airport and Zurich Airport and ends with a consideration of the validity of the necessary assumptions to apply the regression. The applicability of the model and the discussion of the results can be found in later sections.

The goal of the multiple regression approach was to find the most important factors explaining the variability of the real datasets.

Extensive analysis was performed using different stepwise selection methods based on the factors described in Section 4.3.3 (depending on p-values, Akaike's Information Criterion (AIC) and the Bayesian Information Criterion (BIC)). We decided to present models which are as practical as possible for use at airports (requiring less information) and which are easy to interpret. The following models fulfil this aim and are less than 2.2% away from the best found models (according to the R^2_{Adj} value).

4.3.4.1 Stockholm-Arlanda Airport

The final regression model for Stockholm-Arlanda Airport is given in Table 4.1. The first column indicates the variables, the second column represents the estimated unstandardised coefficients and the third column presents the corresponding estimated standard errors. The fourth column shows the estimated standardised coefficients for all non-dummy variables (i.e. the estimated coefficients if the variables were standardised so that their variance was 1). This measure can be used to analyse which factor has the largest positive or negative impact on $\log_{10}(Speed)$. In contrast to the unstandardised coefficients, they have no units and can therefore be compared directly. The last column shows the significance of each variable based on a t-test.

Table 4.1: Coefficients for Stockholm-Arlanda Airport, Sig. indicates if the p-value is < 0.05 (*), < 0.01 (**) or < 0.001 (***)

	Coefficient $\hat{\beta}_i$	Std. Dev.	Standardised Coefficient	Sig.
(Constant)	-2.349	0.091		***
$\log_{10}$(Distance)	0.922	0.029	0.842	***
ARR	0.211	0.015		***
$N_{DEP,\#DEP}$	0.031	0.005	0.260	***
$N_{ARR,\#DEP}$	0.029	0.004	0.230	***
$N_{ARR,\#ARR}$	0.049	0.011	0.209	***
$N_{DEP,\#ARR}$	0.036	0.006	0.176	***
Distance2	-5e-05	8e-06	-0.188	***
$Q_{DEP,\#ARR}$	-0.034	0.004	-0.268	***
$Q_{ARR,\#ARR}$	-0.066	0.011	-0.279	***
$Q_{ARR,\#DEP}$	-0.052	0.006	-0.280	***
$Q_{DEP,\#DEP}$	-0.044	0.005	-0.397	***

The model has a good R^2_{Adj} value of 0.863 (p-value $<$ 2.2e-16). This means that around 86% of the variance of the $\log_{10}$(Speed) values can be explained by the model. The fit of the prediction can be seen in Figure 4.4.

4.3.4.2 Zurich Airport

As indicated in Section 4.3.3.5, a significant factor at Zurich Airport is the current operational mode of the runways. As long as no heavy winds occur, Zurich Airport operates strictly with three operational modes (see Section 3.2.1). We modelled the three operational modes using two dummy variables, $O_{Morning}$ to represent the morning period and $O_{Evening}$ to represent the

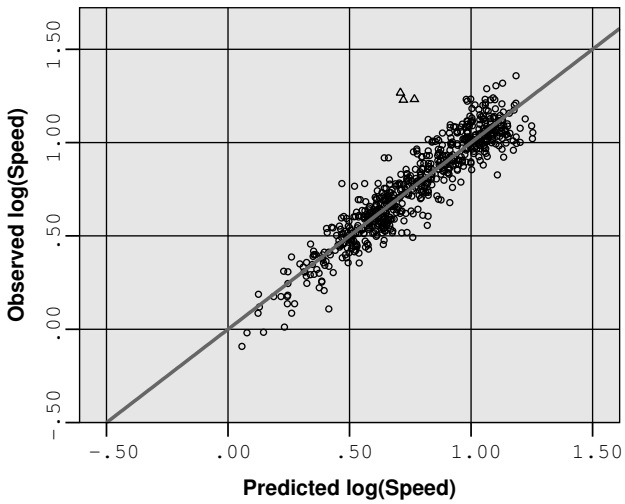

Figure 4.4: Scatterplot showing the linear fit of the regression model in Table 4.1 for Stockholm-Arlanda Airport

evening period. Each variable was set to 1 during the corresponding period and 0 otherwise, so during the day period both variables were set to 0.

In contrast to Stockholm-Arlanda Airport, statistical analysis showed only small improvements by classifying the total distances into different components, so they were excluded from the final model. This was expected from the airport layout since it has fewer straight sub paths without junctions.

The fit for Zurich Airport is given in Table 4.2, and shows an even better fit than for Stockholm-Arlanda Airport, with an R^2_{Adj} value of 0.878 (p-value $<$ 2.2e-16). The scatterplot of the relationship between the observed values and the predicted values can be seen in Figure 4.5.

4.3.4.3 Validation of statistical assumptions

The estimated regression coefficients are unbiased if $E(\epsilon_i) = 0$ for all $i = 1, \ldots, n$. The residual plots in Figure 4.6 indicate that this assumption is approximately valid (with perhaps a slight lack of fit for small speeds). Hence, one can be confident that the estimated regression coefficients and resulting predictions are (almost) unbiased.

The standard errors for the estimated coefficients are valid if the following three assumptions hold: $E(\epsilon_i) = 0$ and $Var(\epsilon_i) = \sigma^2$ for all $i = 1, \ldots, n$, and $Cov(\epsilon_i, \epsilon_j) = 0$ for all $i \neq j$. The

Table 4.2: Coefficients for Zurich Airport, Sig. indicates if the p-value is < 0.05 (*), < 0.01 (**) or < 0.001 (***)

	Coefficient $\hat{\beta}_i$	Std. Dev.	Standardised Coefficient	Sig.
(Constant)	-2.601	0.250		***
$\log_{10}$(Distance)	1.161	0.091	0.731	***
ARR	0.260	0.018		***
$N_{DEP,\#DEP}$	0.025	0.005	0.234	***
$N_{ARR,\#ARR}$	0.054	0.008	0.208	***
$N_{ARR,\#DEP}$	0.019	0.004	0.143	***
$N_{DEP,\#ARR}$	0.029	0.007	0.101	***
$O_{Evening}$	0.049	0.013		***
$O_{Morning}$	-0.075	0.019		***
$\log_{10}$(Angle)	-0.143	0.039	-0.083	***
Distance	-7e-05	2e-05	-0.181	**
$Q_{DEP,\#ARR}$	-0.032	0.004	-0.208	***
$Q_{ARR,\#DEP}$	-0.067	0.006	-0.285	***
$Q_{ARR,\#ARR}$	-0.081	0.007	-0.318	***
$Q_{DEP,\#DEP}$	-0.046	0.004	-0.466	***

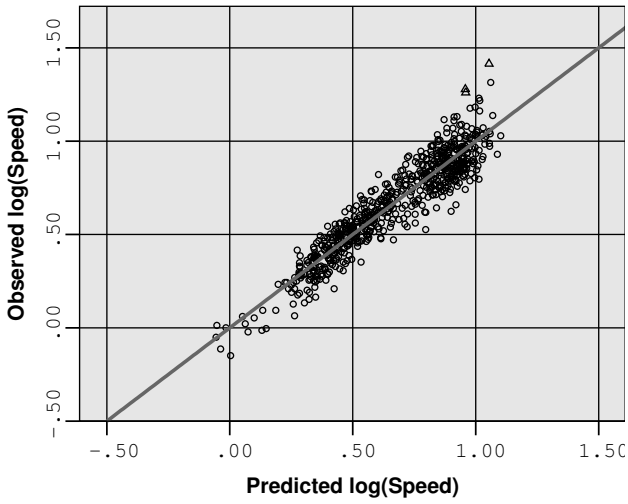

Figure 4.5: Scatterplot showing the linear fit of the regression model in Table 4.2 for Zurich Airport

residual plot in Figure 4.6(a) indicates that the constant variance assumption is approximately valid for Stockholm-Arlanda Airport. For Zurich Airport, there seems to be some increase in the variance with increasing predicted speeds. Due to the time dependent nature of the data, it is likely that there is some correlation in the statistical errors. The Durbin-Watson test (Durbin and Watson 1950, 1951) indicated positive serial correlation for both airports. Generalised least squares models using autoregressive AR(1) and AR(2) models (Fox 2002; Venables and Ripley

2002) for the residuals were fitted to account for this correlation, and the results were compared to Tables 4.1 and 4.2. Estimates of the coefficients and standard errors at both airports are very consistent (see Appendix A).

Finally, the p-values are valid if, in addition to the assumptions above, the statistical errors have a normal distribution. Moreover, even without the normality assumption they hold approximately if the sample size is sufficiently large, due to the central limit theorem. The Q-Q-plots in Figure 4.7 show that the residuals are approximately normally distributed. A discussion about the outliers (indicated with triangles) is presented in Section 4.4.2. Formal Shapiro-Wilk tests (Shapiro and Wilk 1965) were also performed to test the normality assumption, where the outliers were excluded. These tests supported the findings from the figures and indicated no evidence for departure from normality (p-values 0.083 and 0.463 for Stockholm-Arlanda Airport and Zurich Airport, respectively). However, due to potential (small) violations of the assumptions of constant variance and uncorrelated errors, the p-values for Zurich Airport might be slightly off.

The taxi distance appears on both sides of the multiple linear regression models, due to the decision to use speed as the dependent variable. However, since it seems clear that distance might influence speed but not the other way around, we assume that there are no endogeneity problems.

4.3.5 Cross-validation

A common way of testing how well a model performs in predicting new data is the so called PRESS statistic, suggested by Allen (1971):

$$\text{PRESS} = \sum_{i=1}^{n} (y_i - \hat{y}_{(i)})^2. \tag{4.8}$$

It sums the squared differences between the observed variables y_i and the predicted variables $\hat{y}_{(i)}$ for each of the sample points i, where the prediction $\hat{y}_{(i)}$ only uses the data of the remaining

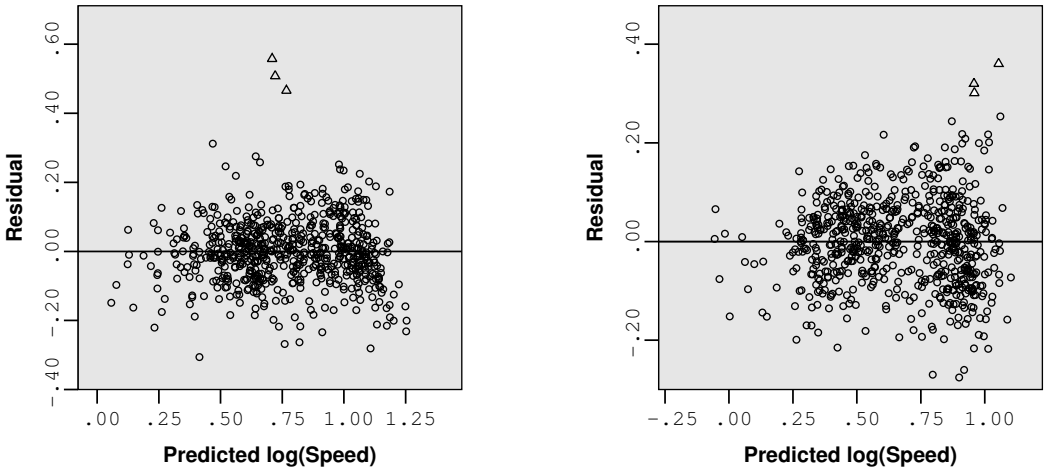

(a) Stockholm-Arlanda Airport (b) Zurich Airport

Figure 4.6: Residual plots showing the validation of the assumptions

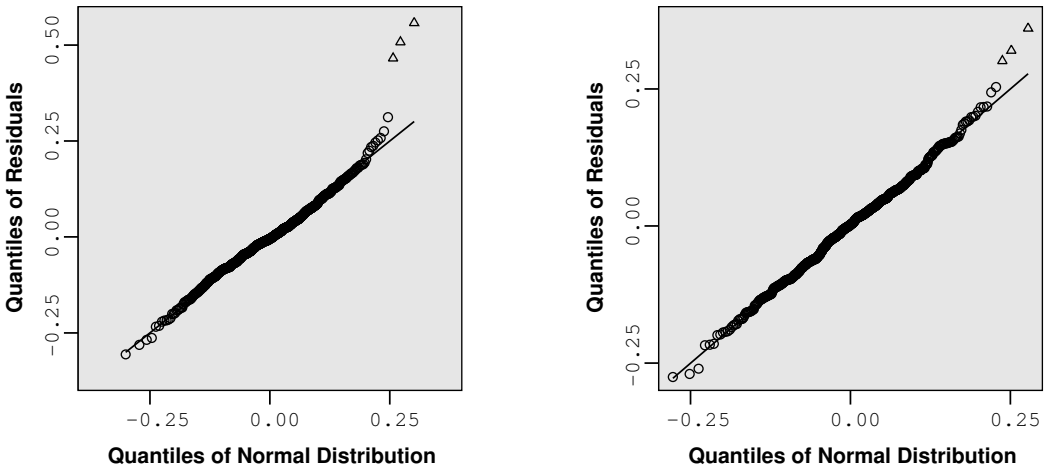

(a) Stockholm-Arlanda Airport (b) Zurich Airport

Figure 4.7: Normal Q-Q-plots showing the validation of the assumptions

observations. It can be categorised as a leave-one-out cross-validation. The PRESS statistic can be used to calculate an R^2 value for a prediction:

$$R^2_{Pred} = 1 - \frac{\sum_{i=1}^{n} (y_i - \hat{y}_{(i)})^2}{\sum_{i=1}^{n} (y_i - \overline{y})^2}. \tag{4.9}$$

The R^2_{Pred} value was 0.860 for Stockholm-Arlanda Airport and 0.875 for Zurich Airport. This means that, for similar settings at the airport (the same operational modes, similar weather conditions and so on), these models could explain around 86% and 87.5%, respectively, of the variability in predicting new observations due to the combination of the statistical analysis with the incorporation of the ground layout model.

4.3.6 Prediction Accuracy

A second dataset was made available for Zurich Airport after the model had been fitted to the existing dataset. The second dataset consisted of 5613 aircraft movements which occurred during one week's operation (dataset "ZRH 2011" from Section 3.2.1). Even though we used the same coefficients as reported in Table 4.2, and they were generated using the old data (dataset "ZRH 2007"), the approach was still able to demonstrate a high R^2_{Adj} value of 0.864 for the prediction. Keeping the same factors as in Table 4.2, but re-estimating the coefficients for the new dataset, the R^2_{Adj} could only be improved to 0.899. These results demonstrate that the model was not only able to fit historic data well but that it can also be used to make accurate taxi speed predictions, especially when keeping in mind that the two datasets were from periods which were almost four years apart.

4.4 Interpretation of the Models

First of all, it can be seen from Tables 4.1 and 4.2 that the two fitted regression models are very similar and have the same general structure, indicating their potential usage for other airports. All the factors in the tables are highly significant (p-value < 0.01).

4.4.1 Coefficient Meanings

We now interpret some of the coefficients, to gain insight into the effects of specific factors. The straightforward interpretation of this model could possibly encourage airport operators to use this approach to support their needs.

4.4.1.1 Distances

The most important factor for both airports was the logarithmic transformation of the total distance. In general, the average taxi speed was higher the further an aircraft had to taxi. This finding is new compared to the results from other research, where the focus was upon airports with longer queues, which probably dominated the effect of the distance. Even with the assumption of using the shortest path for each aircraft, the results look promising and would probably look even better by utilising the actual distance rather than the shortest path.

4.4.1.2 Departures vs. arrivals

Another important factor in the models for both airports was the differentiation between arriving and departing aircraft. Since departures often need to wait in a queue, their average speed is smaller in comparison to arriving aircraft, which are forced to clear the runway as soon as possible and taxi directly to the stands.

4.4.1.3 Angle

The logarithmic transformation of the total turning angle which an aircraft had to complete was observed to be a significant slowing factor at Zurich Airport. The inclusion of this factor significantly improved the accuracy of the prediction.

4.4.1.4 Amount of traffic

All of the different Q values were observed to have a negative effect upon the taxi speed. In general, more aircraft travelling around the airport means that each individual aircraft's speed is reduced. Factors which particularly slowed taxi speeds were $Q_{DEP,\#DEP}$ and $Q_{ARR,\#ARR}$, representing the number of aircraft which have the same target (runways or stands) but end their taxi operation first. The N variables were found to counteract some of the effect of the Q variables, together modelling those aircraft which both start to taxi earlier and which reach their destination earlier. Our results showed differences between the North American airport studied by Idris et al. (2002) and the European airports considered in this research, since the

number of arrivals did not affect the taxi out time in their study whereas there was a strong correlation in our analysis. This may be related to the airport layouts or the runway queue lengths.

4.4.1.5 Operational mode

In the case of Zurich Airport, the influence of the different operational runway modes was incorporated into the model. It can be observed that aircraft taxi faster in the evening than during the day, and faster during the day than in the morning. There is insufficient information at the moment to determine whether the effect is due to the different runway modes or whether other elements such as visibility or different aircraft mixes at different times of the day are affecting the taxi speeds.

4.4.2 Unexplained Variability

Around 13% of the variability in taxi speeds cannot be explained by our models. Some potential explanations are listed below:

- The taxi behaviour can vary between different airlines and pilots. Additional data should allow this to be analysed in more detail in the future.
- In the case of Stockholm-Arlanda Airport the taxi time information was only to the minute rather than to the second, but the model uses continuous time for the speed predictions. The data of Zurich Airport had detailed times at the runway, but again the times at the stands/gates were only to the minute. This matching of continuous time to discrete values is unlikely to provide extremely accurate predictions.
- We assumed that aircraft travelled along the shortest path and that there were no unexpected changes. This assumption will be valid in general but can lead to occasional errors.

An analysis of the outliers at Stockholm-Arlanda Airport showed that the three worst fits (the three triangles in Figures 4.4, 4.6(a) and 4.7(a)) were for aircraft landing at runway 26 and

taxiing to pier F. The indicated taxi times in the data were 1 minute for one of the aircraft and 2 minutes for the other two - showing extremely short taxi times. Given the minute granularity on the data, it is perhaps unsurprising that the estimations were least accurate for these aircraft. Removing these three aircraft from consideration resulted in an improvement to R^2_{Adj} of about 0.01. Similarly, the most extreme outliers at Zurich Airport (the three triangles in Figures 4.5, 4.6(b) and 4.7(b)) were also related to very short taxi times.

4.5 Applicability of this Research

The two main applications for this research are for total taxi time prediction and for use in a ground movement decision support system. We consider both of these in this section.

4.5.1 Improved Total Taxi Time Prediction

To the best of our knowledge, there is no existing taxi time prediction function to compare against for both departing and arriving aircraft, but we have the lookup table which is used for Zurich Airport. This considers only the sources and destinations and gives average taxi-in and taxi-out times. However, it has a granularity of one minute and deliberately underestimates times. In order to eliminate the deliberate underestimates, we used linear regression to find a linear scaling which best fitted their table to the observed data. This resulted in an improved R^2_{Adj} value of 0.180, with a scaling of $ax + b$, where a is 0.883 and b is 2.210. In contrast, the approach presented in this chapter, when applied to taxi times (rather than $\log_{10}(Speed)$) resulted in an R^2_{Adj} value of 0.793, thus explaining the variability in taxi times at this airport to a much greater extent than the lookup table and indicating the benefits of the consideration of more factors. The function generated by our multiple linear regression is, therefore, more appropriate for predicting total taxi time.

The results were also compared to the results from the application of a reinforcement learning algorithm by Balakrishna et al. (2009) at other airports. They presented results for the $\pm$ 3 or $\pm$ 5 minute prediction accuracy for the taxi-out times (see Table 4.3), measuring the percentage of departing aircraft with a time difference between the predicted time and the observed time

Table 4.3: Comparison of prediction accuracy; The first block shows the result for the two airports which were studied where the prediction model is simplified by not considering the airport layout (particularly not considering the factors about the distances and the turning angles). The results for the best found models are indicated in the third block.

	within ± 3 min	within ± 5 min
Stockholm-Arlanda Airport (simplified)	73.2%	88.4%
Zurich Airport (simplified)	82.6%	91.6%
Stockholm-Arlanda Airport	94.4%	98.9%
Zurich Airport	95.6%	99.4%
Stockholm-Arlanda Airport (full)	96.1%	99.2%
Zurich Airport (full)	96.8%	99.7%
Detroit International Airport	89.9% - 97.1%	-
Tampa International Airport	89.9% - 95.7%	-
John F. Kennedy International Airport	-	20.7% - 100%

which is smaller than the given threshold value. An average of 95.7% was found for Detroit International Airport (DTW) and an average of 93.8% for Tampa International Airport (TPA) for ± 3 minute accuracy. The results for John F. Kennedy International Airport (JFK) were not very consistent and much less promising, showing ± 5 minute prediction accuracy between 20.7% and 100% for different days and parts of the day. Additionally, Idris et al. (2002) predicted 65.6% of the taxi-out times at Boston Logan International Airport within ± 5 minutes of the actual time. In contrast, our regression model found an average ± 3 minute accuracy of 94.4% for Stockholm-Arlanda Airport and 95.6% for Zurich Airport, considering both departures and arrivals simultaneously.

Reported taxi times at Stockholm-Arlanda Airport were from 1 to 16 minutes for arrivals and 3 to 20 minutes for departures. The seven cases which were not predicted within ± 5 minute accuracy were all departures with very long taxi times with the highest deviation being 7.40 minutes. Figure 4.8 shows the deviations of the estimated to the actual taxi times where the deviations are ordered. The rounded deviations are also shown (the step function), where the estimated taxi times are rounded to the nearest minute, to match the accuracy of the historic input data from Stockholm-Arlanda Airport, since many stakeholders are only interested to this level of accuracy. Taxi times at Zurich Airport ranged from 1 to 12 minutes for arrivals and 4 to 24 minutes for departures. Again, the four worst predictions were for aircraft with long taxi times and only one prediction was not within ± 6 minutes accuracy (but this has less than 8 minutes deviation).

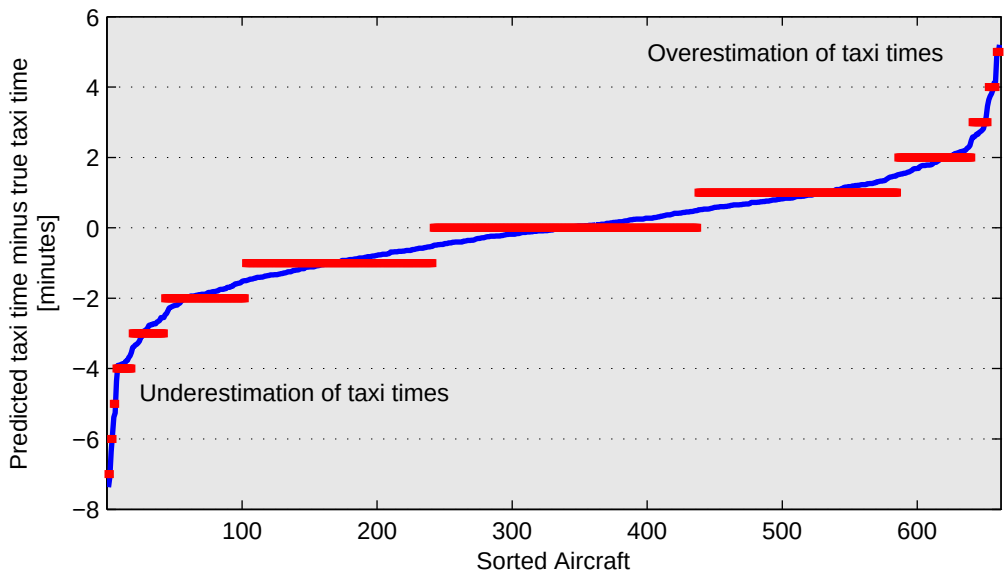

Figure 4.8: Taxi time prediction accuracy at Stockholm-Arlanda Airport

The results labelled '(simplified)' in Table 4.3 also show the prediction accuracy of our approach for both Stockholm-Arlanda Airport and Zurich Airport without taking the actual graph layout of the airports into account. A simplified regression analysis was performed without the different distance measures and the measures related to the turning angle. The significant improvements when the layout is considered emphasise the need for layout-based factors for airports where queuing is not dominating the whole ground movement process.

In contrast, the results labelled '(full)' in Table 4.3 correspond to the model with the best R^2_{Adj} value when considering all possible factors, rather than attempting to simplify the model. These indicate that the R^2_{Adj} would increase by around 2.2%. However, the aim of this research was to provide a practical model which was easy to interpret and hence the focus was not entirely on getting the model with the best accuracy.

As discussed in the introduction of this chapter, several other airport-related decision support systems as well as a wide variety of stakeholders at an airport (e.g. runway controllers, gate allocators, cleaning crews, de-icing crews, bus drivers, etc.) will benefit from better taxi time predictions.

4.5.2 Use for Ground Movement Decision Support

As discussed at the beginning of this chapter, algorithms that aim to optimise ground movement at airports need a model for predicting taxi times when there are no delays, since the interaction between aircraft would be explicitly considered by the model anyway. Such predicted uninterrupted taxi times can then be used to find a globally good solution by adding some delays or detours to aircraft where contention with other aircraft is indicated by the algorithm. The presented regression model allows such uninterrupted taxi time modelling by setting all N and Q values to 0.

Regression models work well within their range of observed data, but have to be handled with care for predictions at the boundaries and for extrapolations. Importantly, both datasets contain a number of observations with all N and Q values equal to 0 (for 3 departures and 9 arrivals at Stockholm-Arlanda Airport and 6 departures and for 4 arrivals at Zurich Airport) and these values are spread throughout the taxi speed range.

Once the regression approach has been implemented in a ground movement search methodology, it will be interesting to test the new system against the actual operations at the specific airport, and to fine tune the parameters to match the taxi times even more.

4.6 Results for London Heathrow Airport

The same multiple linear regression approach was also used to estimate taxi times at London Heathrow Airport (Atkin et al. 2011c). A dataset for one week's operations (9391 movements with outliers removed, see Chapter 3.2.3) was considered from summer 2010. The dependent variable was $\log_{10}(Speed)$ and $\log_{10}(Distance)$ and the N and Q values were used as explanatory variables. For Heathrow, it was found to be better to have separate regression models for departures and arrivals, and to separate cases depending upon which runway the aircraft were starting from or landing at (see table with the coefficients in Appendix B). The R^2_{Adj} value was 0.929 for departing aircraft (0.903 for runway 27R and 0.956 for runway 27 L) and 0.835 for arriving aircraft (0.812 for runway 27R and 0.861 for runway 27L), totalling to 0.882. Experiments with leave-one-out cross-validation, as explained in Section 4.3.5, indicated that

the R^2_{Pred} values were at most 0.1% smaller than the R^2_{Adj} values, leaving them very high.

Figure 4.9 shows four scatterplots for the linear fit of the regression models of the four different models. It is also clear from the figures that the linear fit for departing aircraft is better. Validations of the statistical assumptions were tested and they are approximately valid for all of the models.

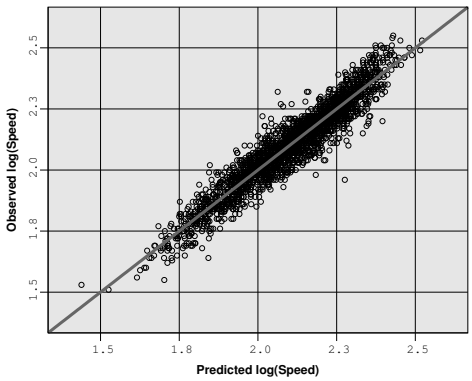

(a) Departures at runway 27R

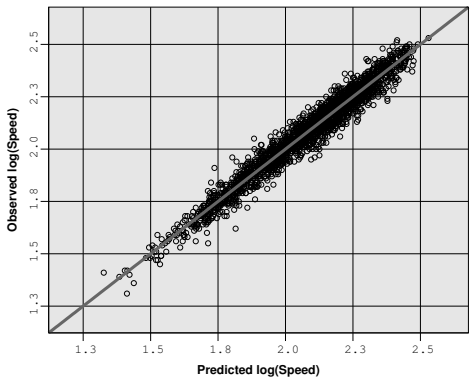

(b) Departures at runway 27L

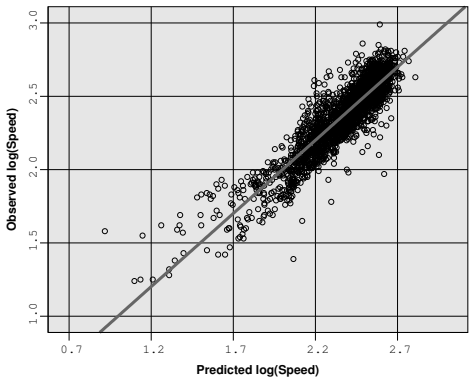

(c) Arrivals at runway 27R

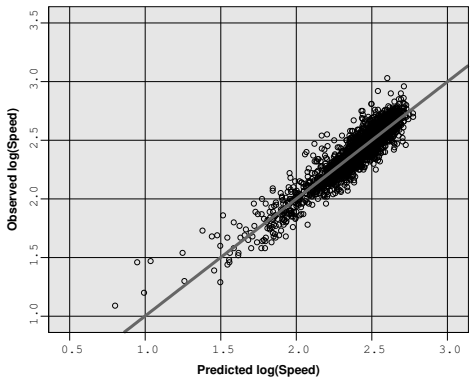

(d) Arrivals at runway 27L

Figure 4.9: Scatterplots showing the linear fit of the regression models for Heathrow Airport

4.7 Conclusions

With the current emphasis upon improving the predictions for on-stand times and take-off times (Eurocontrol 2012), an improved method for taxi time prediction is both important and

timely. This chapter analysed the variation in taxi speed and, consequently, the variations in taxi times, and considered not only departures but, for the first time, also arrivals. Data from Stockholm-Arlanda and Zurich Airport, both major European hub airports, was used for this research and the potential significant factors were identified and individually tested. In addition, similar experiments at London Heathrow Airport and at Hartsfield-Jackson Atlanta International Airport (see Appendix D) strengthened the findings. Multiple linear regression was used to find a function which could more accurately predict the taxi times than existing methods. An emphasis was placed upon ensuring that the function was easy to interpret and simple to use for operators at airports and researchers. Key for the analysis was the incorporation of information about the surface layout, since, in contrast to other airports which have previously been studied, the runway queuing was not dominating the entire taxi time.

The average speed between the gate and runway (and between the runway and gate) was found to be highly correlated to the taxi distance, with higher speeds being expected for longer distances. Arrivals had higher taxi speeds than departures, due to departure queues at the runway, and the quantity of traffic at the airport was also found to have a significant impact upon the average taxi speed, as identified by several variables in the resulting model. Finally, the total turning angle and the operating mode (which runways were in use) were also highly correlated to the average taxi speed.

Consideration of taxi time accuracy does not appear to have been sufficiently incorporated into the current state-of-the-art research in ground movement decision support systems at airports. Better predictions would, if nothing else, reduce the amount of slack which had to be used to allow for taxi time inaccuracies, allowing tighter schedules to be created. Historic data is vital for model calibration, but such data usually includes the effects of various inter-aircraft dependencies. When a decision support system takes care of the dependencies between the aircraft, predicted taxi speeds should not themselves include the effects of these dependencies. However, it is not usually obvious how to quantify and eliminate these effects. Amongst other uses, the approach which has been presented here could potentially be employed for such situations, allowing individual effects to be removed from consideration. The development of such a facility was the prime motivation for this research.

Since this work considers a combined statistical and ground movement model, which seems to accurately predict the effects of turns and congestion as well as total travel distances, we note here that these results can also feed into ground movement models, to improve the accuracy of the predictions for the effects of re-routing or delays.

Further research should explore more sophisticated ways of fine-tuning the parameters to increase the value of the approach for decision support systems for ground movement at airports, or other prediction approaches such as fuzzy rule-based systems (see next Chapter) or time series analysis (Chatfield 2003; Box et al. 2008).

5

Aircraft Taxi Time Prediction: Comparisons and Insights

> Prediction is very difficult,
> especially if it's about the future.
>
> Nils Bohr,
> Nobel laureate in Physics

5.1 Introduction

The latest vision for air transportation in Europe is predicting marked growth in this sector. The European Commission (2011) assumes an increase in the global volume of air traffic from 2.5 billion passengers in 2011 to 16 billion passengers in 2050. Thus, the number of commercial flights in Europe per year is expected to increase from 9.4 million to 25 million during the same time period. Nevertheless, one of the formulated goals by 2050 is also that on-time performance of flights is within 1 minute.

Efficient ground movement operations are key to successful operations of air transportation networks (Atkin et al. 2010b). For example, the benefits for take-off sequence of having accurate taxi times was shown in Atkin et al. (2008b) and recent developments for Heathrow (Atkin et al.

2012) require accurate taxi times both for take-off sequencing and for allocating pushback times to aircraft, at which they should leave the stands. A significant proportion of the actual travel time can be spent on airport's surfaces especially with short-haul flights. To achieve the stated on-time performance from the European Commission, it is crucial to more accurately predict taxi times at European airports.

Idris et al. (2002) published the first paper on taxi-out time estimation based on multiple linear regression. With the introduction of Collaborative Decision Making (CDM) systems at airports within the last few years (Pina et al. 2005; Pina and Pablo 2005; Eurocontrol 2012; Brinton et al. 2011), practitioners at airports realised the need for having more accurate taxi times and, driven by that, more researchers have analysed the problem of taxi time prediction. Several authors have published their results about taxi-out time prediction at US airports (Balakrishna et al. 2008a,b, 2009, 2010; Balakrishna 2009; Clewlow et al. 2010; Ganesan et al. 2010; Zhang et al. 2010; Srivastava 2011). Balakrishna et al. used a reinforcement learning algorithm which showed good results for data from Detroit International Airport (DTW) and Tampa International Airport (TPA), but the results were not very consistent for data from John F. Kennedy International Airport (JFK) (Balakrishna et al. 2008b, 2009, 2010; Ganesan et al. 2010). However, this approach cannot provide the same insights into the problem as some other approaches. Clewlow et al. (2010) highlighted that the number of arrivals does affect the taxi-out times, which was not sufficiently taken into account prior to that. Their multiple regression approach was based on John F. Kennedy International Airport and Boston Logan International Airport. Jordan et al. (2010) developed a sequential forward floating subset selection method with the aim of selecting the most influential explanatory variables from a set. It seems to be one of the few sources which analysed not only taxi-out times, but also taxi-in times. The analysis was performed with data from Dallas/Fort Worth International Airport (DFW). Kistler and Gupta (2009) developed a multiple linear regression approach, for the same airport, with several different explanatory variables, to predict taxi-in and taxi-out times.

All of the aforementioned publications were based on data from US airports. One problem of adopting these findings for Europe is that US airports are usually structurally different from European airports. For example, they distinguish between gate-ramps which are operated by

airlines, and taxiways, which are controlled by tower ground controllers. In addition, it seems that the problem of taxi time prediction in the US is dominated by the runway queue size and is less related to the actual distance that an aircraft has to taxi (Ravizza et al. (2012a) and Chapter 4). Furthermore, since no cross-validation details were often given in the papers, and the assumptions for multiple linear regression were not discussed, the importance of some of the findings was not clear.

Chapter 4 identified which explanatory variables affect the taxi time the most at two major European airports, Stockholm-Arlanda Airport and Zurich Airport. The utilised multiple linear regression approach incorporated explanatory variables based on the airport layout and not only fitted historic data well, but also predicted taxi times accurately. The assumptions for multiple linear regression were also tested, making the findings more reliable. Chen et al. (2011)[1] further improved the accuracy of the prediction by using a Mamdani fuzzy rule-based system based on the same explanatory variables which had been identified for Zurich Airport.

This chapter uses the same explanatory variables as in the research by Ravizza et al. (2012a) and Chapter 4, on datasets from the same airports, but with considerably longer operational periods. The aim is to test different regression approaches to more accurately predict taxi times, to demonstrate the advantages and disadvantages of these approaches and to give further insights into the problem, especially about taxi-in times. Such predictions can be used to make better overall decisions at airports and also to improve the quality of decision support systems for the ground movement problem at airports, by applying the findings and integrating the different aircraft speeds into such models (see Chapter 6).

The remainder of the chapter discusses the utilised datasets from Stockholm-Arlanda Airport and Zurich Airport in Section 5.2. Section 5.3 introduces six different regression approaches, which are tested in Section 5.4. This section also presents insights from the best performing approach, before Section 5.5 ends with the conclusions.

[1] Joint work between the University of Lincoln and the University of Nottingham

5.2 Considered Airport Data

Historic data from two European airports was utilised. All available data from each airport was combined into one dataset each and they were tested separately. This approach was used, since, as discussed by Demšar (2006), no statistical test exists which could compare different prediction methods based on different datasets where each prediction method is utilised for several repetitions of 10-fold cross-validation, due to the overlaps of the training data in different random samples.

Data from two entire days' operations were used within the analysis of Stockholm-Arlanda Airport (661 movements in datatset "ARN 1" and 656 movements in dataset "ARN 2"). The dataset for Zurich Airport consists of an entire day's operations (679 movements in dataset "ZRH 2007") and an entire week's operation (5611 movements in dataset "ZRH 2011"). More details of the datasets can be found in Chapter 3. The reported taxi time information is only to the minute rather than to the second. The only exception is the information about landing times on the runway at Zurich Airport, where detailed times have been recorded.

This research aims to compare various prediction methods and to find further insights into taxi time prediction at airports. Thereby, it extends the research by Ravizza et al. (2012a) and Chapter 4 which highlighted the statistically significant explanatory variables of this problem. The same explanatory variables were used in this study, which is based on more data from the same airports. All of the explanatory variables and their ranges for both datasets are shown in Table 5.1.

Appendix D presents the same analysis for Hartsfield-Jackson Atlanta International Airport as in this chapter for Stockholm-Arlanda and Zurich.

5.3 Regression Approaches to Predict Taxi Time

The aim of this research is to compare a wide range of different regression approaches for the problem of predicting taxi times at airports. WEKA (Hall et al. 2009) is an open source collection of machine learning algorithms for data mining tasks. It was used to explore which

Table 5.1: Overview of datasets from Stockholm-Arlanda Airport and Zurich Airport

	Range for ARN	Range for ZRH	Type
Taxi Time	[1,30]	[0.2,34.0]	Ordinal for ARN, Scale/Ordinal for ZRH
ARR	{0,1}	{0,1}	Nominal
Distance	-	[119,4320]	Scale
$\log_{10}$(Distance)	[2.8,3.7]	[2.1,3.6]	Scale
$\log_{10}$(Angle)	-	[1.9,3.0]	Scale
Distance2	[0,3373]	-	Scale
$Q_{DEP,\#DEP}$	[0,16]	[0,23]	Ordinal
$Q_{DEP,\#ARR}$	[0,13]	[0,26]	Ordinal
$Q_{ARR,\#DEP}$	[0,9]	[0,8]	Ordinal
$Q_{ARR,\#ARR}$	[0,9]	[0,7]	Ordinal
$N_{DEP,\#DEP}$	[0,12]	[0,15]	Ordinal
$N_{DEP,\#ARR}$	[0,6]	[0,6]	Ordinal
$N_{ARR,\#DEP}$	[0,12]	[0,16]	Ordinal
$N_{ARR,\#ARR}$	[0,6]	[0,6]	Ordinal
Mode 1	-	{0,1}	Nominal
Mode 3	-	{0,1}	Nominal
Mode 5	-	{0,1}	Nominal

regression models were promising for delivering good taxi time predictions, as this software contains a large selection of approaches to compare against our existing approaches. During an initial selection analysis, several regression approaches were not showing promising results and thus are not included in the rest of this study. These regression methods included, among others: a decision tree learner using reduced-error pruning, nearest neighbourhood methods, Pace regression, multilayer perceptron (back propagation neural network) and Gaussian processes (Witten et al. 2011).

Four regression techniques from WEKA which are showing promising results are explained in the rest of this section. These are: multiple linear regression, least median squared linear regression, support vector regression and M5 model trees. In addition, the approach published by Chen et al. (2011) with a Mamdani fuzzy rule-based system is used for comparison. Finally, a very promising extended version of another fuzzy rule-based system is also explained and utilised.

5.3.1 Multiple Linear Regression

Multiple linear regression is a very widely used regression methodology. It is not only very well studied, but it also has the advantage of determining which explanatory variables of a

model are significant. In this study, this technique is mainly used to act as a baseline for comparison with the other approaches. As discussed in Chapter 4, multiple linear regression is a statistical method attempting to model the dependent variable as a linear weighted function of the explanatory variables. The weights, or regression coefficients, can be estimated by using the least square approach. More in-depth coverage of multiple linear regression can be found in the book by Montgomery et al. (2001) and Chapter 4.

5.3.2 Least Median Squared Linear Regression

Least median squared linear regression is a more robust linear regression approach than multiple linear regression (see Figure 5.1). Instead of minimising the mean of the squares of the errors, this approach aims to minimise the median of these squares. Standard linear regression is applied iteratively to subsamples of the data and the solution with the smallest median of the squared errors is output (Rousseeuw and Leroy 1987). The advantage of being robust against the effects of outliers comes with the disadvantage of higher computational costs.

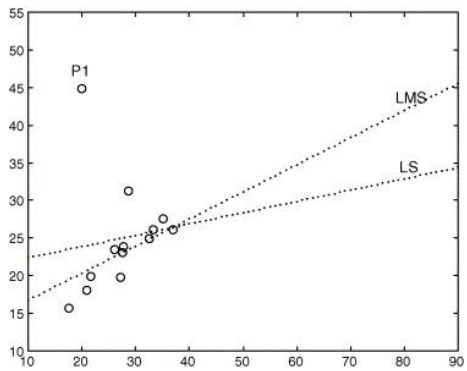

Figure 5.1: Effect of an outlier (P1) on least median squared (LMS) and least squared (LS) regression (source: Ortiz et al. (2006))

5.3.3 Support Vector Regression

Support vector machines are supervised learning methods and can be used for classification and regression analysis. Support vector regression ignores training data within a specified threshold ϵ of the model prediction (see Figure 5.2).

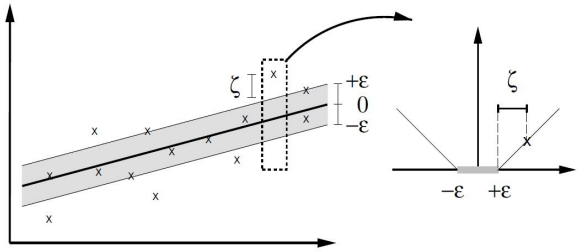

Figure 5.2: Linear support vector regression (source: Smola (1998))

The objective is to minimise the norm of the weights of the explanatory variables together with the error ζ for the training data which is further away from the prediction than the set threshold. A value C is normally defined to weight the trade-off between the two objective functions. The dual formulation of this optimisation model is often solved by preference. Support vector regression can be extended to non-linear models by incorporating a kernel function which transforms the original training data into a higher dimensional space. The best kernel found for this particular problem and these datasets turned out to be a normalised polynomial kernel (see Section 5.4.1). A tutorial for this approach can be found in Smola and Schoelkopf (2004).

5.3.4 M5 Model Trees

Another way of predicting numeric values is by using decision trees which store linear regression models on their leaves (see Figure 5.3). Such trees are called model trees and are similar to piecewise linear functions for the entire model. Model trees are usually smaller and more accurate than regression trees which have only an average value on their leaves (Quinlan 1992). The tree is constructed by the divide-and-conquer method, where the splitting criterion determines the best explanatory variables to split on, based on the expected error reduction. The splitting process finishes when the standard deviation of the subset of the training data is below a certain threshold or the size of this subset is too small. Afterwards, a linear regression model is calculated for each leaf. Pruning can be applied in a second stage: all non-leaves are tested for whether it is better to keep the subtree or whether a linear model could replace the subtree. Additionally, a smoothing stage can be added to reduce the discontinuities between the linear

models for different leaves. More details of the M5 model tree can be found in the paper by Quinlan (1992).

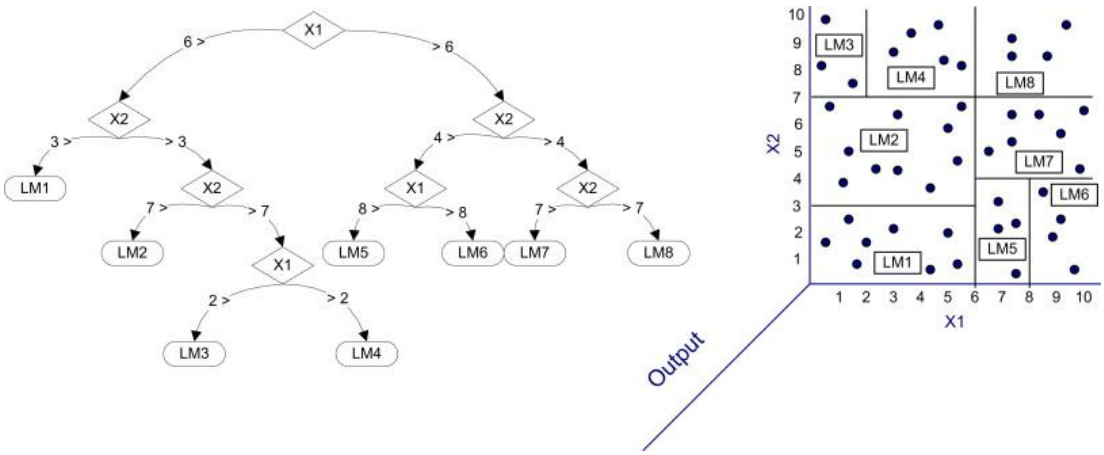

Figure 5.3: Example of M5 model tree (source: Bonakdar and Etemad-Shahidi (2011))

5.3.5 Mamdani Fuzzy Rule-Based Systems

Fuzzy Rule-Based Systems (FRBSs) are a way of modeling processes which have the ability to be interpreted with linguistic statements. First introduced by Zadeh (1965), they give the possibility of combining human expertise together with mathematical models. In addition, FRBSs, with the proven ability to approximate any real continuous function on a compact set to an arbitrary accuracy (Wang and Mendel 1992; Kosko 1994), should be very competent for modelling the non-linearity which is present in airport data. The input is first mapped with a fuzzification interface, then decisions can be made before these are mapped to a single crisp output using a defuzzification interface. The general process of fuzzy inference and its schematic diagram is shown in Figure 5.4.

The concept behind Mamdani FRBSs (Mamdani 1974) is introduced first here, before discussing another concept based on a different type of FRBS in the next section. The general 'rule-base' of a FRBS has the following form, with a number of fuzzy if-then rules R_i:

$$R_i: \text{ If } x_1 \text{ is } A_i^1 \text{ and } x_2 \text{ is } A_i^2, \ldots, \text{ and } x_j \text{ is } A_i^j \quad \text{Then } y_i = Z_i. \tag{5.1}$$

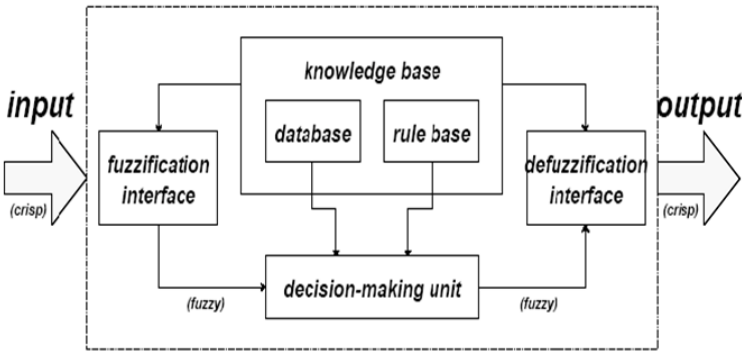

Figure 5.4: Fuzzy Inference Systems (source: Jang (1993))

The values x_l for all $l = 1, \ldots, j$ are the explanatory variables, y_i is the output of the ith rule and the A_i^l are the ith linguistic values (fuzzy sets). For each A_i^l, there is a membership function $\mu_{A_i^l}(x_l)$ associated with it which maps the universe of discourse to the range [0,1]. They take the form of Gaussian functions of the following form in this work:

$$\mu_{A_i^l}(x_l) = \exp\left[-\frac{1}{2} \cdot \frac{(x_l - c_i^l)^2}{(\sigma_i^l)^2}\right], \tag{5.2}$$

where c_i^l denotes the centre of the bell-shape curve and σ_i^l denotes the standard deviation. Figure 5.5 shows, for illustration, an example of a Gaussian membership function with its centre at 0.5 and its standard deviation at 0.2.

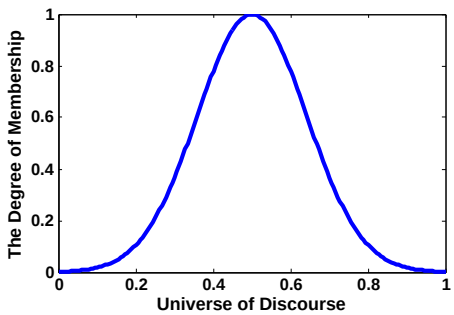

Figure 5.5: The shape of a Gaussian membership function for the explanatory variables

The consequence part Z_i is a fuzzy set for a Mamdani FRBS, which is here modelled as a bell-shaped membership function.

Each of the resulting rules can be expressed to the end user using linguistic terms, without showing the mathematical details of A_i^l and $\mu_{A_i^l}(x_l)$. For example, R_i, could also be rewritten as follows:

$$R_i: \text{ If } x_1 \text{ is big and } x_2 \text{ is small, } \ldots, \text{ and } x_j \text{ is medium} \quad \text{Then } y_i = Z_i,$$

where 'big', 'small' and 'medium' are linguistic values defined by $\mu_{A_i^l}(x_l)$.

The 'database' in Figure 5.4 contains all such membership functions for the fuzzy sets used in the fuzzy rules. Usually, the rule base and the database are jointly referred to as the 'knowledge base'. The 'decision-making unit' performs the inference operations on the rules and two interfaces perform fuzzification and defuzzification, respectively. Defuzzification is an important module since it converts a set of output values or output membership functions from different fuzzy rules into a single crisp output value.

As Jun Chen from the University of Lincoln developed the MATLAB code for this joint work (Chen and Mahfouf 2012), we point the interested reader to our paper (Chen et al. 2011) for more details about the approach, the tuning for the problem of estimating aircraft taxi times, and preliminary results.

Some of the key features of Mamdani FRBSs highlighted in Chen et al. (2011) are:

1. the ability to approximate complex non-linear systems,
2. the ability for rules to differ in different regions,
3. the ability to integrate human expertise, and
4. the ability to interpret the underlying system.

5.3.6 TSK Fuzzy Rule-Based Systems

Another form of fuzzy inference system, originally proposed by Takagi and Sugeno (1985), has fuzzy sets involved only in the premise part. By using Takagi and Sugeno's fuzzy inference

scheme (TSK), one can describe the fuzzy if-then rules as follows:

$$R_i: \text{ If } x_1 \text{ is } A_i^1 \text{ and } x_2 \text{ is } A_i^2, \ldots, \text{ and } x_j \text{ is } A_i^j \quad \text{Then } y_i = g_i(x_1, x_2, \ldots, x_j). \tag{5.3}$$

R_i denotes the ith rule to be considered. $g_i()$ is any function, and could, for example, be linear or quadratic. Normally, using a linear function for $g_i()$ is enough, since the fuzzy if-then rule has already embedded non-linearity inherently. When a linear model structure is assumed then a rule base with k rules takes the following format:

$$\begin{aligned} &R_1: \text{ If } x_1 \text{ is } A_1^1 \text{ and } x_2 \text{ is } A_1^2, \ldots, \text{ and } x_j \text{ is } A_1^j \quad \text{Then } y_1 = b_1^0 + b_1^1 \cdot x_1 + \ldots + b_1^j \cdot x_j \\ &\ldots \\ &R_k: \text{ If } x_1 \text{ is } A_k^1 \text{ and } x_2 \text{ is } A_k^2, \ldots, \text{ and } x_j \text{ is } A_k^j \quad \text{Then } y_k = b_k^0 + b_k^1 \cdot x_1 + \ldots + b_k^j \cdot x_j. \end{aligned} \tag{5.4}$$

The crisp output from the input $(x_1, x_2, \ldots, x_j)$ is obtained as the weighted sum of the consequences of the k rules:

$$y = \sum_{i=1}^{k} \beta_i \cdot y_i = \sum_{i=1}^{k} \beta_i \cdot (b_i^0 + b_i^1 \cdot x_1 + \ldots + b_i^j \cdot x_j), \tag{5.5}$$

where β_i actually represents the certainty of each rule contributed by the premise of the corresponding rule:

$$\beta_i = \frac{\mu_{A_i^1}(x_1) \cdot \mu_{A_i^2}(x_2) \cdot \ldots \cdot \mu_{A_i^j}(x_j)}{\sum_{l=1}^{k} \mu_{A_l^1}(x_1) \cdot \mu_{A_l^2}(x_2) \cdot \ldots \cdot \mu_{A_l^j}(x_j)}. \tag{5.6}$$

The membership function $\mu_{A_i^l}(x_l)$ of the premise part is again a Gaussian function as in Equation (5.2).

When a set of input-output data is given, one can obtain the consequent parameters $b_i^0, b_i^1, \ldots, b_i^j$ for all $i = 1, \ldots, k$ via some learning algorithms. As used for the Mamdani FRBS, Chen et al. (2011) utilised a combined k-means algorithm and genetic algorithm (Chen 2009) to automatically identify the initial values of the parameters both in the premise and consequent

parts, the same method is used to determine the parameters associated with the premise part. The least square approach is then used to determine the initial values of $b_i^0, b_i^1, \ldots, b_i^j$ for all $i = 1, \ldots, k$. To further refine the initial fuzzy system which is obtained, a genetic algorithm, namely G3PCX (Deb et al. 2002), is incorporated into TSK to fine-tune the premise part, followed by a least square approach to obtain the consequent part. This process continues iteratively until a pre-specified condition is met in order to reach a more accurate fuzzy system. G3PCX is a real-parameter genetic algorithm using a parent-centric recombination operator (PCX) and an elite-preserving, computationally fast evolutionary model (G3).

The Figures 5.7 and 5.10 together with the Tables 5.4 and 5.5 help the reader to better understand and visualise a possible model from a TSK fuzzy rule-based system.

As mentioned in Chen et al. (2011), in comparison to the Mamdani FRBS, the following distinctive features associated with the TSK FRBS can be identified:

- The TSK FRBS could in some ways be viewed as an extension of multiple linear regression. Each rule in the rule base resembles a multiple linear regression model for a decomposed explanatory variable region. Hence, the explanatory ability associated with multiple linear regression automatically applies to the TSK FRBS.
- These rules work cooperatively to produce estimations, which may result in more accurate estimations.
- Although one could lose certain linguistic meanings in the consequent part in comparison to the Mamdani FRBS, due to the function form of the TSK consequent part, such a form should be able to approximate the sub region more accurately than a fuzzy set.

5.4 Comparisons and Insights

This section analyses the different regression approaches for predicting taxi times at airports and shows comparisons between approaches and insights from the best performing approach. First, it is necessary to specify the experimental setup and the performance measures used.

5.4.1 Experiment Setup

All of the experiments were executed on a standard desktop PC (Intel Core 2 Duo, 3GHz, 2GB RAM). WEKA was used to perform all of the experiments apart from the last two models, which are related to the fuzzy rule-based systems (Hall et al. 2009; Witten et al. 2011). Analysis with the support vector regression method identified that the best parameter for the model is to use the value C equal to 2 and to employ a normalised polynomial kernel with exponent 3 (tested with discretised values over a reasonable value range and with different kernels). Both of the fuzzy rule-based systems were implemented by Jun Chen from the University of Lincoln, who provided the source code, and were tested in MATLAB R2010a. Mamdani FRBS was based on 12 rules, as published in Chen et al. (2011). TSK FRBS was analysed in detail and preliminary results showed that 4 rules for Stockholm-Arlanda Airport and 8 rules for Zurich Airport were the most promising settings. All experiments were based on 10-fold cross-validation if not otherwise stated. This is suggested to be the recommended setting and leads to relative low bias and variance (Nadeau and Bengio 2003; Han et al. 2011). Furthermore, 15 repetitions were done for each experiment, as recommended by Nadeau and Bengio (2003). They also recommend using the corrected resample t-test to test whether the difference between the two prediction models is significant. The corrected resample t-test should be preferred over a normal paired t-test, because the suggested test adjusts the variance in relation to the overlaps between subsets of the data (Demšar 2006). To compare the different models the same seeds were used to generate the subsets for cross-validation for the different repetitions for both utilised software packages. The significance level α was set to 0.05.

5.4.2 Performance Measures

This research aims to compare the different prediction methods using various performance measures to provide as much insight as possible and to enable further comparisons. The utilised measures are discussed below.

5.4.2.1 Root mean-squared error

The root mean-squared error (RMSE) is a very commonly used measure and was preferred over the mean-squared error in this chapter since it gives values in the same dimension as the predicted values. The formula is

$$\text{RMSE} = \sqrt{\frac{(y_1 - \hat{y}_1)^2 + \ldots + (y_n - \hat{y}_n)^2}{n}}, \tag{5.7}$$

where y_i is the actual taxi time of aircraft i and $\hat{y}_i$ is the corresponding predicted taxi time. The value n represents the number of aircraft in the dataset.

5.4.2.2 Mean-absolute error

A second measure, which is also in the same dimension as the predicted value, is the mean-absolute error (MAE). This performance measure averages the individual errors by neglecting their sign. It is defined as follows:

$$\text{MAE} = \frac{|y_1 - \hat{y}_1| + \ldots + |y_n - \hat{y}_n|}{n}. \tag{5.8}$$

5.4.2.3 Root relative-squared error

In the root relative-squared error (RRSE) the total squared errors are divided by the total squared errors when using the simplest prediction model (which just outputs the average value $\bar{y}$) as can be seen below:

$$\text{RRSE} = \sqrt{\frac{(y_1 - \hat{y}_1)^2 + \ldots + (y_n - \hat{y}_n)^2}{(y_1 - \bar{y})^2 + \ldots + (y_n - \bar{y})^2}}. \tag{5.9}$$

5.4.2.4 Relative-absolute error

The relative-absolute error (RAE) is the total absolute error, which is normalised as the root relative-squared error by using a simple average predictor:

$$\text{RAE} = \frac{|y_1 - \hat{y}_1| + \ldots + |y_n - \hat{y}_n|}{|y_1 - \bar{y}| + \ldots + |y_n - \bar{y}|}. \tag{5.10}$$

5.4.2.5 Coefficient of determination

A commonly used performance measure related to linear regression is the coefficient of determination R^2. It can be determined from the root relative-squared error and takes values between 0 and 1 for linear regression models, with values closer to 1 indicating a better fit. R^2 is defined as follows:

$$R^2 = 1 - \frac{(y_1 - \hat{y}_1)^2 + \ldots + (y_n - \hat{y}_n)^2}{(y_1 - \bar{y})^2 + \ldots + (y_n - \bar{y})^2}. \tag{5.11}$$

Sometimes an adjusted coefficient of determination is used, which penalises models with many explanatory variables. In this study, the explanatory variables are fixed and the size of the datasets is much larger than the number of explanatory variables, making such a correction term unnecessary.

5.4.2.6 Prediction accuracy

The last set of performance measures is used to show practitioners the accuracy of the models and is of a form that they will be familiar with. The percentage of the prediction accuracy measure indicates what percentage of the flights in the dataset are predicted within ± 1, 2, 3, 5 or 10 minutes.

5.4.3 Visual Comparisons

Figure 5.6 shows the predication accuracy of the 6 different regression approaches for Zurich Airport. The x-axis represents the aircraft, which were sorted from underestimated to over-

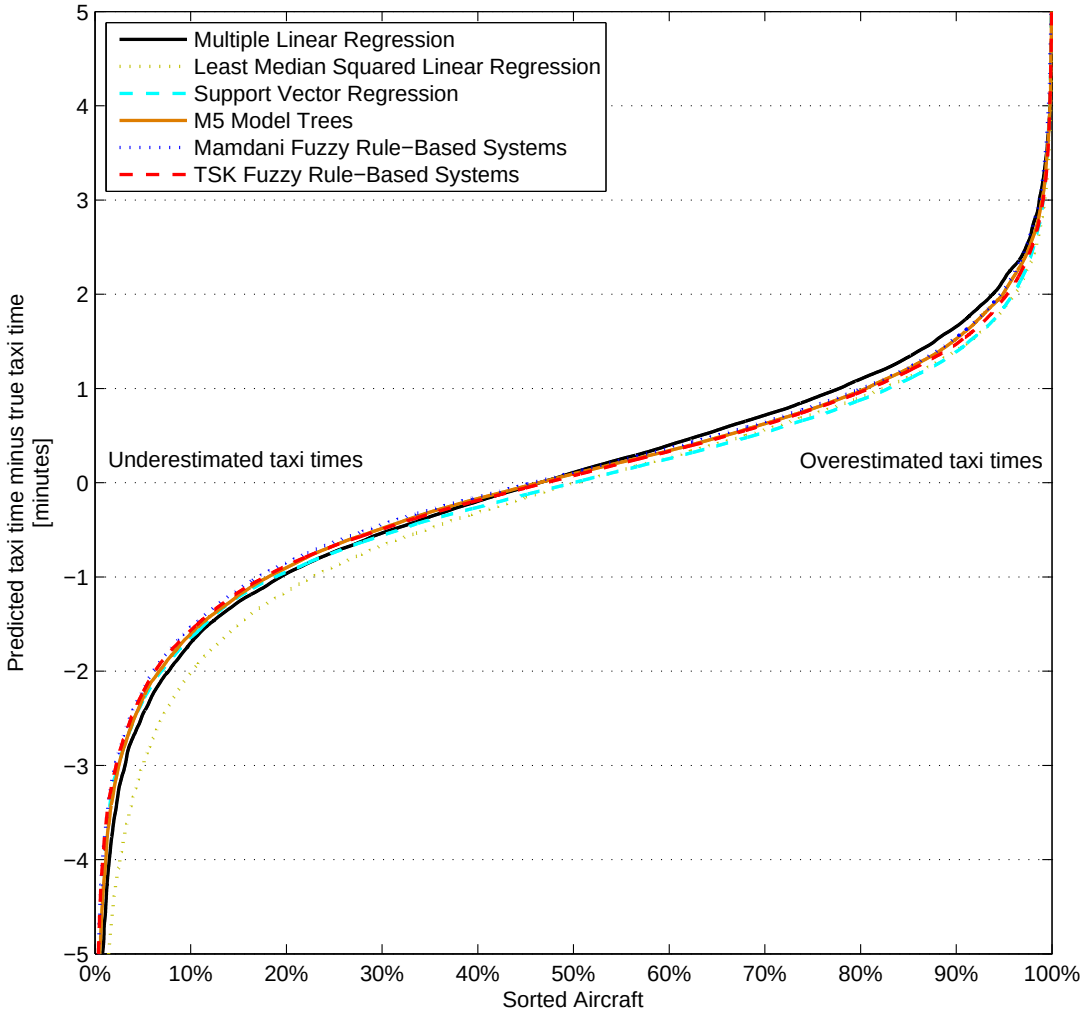

Figure 5.6: Taxi time prediction accuracy at Zurich Airport

estimated taxi times within each approach. The analysis is based on 15 repetitions of 10-fold cross-validation and shows each single error value. The range on the y-axis is only shown within the interval of ± 5 minutes. The solid black line visualises the multiple linear regression approach (LinReg) which is used as a baseline analysis. It is clear that least median square linear regression performs (LMS) poorly for predictions which underestimate the actual taxi time. Support vector regression (SMOreg), Mamdani FRBS and TSK FRBS seem to perform the best, but it is hard to distinguish clearly based on this figure, so a numerical comparison will now be presented.

5.4.4 Numeric Comparisons

Table 5.2 shows the first four performance measures for both airport datasets. Bold numbers highlight the best (smallest) result for each performance measure at each airport. The newly introduced TSK FRBS outperforms the other approaches in almost all cases. Only in the case of Zurich Airport does support vector regression have the same result for the mean-absolute error and be slightly better in terms of the relative-absolute error. Tests with the corrected resample t-test showed that there is always a significant improvement between the multiple linear regression approach and TSK FRBS at Zurich Airport. For this dataset, TSK FRBS also significantly outperformed least median square linear regression and, apart from the relative-absolute error, also outperform the M5 model trees. Although the numeric results are better for the TSK FRBS, it only significantly outperformed the Mamdani approach in terms of the root mean-square error and the root relative-squared error and did not outperform the support vector regression. The results for Stockholm-Arlanda Airport are very similar, but fewer tests identify significant differences. The best values found for the coefficient of determination R^2 were 80.85% and 93.25% for Stockholm-Arlanda Airport and Zurich Airport, respectively, using the TSK FRBS.

Table 5.2: Comparisons of performance measures for Stockholm-Arlanda Airport and Zurich Airport

Performance Measure	Airport	LinReg	LMS	SMOreg	M5P	Mamdani	TSK
Root mean-squared error	ARN	1.52	1.57	1.50	1.51	1.46	**1.44**
	ZRH	1.47	1.60	1.32	1.36	1.33	**1.30**
Mean-absolute error	ARN	1.14	1.14	1.09	1.13	1.07	**1.06**
	ZRH	1.08	1.10	**0.96**	0.99	0.97	**0.96**
Root relative-squared error	ARN	45.70%	47.47%	45.27%	45.54%	44.19%	**43.53%**
	ZRH	29.29%	31.76%	26.28%	27.00%	26.41%	**25.89%**
Relative-absolute error	ARN	45.80%	45.92%	43.98%	45.61%	43.28%	**42.83%**
	ZRH	26.85%	27.52%	**23.87%**	24.55%	24.30%	23.93%

The prediction accuracy within ± 1, 2, 3, 5 and 10 minutes can be found in Table 5.3. Again, TSK FRBS outperformed, in most cases, the other regression approaches in both datasets. Mamdani FRBS also did very well in comparison to the others. Our findings are in line with the analysis by Wu et al. (2011) where TSK FRBSs with fewer rules were more successful than Mamdani FRBSs. Although support vector regression often seemed a good alternative, the ± 2 and 3 minutes accuracy in the case of Stockholm-Arlanda Airport reported the worst values.

Table 5.3: Comparisons of accuracies for Stockholm-Arlanda Airport and Zurich Airport

Accuracy	Airport	LinReg	LMS	SMOreg	M5P	Mamdani	TSK
Accuracy within ± 1 min	ARN	54.28%	56.02%	57.86%	54.61%	58.21%	**58.80%**
	ZRH	58.38%	59.66%	**64.05%**	62.49%	62.97%	63.33%
Accuracy within ± 2 min	ARN	85.30%	85.18%	84.91%	85.19%	86.73%	**86.81%**
	ZRH	86.12%	85.99%	88.98%	88.15%	88.55%	**89.07%**
Accuracy within ± 3 min	ARN	95.40%	94.80%	94.32%	95.43%	95.72%	**96.16%**
	ZRH	95.55%	94.26%	96.60%	96.46%	96.54%	**96.89%**
Accuracy within ± 5 min	ARN	99.16%	98.81%	99.16%	**99.18%**	98.97%	99.08%
	ZRH	99.21%	98.56%	99.45%	99.46%	99.53%	**99.62%**
Accuracy within ± 10 min	ARN	**99.92%**	**99.92%**	**99.92%**	**99.92%**	**99.92%**	**99.92%**
	ZRH	99.92%	99.87%	99.97%	99.97%	**99.98%**	99.97%

Appendix D shows a case study of Hartsfield-Jackson Atlanta International Airport. TSK FRBS was again the best approach to predict taxi times and it can be highlighted that the ± 1 minute accuracy can be improved by 26% when using a TSK FRBS instead of the baseline experiment with the multiple linear regression approach.

5.4.5 Insights from Prediction Models

As the comparisons in the last section showed, TSK FRBS seems to perform the best for the analysed prediction problem. Therefore, this section focuses on the results of that particular approach. Earlier results and analysis for multiple linear regression and Mamdani FRBS can be found in the papers by Ravizza et al. (2012a) and by Chen et al. (2011), respectively.

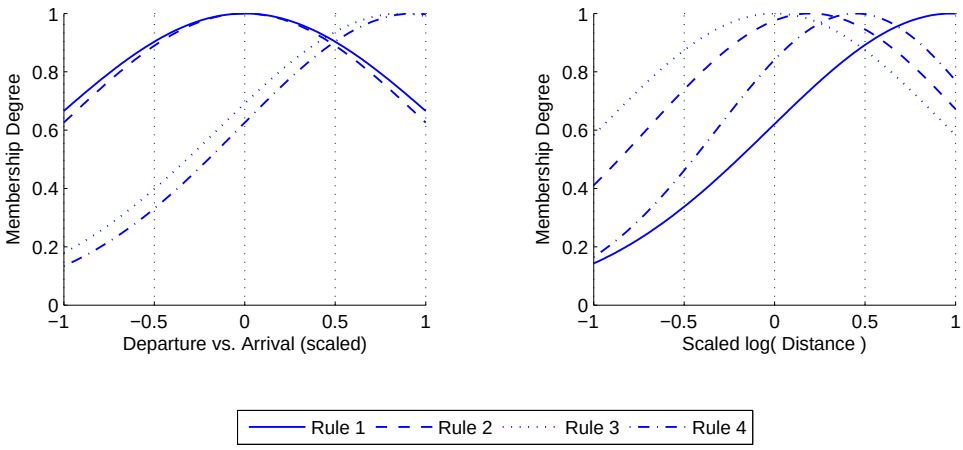

Figure 5.7: Four fuzzy rules extracted from the TSK FRBS analysis for Stockholm-Arlanda Airport

Figure 5.7 illustrates the membership functions of two explanatory variables for the TSK FRBS at Stockholm-Arlanda Airport. The first figure is related to the indication of whether an aircraft is arriving or departing and the second figure shows the final model related to the logarithmic transformation of the total taxi distance. Input values are scaled to the range [-1,1] using a linear scaling from their range (see Table 5.1). The other explanatory variables had less distinct membership functions for the different rules and were omitted. The four rules are represented with different lines, showing rules 1 and 2 are more focusing on departures and rules 3 and 4 on arrivals. Furthermore, the rules cover different total taxi distances starting with rule 3 for smaller distances, following by rules 2, 4 and 1. The consequence part of the four rules can be seen in Table 5.4. Each rule has the form of a multiple linear regression approach with the coefficients stated in the table. This model is based upon using the approach with the entire dataset as training data.

Table 5.4: Consequence part of the TSK FRBS analysis for Stockholm-Arlanda Airport

	Rule 1	Rule 2	Rule 3	Rule 4
(Constant)	0.793	-0.146	-0.964	1.666
ARR	0.772	0.339	-1.452	-1.080
Distance2	0.082	-0.208	-0.587	-0.224
$\log_{10}$(Distance)	-0.200	0.091	-0.871	-0.368
$Q_{DEP,\#DEP}$	0.594	2.891	-1.554	-2.144
$Q_{DEP,\#ARR}$	0.550	-6.788	10.588	5.890
$Q_{ARR,\#DEP}$	-0.415	-0.569	1.118	1.042
$Q_{ARR,\#ARR}$	0.041	3.828	-1.002	-0.080
$N_{DEP,\#DEP}$	-0.019	-0.273	-1.767	0.575
$N_{DEP,\#ARR}$	-0.127	5.007	-7.334	-4.586
$N_{ARR,\#DEP}$	-0.169	0.455	-0.472	-0.038
$N_{ARR,\#ARR}$	-0.139	-2.314	0.906	0.089

TSK FRBS has the ability to model non-linear relationships, which gives such a method an advantage over multiple linear regression, least median squared linear regression or M5 model trees, where automatic transformations of linear functions with a polynomial transformation are only possible to a certain extent (Box and Cox 1964; Weisberg 2005; Osborne 2010). Figure 5.8 shows how the functions look for the TSK FRBS and multiple linear regression on the example of the total taxi distance at Stockholm-Arlanda Airport. Two functions are plotted for each approach, one showing the predicted taxi times for departing aircraft (solid lines) and one for arriving aircraft (dashed lines). In the case of multiple linear regression, the two functions

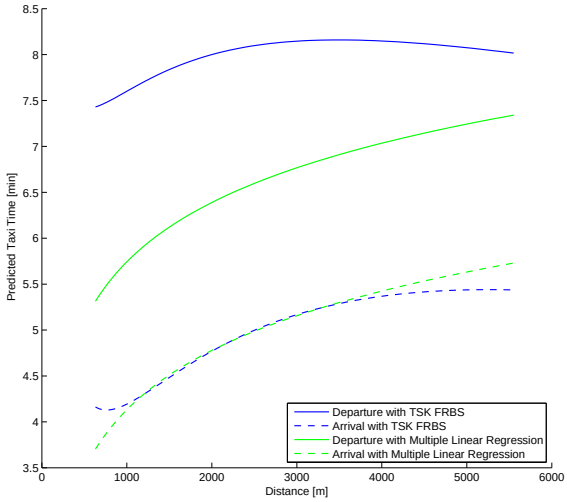

Figure 5.8: Non-linearity of TSK FRBS models shown on data from Stockholm-Arlanda Airport

always differ by the same taxi time, which is equal to the coefficient of the explanatory variable indicating the difference between arrivals and departures of the multiple linear regression model. The two functions have a logarithmic shape due to the transformation of the explanatory variable related to distance. On the other hand, the shapes of the two functions of the TSK FRBS can differ and can model non-linear behaviour. This is demonstrated in Figure 5.8, where the model only consists of four fuzzy rules.

Figure 5.9 shows the predicted taxi times of arriving aircraft at Stockholm-Arlanda Airport with the TSK FRBS. One axis again represents the different total taxi distances and the other axis the amount of traffic at the airport. The analysis was set as following: the Q and N values counting other arrivals are indicated on the axis while the Q and N values counting other departures were set to 1 for the entire analysis. An increase in the total taxi distance increases the predicted taxi time and the curve flattens for longer distances. Analysing the amount of traffic on the surface shows that the slope of the curve grows with higher Q and N values. In addition, the influence of the total taxi distance decreases with more traffic on the surface. Again, it should be highlighted that such behaviours cannot be modelled with a linear regression approach, suggesting that one reason for TSK FRBS outperforming the other approaches is due to its ability to model complex non-linear systems. Figures for departures

and for Zurich Airport can be found in Appendix C. These figures show that regression models are airport dependent and linear or quadratic models are less likely to predict taxi times to the same accuracy.

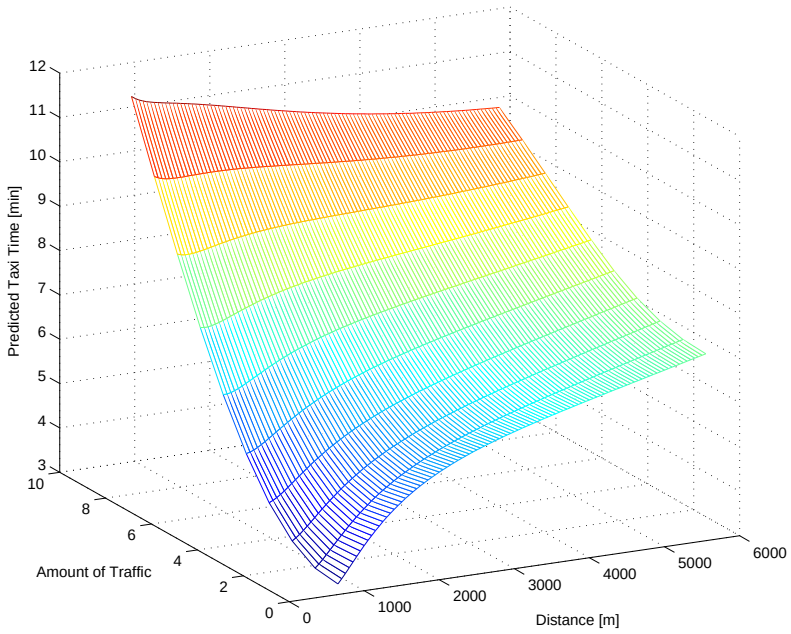

Figure 5.9: Analysis of predicted taxi-in times at Stockholm-Arlanda Airport with the TSK FRBS

The model for the TSK FRBS for Zurich Airport is shown in Figure 5.10 in the same way as in Figure 5.7. As mentioned earlier, the best results at Zurich Airport were found with eight fuzzy rules. The most interesting membership functions were for the indication of arrivals and departures and the logarithmic transformation of the total taxi distance as it was before for Stockholm-Arlanda Airport. In addition, the logarithmic transformation of the total angles and the three operational modes also have distinct functions. The corresponding multiple linear regression models for the eight different fuzzy rules can be found in Table 5.5.

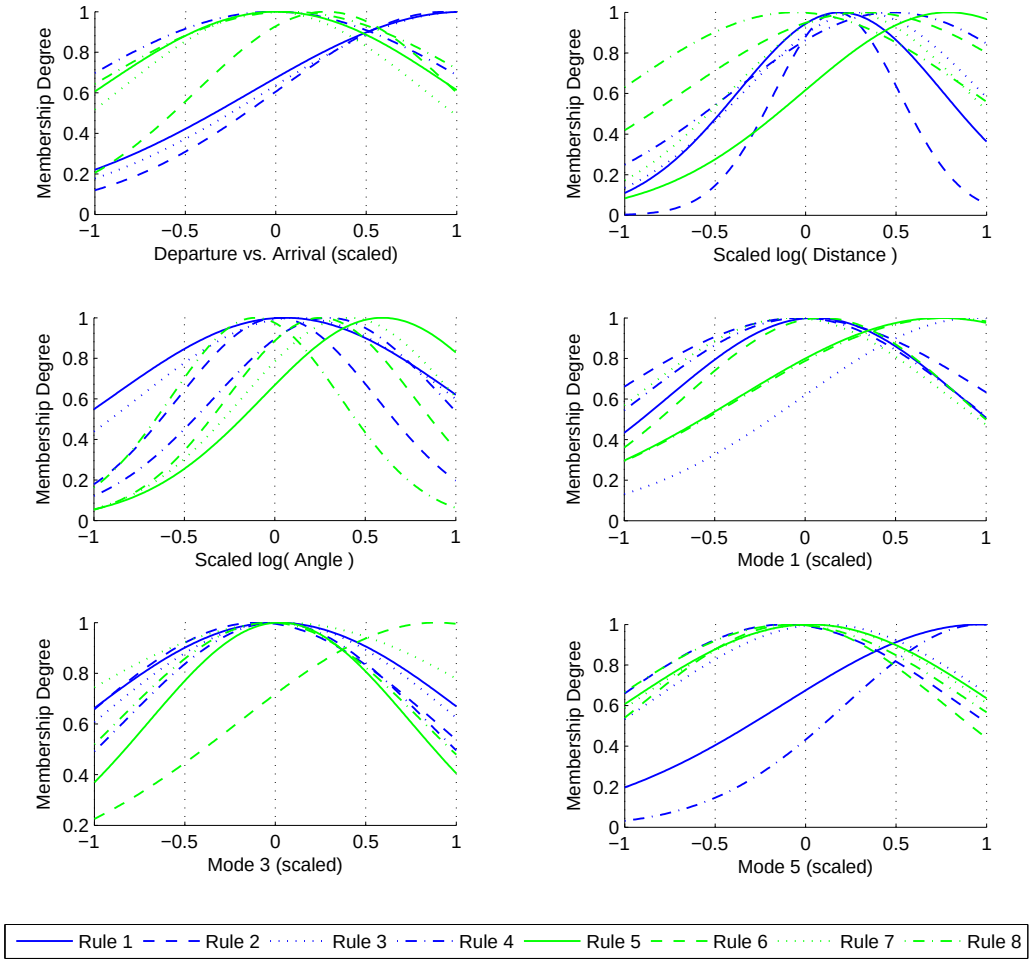

Figure 5.10: Eight fuzzy rules extracted from the TSK FRBS analysis for Zurich Airport

5.5 Conclusions

Airports commonly have many objectives and most will already be automating some of their operations or decision support, or be planning to do so in the future. Some of the aims of doing so are to improve predictability, improve on-time performance, reduce ground movement costs, enhance the use of ground handling resources, stands, gates and terminals and reduce apron and taxiway congestion (Eurocontrol 2012). Key to such improvements are better taxi time predictions, which can help many different decision support systems. This research should also help the development of new and more accurate decision support systems for the ground

Table 5.5: Consequence part of the TSK FRBS analysis for Zurich Airport

	Rule 1	Rule 2	Rule 3	Rule 4	Rule 5	Rule 6	Rule 7	Rule 8
(Constant)	4.983	-4.299	0.424	-19.232	0.923	-0.203	3.008	2.233
ARR	2.608	0.398	-0.273	-2.787	-0.037	-0.062	-0.334	-0.169
Distance	-1.294	-2.751	0.466	2.004	0.954	-2.952	-0.342	1.694
$\log_{10}$(Distance)	0.374	2.778	-0.321	-1.031	-2.494	3.096	0.280	-0.582
$\log_{10}$(Angle)	1.598	-0.443	0.060	-0.579	0.688	0.155	-0.603	-0.270
$Q_{DEP,\#DEP}$	-0.215	-1.108	0.825	1.096	0.294	-2.289	1.952	0.936
$Q_{DEP,\#ARR}$	7.583	0.063	0.232	-1.360	0.039	0.528	1.633	0.827
$Q_{ARR,\#DEP}$	-0.028	-0.247	0.168	-1.639	0.025	-0.941	0.545	0.887
$Q_{ARR,\#ARR}$	1.422	0.351	0.009	-0.333	1.023	1.772	-0.448	1.069
$N_{DEP,\#DEP}$	1.727	1.229	-0.707	-0.176	0.283	-0.763	-0.367	0.194
$N_{DEP,\#ARR}$	0.764	-0.213	-0.240	-0.377	0.074	1.022	-0.379	-0.077
$N_{ARR,\#DEP}$	0.296	0.001	0.061	1.111	-0.955	-0.818	0.542	-0.369
$N_{ARR,\#ARR}$	-0.643	-0.810	0.088	-1.055	-1.326	-0.508	0.568	-1.157
Mode 1	1.072	0.214	-0.034	5.916	-0.797	-2.149	0.201	-0.166
Mode 3	-2.264	-0.317	0.154	2.814	0.122	-2.580	-0.099	0.074
Mode 5	1.342	-0.035	0.056	21.321	0.251	1.665	0.086	0.127

movement problem.

Chapter 4 focused on finding the explanatory variables for taxi time prediction for both arrivals and departures, using multiple linear regression to highlight their statistical significance. This chapter uses the same explanatory variables and shows an analysis of different regression approaches for predicting taxi times at airports to demonstrate the performance of each. Six different approaches were analysed in detail: multiple linear regression, least median squared linear regression, support vector regression, M5 model trees, Mamdani fuzzy rule-based systems and TSK fuzzy rule-based systems. The latter outperformed the other approaches on datasets from two European hub airports and the world's busiest airport (see Appendix D). TSK fuzzy rule-based systems use fuzzy membership functions to subdivide the input space in the premise part and a weighted sum of multiple linear regression approaches in the consequent part. As the different fuzzy rules work cooperatively, in contrast to approaches such as M5 model trees, the approach may potentially give more accurate estimates and can also model non-linear patterns in the data. Furthermore, this chapter gave insights into the different rules found by the TSK fuzzy rule-based system and considered taxi-in times, which seems to be a less understood problem in this field.

It would be interesting to also compare these regression approaches for other busy airports

to see whether these findings can be extended into settings where the airport operations are managed differently or are operated under differing constraints. In addition, this research could be integrated into decision support systems which help controllers in the towers, followed by a fine-tuning phase of the models and the decision support systems to provide more valuable decision-making aids.

6

A More Realistic Approach for Airport Ground Movement Optimisation with Stand Holding

> Do not worry about your difficulties in Mathematics. I can assure you mine are still greater.
>
> Albert Einstein

6.1 Introduction

European airports face several challenges in the 21st century, including the capacity challenge (with demands for air travel still increasing year on year) and the environmental challenge (ACI EUROPE 2010). To avoid forming huge bottlenecks in the air transportation system, airports have either to be enlarged, or (since enlargement is either not possible or prohibitively expensive in most cases), to utilise the existing resources as efficiently as possible. De-peaking hub-and-spoke flight schedules would be an alternative, but can cause revenue decreases for airlines, as it was the case for Delta Air Lines with their project "Operation Clockwork" in

2005 (Petroccione 2007). In addition, the increasing focus upon environmental issues is likely to further grow over time. As airports work closer to their maximum capacity, airside airport operations become much harder to address. As a result, decision support systems have to be increasingly advanced and they need to integrate different airside airport operations with each other and to model each process increasingly realistically.

From an optimisation point of view, ground movement of aircraft can be considered to be one of the most important airside operations at an airport, since it links several other problems together, such as the runway sequencing problems for arrivals and/or departures (Atkin et al. 2007), the stand holding problem (Atkin et al. 2011a) and the gate assignment problem (Dorndorf et al. 2007). A comprehensive literature review of ground movement research and the integration with other operations can be found in Chapter 2.

This chapter presents a decision support framework for environmentally friendly ground movement, along with promising experimental results which utilise more realistic taxi time predictions for a European hub airport. A framework is described for integrating a graph-based sequential movement algorithm into a larger decision support system which can also consider the runway sequencing problem and the stand holding problem. A Fuzzy Rule-Based System (FRBS) has been used to more accurately estimate taxi and pushback times for aircraft than a standard lookup table may allow. This utilises the same graph which is employed for the ground movement model. This integrated approach allows the effects of ground plan changes to be modelled more accurately, changing both taxi time predictions and routing information. In addition, several concepts have been included in the model which allow airport layouts to be modelled in a more realistic manner, such as restricting certain taxiways to be used only by certain aircraft and coping with the required separations between aircraft. Finally, the absorption of delay at the stand, before to starting the engines, has been considered. This reduces the waiting times at the runway and is further extending previous stand holding ideas (Burgain et al. 2009; Atkin et al. 2010a, 2011a). The potential benefits of such a system have been quantified.

Section 6.2 provides a description of the airport ground movement problem and how it can be embedded into the larger combined sequencing/routing/stand holding framework. Details

of the dataset which were provided by the airport are then presented in Section 6.3 together with the method for estimating taxi times. Following this, the sequential ground movement algorithm which has been developed, and was utilised for these experiments, is detailed in Sections 6.4 and 6.5. The results of the application of the algorithm to the dataset are then shown in Section 6.6. The chapter ends with some conclusions in Section 6.7.

6.2 Problem Description

The links between the ground movement problem and runway sequencing are considered first in this section, before the ground movement problem itself is discussed in more detail. The section ends with a consideration of the stand holding benefits which can result from the appropriate solution of the ground movement problem.

6.2.1 The Links with Runway Sequencing

Atkin et al. (2010b) highlighted the importance of integrating the ground movement problem with other airside airport operations, such as the problems of finding good departure and arrival sequences. Supporting controllers in these tasks is a challenge, especially when departures and arrivals have common restrictions and interactions due to the airport layout. For this chapter, we assume that the runway sequencing and ground movement problems are solved as two distinct stages. The integrated (departures and arrivals) runway sequencing problem is assumed to be solved in a first stage, then the consequent landing and take-off times are used in the second stage, within the consideration of the ground movement problem. Thus, the wheels-on time at the runway (for arrivals) and the wheels-off time at the runway (for departures) are both assumed to be fixed within the ground movement problem. Issues such as conformance with take-off time slots are assumed to be taken into account by the runway sequencing stage. This decomposition has been found to be effective, but further research will analyse the benefits of providing a feedback loop from the ground movement problem to the integrated runway sequencing problem and of closer integration between the two problems.

6.2.2 Problem Description of the Ground Movement Problem

This chapter considers ground movement at an airport. The ground movement problem is a combined routing and scheduling problem. It involves guiding aircraft on the surface of an airport to their destinations in a timely manner, where the goal is to reduce the overall travel time and to enable the target take-off times at the runway to be met. It is important that two aircraft never conflict with each other throughout the ground movement process.

In the model which is considered in this chapter, the route of the aircraft is not pre-determined, allowing greater flexibility for solutions. However, the utilised solution method provides the possibility to restrict certain aircraft to specific taxiways and/or to avoid routes which involve tight turns. The airport layout is represented as a directed graph, where the edges represent the taxiways and the vertices represent the junctions or intermediate points. Aircraft are considered to occupy edges, and conflicts are avoided by preventing any two aircraft from using the same edge simultaneously, or from simultaneously using edges which are too close together.

The sequential approach to ground movement will then minimise the taxi time for each individual aircraft given the planned movement for the other aircraft which have already been routed. Hence, the approach will attempt to absorb as much of the waiting time as possible at the gate/stand, allowing the departures to start their engines as late as possible, reducing fuel burn and environmental impact. Thus, the solution method could be considered to be not only reducing the ground movement time, but also solving the stand holding problem (Burgain et al. 2009; Atkin et al. 2010a, 2011a) for a given runway sequence.

6.3 Analysed Case: Zurich Airport

This analysis utilised data from Zurich Airport. The major part of the analysis is based on the dataset “ZRH 2011” for an entire week’s operations with 5613 movements in total (2806 arrivals and 2807 departures). A preliminary study is based on the smaller dataset from Zurich Airport “ZRH 2007” which was available at the time of the study. In addition, this preliminary study needed very long experimental runtimes, which was only reasonable to analyse for one day of operation.

6.3.1 Taxi Time Prediction

Ground movement models need accurate taxi time predictions, but sufficiently accurate values are rarely available. Comparisons between ground movement tool results and the status quo at airports have previously been hard to analyse, due to the need for accurate taxi speed data. The historic data which has to be used usually includes the effects of any delays or re-routing due to conflicts between aircraft, so the effects of taxi time variability and the benefits from the ground movement decision support system were often intermingled. This research confronts that challenge.

An approach to more accurately predict taxi times for aircraft or, equivalently, their average speeds, was proposed in Ravizza et al. (2012a) and Chapter 4 with a multiple linear regression approach. The aim was to be able to eliminate the effects of factors which represented the actual amount of traffic at the airport (by zeroing the factors related to airport load), with the goal being to predict the taxi times for unimpeded aircraft. These predictions could then be used in a more advanced ground movement decision support system, such as the one described in this chapter, which would itself model the effects of the interaction between aircraft (so these should not already be included in the taxi speed data). Chapter 5 introduced a Mamdani FRBS approach to estimate taxi times at airports and was adopted and extended for this research.

It was observed for Zurich that some aircraft have to push back from their allocated gates, taking additional time to do so, whereas other gates allow aircraft to immediately start their engines. The work by Ravizza et al. (2012a) and Chapter 4 was extended to include a pushback duration and the multiple linear regression approach indicated that this factor was significant for Zurich. The resulting taxi time prediction functions by Chen et al. (2011) were therefore further enhanced for this work adding a predicted pushback duration to the taxi time for the first edge for departures where the gate requires it, before being utilised to predict the taxi times.

Finally, depending upon the terminal and the operating mode (which runways are in use), runway crossings may be necessary during the taxi process. For the moment, these are included only in the prediction model for taxi times (having influenced the historic data), but we plan to integrate these effects into the combined ground movement and sequencing model later.

6.4 Ground Movement Decision Support System

Figure 6.1 provides an overview flowchart describing the ground movement algorithm. Further details are provided later. The aircraft are routed sequentially in this approach. When an aircraft is ready, it has to be routed respecting all previous reservations by other aircraft using the taxiways. The routes which have been previously calculated for other aircraft do not normally change as new aircraft are taken into consideration (the exceptions are discussed in Section 6.5). This has an advantage for the dynamic case, where some aircraft will have prior instructions, and acknowledges the difficulty and time costs associated with communicating changes to pilots and reducing the quantity of communication needed between the surface controllers and pilots. The objective for each of the sequential routings is to find the routing with minimal taxi time among all remaining conflict-free routings.

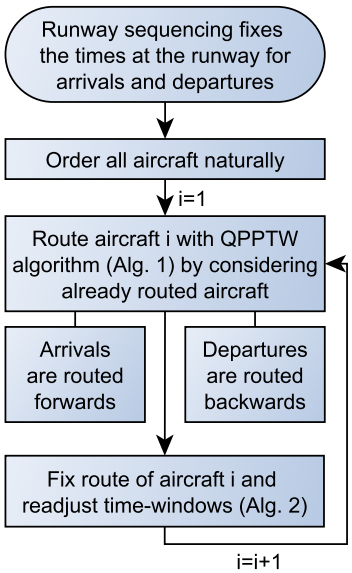

Figure 6.1: Flow chart of general concept of the approach

The approach described here is based on research by Gawrilow et al. (2008) and the PhD thesis of Stenzel (2008) which advances earlier work of Desrochers and Soumis (1988) and Sancho (1994). Ravizza modified the approach for his Master's dissertation (Ravizza 2009) to label the vertices instead of the edges, to simplify their interpretation. The original aim of this approach

was to control automated guided vehicles in container terminals in harbours or in storage areas, but it is here applied instead to routing aircraft. The approach has been further modified for this work. The approach has been extended that it can be applied forwards and in addition also backwards to meet a specific end time rather than a specific starting time (see Section 6.4.5). Furthermore, different heuristics were integrated to improve the solution quality by changing the sequence in which the aircraft are routed (see Section 6.5). The resulting algorithm is described in this section.

The Quickest Path Problem with Time Windows (QPPTW) algorithm is a generalised vertex-based label-setting algorithm based on Dijkstra's algorithm and can sequentially route aircraft on the airport surface, using a directed graph model of the airport. No time discretisation is used in this approach, in contrast to many other ground movement support systems (Balakrishnan and Jung 2007; Marín 2006; Marín and Codina 2008; Roling and Visser 2008). It has similarities to the recently published work by Lesire (2010), which used a sequential A* algorithm, but it provides a better coverage of the solution space, potentially allowing it to find better solutions within comparable execution times - these being short enough for it to be appropriate for real-time decision making. It also provides the possibility to define which edges in the graph are in conflict with each other and hence cannot be used simultaneously. In addition, for each edge incident to a vertex, the set of valid outgoing edges can be manually defined if desired, or can depend upon information about the aircraft. This enables the decision support system to forbid aircraft from making tight turns or to prevent aircraft from using taxiways for which they are too large. Together, these features enable the approach to more realistically model the airport surface while leaving the routing task itself to the algorithm.

The preprocessing of the algorithm is explained in Section 6.4.1, then the key concepts are introduced. The QPPTW algorithm is detailed next and the section ends with a discussion about buffer times and the sequence in which aircraft are routed.

6.4.1 Ground Plan Preprocessing

It is important to maintain separations between aircraft on the ground. The concept of conflicting edges is introduced here for this reason, so that no two conflicting edges can be occupied

simultaneously. The conflicting edges are determined in a preprocessing stage. For this research, we used an approach which assumes straight connecting lines between vertices, since this requires less time in the preprocessing stage and is adequate for the directed graph model which has been used in this research, where the paths are almost straight lines between vertices. Edges in the graph, together with their embedding in the airport plan, are here named segments. In this approach, two segments conflict with each other if they are located closer together than a given threshold distance. To find the minimal Euclidian distance between two segments, the algorithm performs two processing steps. Firstly, it verifies whether the edges are intersecting, then, if they are disjoint, the distance between each end point of one segment and the closest point on the other segment is calculated (see Figure 6.2). The minimum over these four distances corresponds to the minimal distance between the two segments.

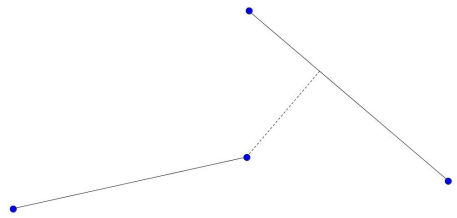

Figure 6.2: Euclidian distance between two segments

6.4.2 Variable Definitions

Definitions of the variables and data structures which are used in the model are given in Table 6.1.

6.4.3 Key Concepts

The QPPTW algorithm with its expansion steps works in a similar way to Dijkstra's algorithm (Dijkstra 1959; Cormen et al. 2001). However, a label can be expanded several times due to the different time-windows and an additional concept of dominance is needed in order to guarantee a polynomial solution time. It is necessary to define some of the concepts upon which the approach is based. Firstly, the algorithm needs information about the times that each part of the taxiway (edge) is free:

Table 6.1: Table of definitions

Variable	Explanation
$confl(e)$	The set of edges which conflict with edge $e \in E$
$F_e^j = [a_e^j, b_e^j]$	jth time-window on edge $e \in E$, from time a_e^j to time b_e^j
$\mathcal{F}(e)$	The sorted set of all the time-windows on edge $e \in E$
$G = (V, E)$	The directed graph representing the airport layout, with vertices $v \in V$ and edges $e \in E$
H	The Fibonacci heap storing the added labels
$I_L = [a_L, b_L]$	The time interval used in a label L
$L = (v_L, I_L, pred_L)$	A label on vertex $v_L \in V$ with time interval I_L and predecessor label $pred_L$
$\mathcal{L}(v)$	The set of all of the labels at vertex $v \in V$
R	A conflict-free route that is being generated
$T = (s, t, time)$	A taxi request to route, from source $s \in V$ at time $time$ to target $t \in V$
w_e	The weight (necessary taxi time) of edge $e \in E$

Definition: Set of sorted time-windows

The set $\mathcal{F}(e)$ contains the sorted set of time intervals $F_e^j = [a_e^j, b_e^j]$ which specify the times when the edge e can be used for a new route. This will exclude the times when e, or an edge which conflicts with e, are in use by previously routed aircraft. These are inputs to the routing algorithm for each aircraft.

The use of labels is an essential concept of the QPPTW algorithm:

Definition: Label

A label $L = (v_L, I_L, pred_L)$ *specifies the time period $I_L = [a_L, b_L]$ within which the current aircraft could reach vertex v_L. It includes a reference to the previous label on the route, $pred_L$, and thus implicitly represents a route (with edge traversal timings) from a source vertex to the specified vertex v_L. These labels are generated as the routing algorithm progresses, together specifying the (undominated) time periods (from time a_L to time b_L) when the current aircraft could reach vertex v_L.*

An ordering relation is defined over the intervals of the labels to allow the definitions of dominance:

Definition: Dominance

A label $L = (v_L, I_L, pred_L)$ dominates *a label* $L' = (v_{L'}, I_{L'}, pred_{L'})$ *on vertex* $v_L = v_{L'}$ *if and only if* $I_{L'} \subseteq I_L$ *(and there are identical route restrictions on the outgoing edges), which implies* $a_L \leq a_{L'}$ *and* $b_L \geq b_{L'}$.

Once the routing has been performed by the QPPTW algorithm, the time-windows are readjusted (as discussed in Section 6.4.6) before the QPPTW algorithm is reapplied to route the next aircraft.

6.4.4 QPPTW Algorithm

The input of the QPPTW algorithm contains the graph $G = (V, E)$ with its weight function w_e, which corresponds to the taxi times for each edge, estimated using the taxi time estimation method which was described in Section 6.3. The sorted set of available time-windows $\mathcal{F}(e)$ also has to be provided for each edge e, specifying when the edge is available. A taxi request $T_i = (s_i, t_i, time_i)$ for aircraft i is then a conflict-free route R from the vertices s_i to t_i with minimal taxi time (w.r.t. w_e) that respects the given time-windows.

The pseudocode of the QPPTW algorithm is shown in Algorithm 1 and is a variant of the QPPTW algorithm described by Stenzel (2008). The main difference is that we allocate the labels to vertices, which helps both to model the process more realistically and to more easily understand the algorithm, since it distinguishes between the use of the labels at the vertices and the input time-windows at the edges.

In summary, the algorithm expends iteratively found quickest routes from the source to vertices in the network until it reaches the target by making sure that all the relevant time-window constraints are fulfilled. The expansion steps of the algorithm work similarly to Dijkstra's algorithm. The main feature of the QPPTW is the ability to take into account when which edge is free or blocked by another aircraft. The complexity of the algorithm is higher and the dominance rules for two labels have to be extended.

Lines 1 and 2 of Algorithm 1 involve the initialization of the Fibonacci heap and the references to this heap which are stored at each vertex. The use of Fibonacci heaps for this algorithm

Algorithm 1: Quickest Path Problem with Time Windows (QPPTW)

Input: Graph $G = (V, E)$ with weights w_e for all $e \in E$, the set of sorted time-windows $\mathcal{F}(e)$ for all $e \in E$, a taxi request $T_i = (s_i, t_i, time_i)$ with the source vertex $s_i \in V$, the target vertex $t_i \in V$ and the start time $time_i$.

Output: Conflict-free route R from s_i to t_i with minimal taxi time that starts at the earliest at time $time_i$, respects the given time-windows $\mathcal{F}(e)$ or returns the message that no such route exists.

```
1  Let H = ∅
2  Let ℒ(v) = ∅   ∀v ∈ V
3  Create new label L such that L = (s_i, [time_i, ∞), nil)
4  Insert L into heap H with key time_i
5  Insert L into set ℒ(s_i)
6  while H ≠ ∅ do
7      Let L = H.getMin(), where L = (v_L, I_L, pred_L) and I_L = [a_L, b_L]
8      if v_L = t_i then
9          Reconstruct the route R from s_i to t_i by working backwards from L
10         return the route R
11     forall the outgoing edges e_L of v_L do
12         foreach F^j_{e_L} ∈ ℱ(e_L), where F^j_{e_L} = [a^j_{e_L}, b^j_{e_L}], in increasing order of a^j_{e_L} do
13             /*Expand labels for edges where time intervals overlap*/
14             if a^j_{e_L} > b_L then
15                 goto 11 /*consider the next outgoing edge*/
16             if b^j_{e_L} < a_L then
17                 goto 12 /*consider the next time-window*/
18             Let time_in = Maximise(a_L, a^j_{e_L})    /*a^j_{e_L} > a_L ⇒ waiting*/
19             Let time_out = time_in + w_{e_L}
20             if time_out ≤ b^j_{e_L} then
21                 Let u = head(e_L)
22                 Let L' = (u, [time_out, b^j_{e_L}], L)
23                 /*dominance check*/
24                 foreach L̂ ∈ ℒ(u) do
25                     if L̂ dominates L' then
26                         goto 12 /*next time-window*/
27                     if L' dominates L̂ then
28                         Remove L̂ from H
29                         Remove L̂ from ℒ(u)
30                 Insert L' into heap H with key a_{L'}
31                 Insert L' into set ℒ(u)
32 return "there is no s_i-t_i route"
```

has the same beneficial effect upon the execution time as it does for Dijkstra's algorithm. The starting label is generated for the source s_i in line 3 and is then inserted into the Fibonacci heap, which is sorted with respect to the earliest possible arrival time (key). A reference is maintained to this label using the $\mathcal{L}(s_i)$ set for each vertex. These references are used as a look-up by the dominance check in lines 23-29, where the algorithm needs fast access to all of the labels associated with a particular vertex.

In each iteration of the while loop, the algorithm checks whether the Fibonacci heap still contains elements. If this is not the case, there is no route which can be enlarged and, therefore, no route from s_i to t_i, starting at $time_i$, exists (line 32). If the Fibonacci heap still contains elements, the algorithm takes a minimal element with respect to the key (line 7), checks whether this label already represents a route to the target t_i (lines 8-10) and, if not, tries to expand the associated route.

The route can usually continue along a number of different outgoing edges from any vertex and can potentially use different time-windows on each edge (lines 11 and 12). In order to use an edge, there must be a time-window available with an overlapping time interval, as expressed by the conditions on lines 14 and 16. The earliest possible point in time that edge e_L can be exited is identified (lines 18 and 19) and the expansion step is executed. When the condition stated in line 20 is true, a new label will be generated (lines 21 and 22). Different cases are possible at this stage. Firstly, the new label may dominate another label (line 27), in which case the dominated label will be erased (lines 28 and 29). Secondly, the new label may be dominated by an older one (line 25), in which case it is not necessary to take this label into account (line 26). The while loop is executed as long as there is a route which can be expanded. Once a route R to the target t_i has been found, the route can be generated by working backwards through the set of labels (line 9) using the references, $pred_L$, to the previous labels.

This generalised vertex-based Dijkstra's algorithm is a variant of that given by Stenzel (2008). His proof that the edge-based algorithm solves the problem in polynomial time (in the number of time-windows) will also hold for this algorithm.

6.4.5 Modifications to the QPPTW Algorithm for Airport Ground Movement

Algorithm 1 is used for arriving aircraft as described above, since their goal is to clear the runway and reach the gate/stand as quickly as possible. In our model, departing aircraft aim to reach the runway in time for their predetermined take-off time and leave the gate/stand as late as possible in order to do so. This allows for more of the waiting time to be absorbed at the gate/stand when the engines are not running. The same algorithm is used for this purpose, computing the route backwards, with the end time fixed instead of the start time and with changes to reverse the time-related steps. Since the algorithm logic remains unchanged, this modified algorithm has not been presented here.

In an attempt to further speed up the execution time of the algorithm, we applied goal-oriented search (Sedgewick and Vitter 1986) to the QPPTW algorithm. Two heuristic measures were investigated for estimating lower bounds for the rest of the partial route: firstly the Euclidean distance was used to measure the linear distance to the target, and secondly the remaining time was estimated using Dijkstra's algorithm to compute the time which would be needed ignoring any interference from other aircraft. Unfortunately, neither approach resulted in a valuable speed-up when applied to this problem. This can possibly be explained by the fact that the graph representing the airport layout is sparse (having on average only a few outgoing edges for each vertex) and routes often start on the border of the graph (see Figure 3.2), so the number of expansions exploring non-promising areas of the airport is relatively small already.

6.4.6 Readjustment of the Time-Windows

When an aircraft has been routed, the time-windows have to be readjusted according to the edge utilisation of the adopted route R, and the edges which conflict with these. It is necessary to consider edge conflicts only during this stage and not during the routing process (Algorithm 1).

Algorithm 2 presents the pseudocode for the readjustment of the time-windows. The input consists of the weighted graph $G = (V, E)$, the set of conflicting edges $confl(e)$ for all $e \in E$,

Algorithm 2: Readjustment of the time-windows

Input: Graph $G = (V, E)$ with weights w_e for all $e \in E$, the route R with reservations $\left[time_f^{in}, time_f^{out}\right]$ for all $f \in R$, the set of sorted time-windows $\mathcal{F}(e)$ for all $e \in E$ and the set of conflicting edges $confl(e)$ for all $e \in E$.

Output: Sorted set of time-windows $\mathcal{F}(e)$ including the reservations of the route R

```
1  foreach f ∈ R do
2      foreach e ∈ confl(f) do
3          foreach F_e^j = [a_e^j, b_e^j] ∈ F(e) do
4              if time_f^out ≤ a_e^j then
5                  goto 2    /*time-window is too late*/
6              if time_f^in < b_e^j then
7                  /*otherwise time-window is too early*/
8                  if time_f^in < a_e^j + w_e then
9                      if b_e^j − w_e < time_f^out then
10                         Remove F_e^j from F(e)
11                     else
12                         /*shorten start of time-window*/
13                         F_e^j = [time_f^out, b_e^j]
14                 else
15                     if b_e^j − w_e < time_f^out then
16                         /*shorten end of time-window*/
17                         F_e^j = [a_e^j, time_f^in]
18                     else
19                         /*split time-window*/
20                         F_e^j = [a_e^j, time_f^in]
21                         Insert [time_f^out, b_e^j] into set F(e)
```

the set of sorted time-windows $\mathcal{F}(e)$ for all $e \in E$, and the route R which was found for the most recent aircraft to be routed. The output is the new sorted set of time-windows $\mathcal{F}(e)$, including the reservations of the new route R.

In summary, all the affected edges are considered one by one and their time-windows are readjusted according to four cases (remove time-window, shorten at the start or the end, respectively, or splitting the time-window). The main features are the specific distinctions of the cases and the procedure to consider all possibly affected edges.

Basically, the algorithm determines which other edges are blocked for each edge of the route R (lines 1 and 2). All affected time-windows on these edges are adjusted (lines 3-7) and four different cases then have to be considered, depending upon the relative positions of the time-windows. The remaining time-window may be removed (lines 9-10) if it becomes too short to allow an aircraft to taxi; be shortened at the start (lines 11-13) or shortened at the end (lines 15-17); or it could be split in two smaller windows (lines 18-21).

Once a route has been allocated to an aircraft, some additional waiting time may be required on edges, beyond the time required to traverse the edge as specified by the time intervals on the labels from Algorithm 1. Time intervals on adjacent edges often overlap sufficiently that there is a choice of which edge the wait can be assigned to. In our implementation, the waiting times are forced to be as late in the corresponding part of the route as possible, apart from the initial waiting time for departures, which is allocated so as to maximise the stand hold. Alternative approaches could use this flexibility to select better and smoother speed profiles for the aircraft. Using a similar approach to that used in Lesire (2010), the aim could be to spread the necessary waiting times for an aircraft in such a way that the speed profiles are as "engine friendly" as possible. Although the effects of such postprocessing are not studied within this chapter, Chapter 7 analyses fuel efficient taxiing.

6.4.7 Buffer Times

The solutions of the approach are conflict-free routings, but it is possible for small delays to affect the entire plan. Buffer times would enable small deviations from the taxi times to be absorbed. To achieve such buffer times the label intervals in the algorithm are lengthened in

the desired direction (before or after) by a certain amount. To reflect growing uncertainties along the route, the amount of time can be made distance-dependent. Buffer times could also depend upon the expected congestion at the time, being increased when delays were expected to be more likely, although at these times the introduction of a buffer time would be more likely to reduce throughput.

6.4.8 Initial Sequencing of Taxiing Aircraft

The order in which aircraft are considered by the sequential routing algorithm can potentially affect the efficiency of the routing. The natural sequencing, of considering aircraft in the order in which they become available, has advantages in terms of perceived fairness and has been adopted in the past (Busacker and Fricke 2002). A more advanced approach using a concept of collaborative virtual queues was presented in Burgain et al. (2009), with the idea being to limit the number of aircraft which were taxiing on the surface to a specified maximum and maintaining a virtual queue of those waiting to start, forcing them to wait until the count allows them to pushback. The natural ordering (the expected wheel-on time on the runway for arrivals and the expected earliest pushback time at the gate/stand for departures) was adopted by default for this chapter, but the potential benefits of using better sequences have also been considered, as explained in the next section.

6.5 Heuristics for Finding Better Aircraft Sequences

The aim of this section is to introduce heuristics which are used to improve the quality of the utilised aircraft sequence. Gotteland et al. (2001) applied the concept of genetic algorithms to attempt to find better orderings. A major drawback of such an approach is that there is no control of the final sequence and a lot of communication between controllers in the tower and pilots is potentially needed to change the routes of all of the affected aircraft as the situation changes.

Our approach attempts to balance the additional communication between controllers and pilots and reduce the total taxi time. The concept of a ‘causer aircraft’ is introduced first in

this section, based on ideas from Ravizza's Master's dissertation (Ravizza 2009). Afterwards, different heuristics are explained in order to improve the solution quality, as far as reduction in total taxi time is concerned, while staying close to the original natural sequencing, to maintain an element of fairness.

6.5.1 Finding a Causer Aircraft

The QPPTW algorithm sequentially routes new aircraft whilst respecting previous reservations by other aircraft. The time needed by each aircraft to complete its route is compared to the time which would have been needed if the aircraft had been routed in isolation (using Dijkstra's algorithm (Dijkstra 1959; Cormen et al. 2001) to find the shortest route). If the difference is bigger than a certain threshold value then the algorithm attempts to find a better sequence. This delay will always have been caused by an already routed aircraft and this aircraft is classified as the causer aircraft. If several aircraft are affecting an aircraft, the one affecting the current aircraft's route the earliest is classified as the causer aircraft.

There are two cases to consider when detecting a causer aircraft. Firstly, an aircraft can need to wait during taxiing because another aircraft is blocking the next part of the route and thus causes a delay. Secondly, an aircraft could be forced to do a detour to avoid a wait, leading to a delay which is longer than the threshold value. In this case, the computed route is compared to the shortest route and from the separation point on, a look-ahead mechanism on the shortest route is used to determine the causer aircraft. The blocking of a part of the taxiway can potentially be further on the shortest route, since the QPPTW algorithm finds a way to detour which leads to the destination the fastest, so a detour may diverge earlier than the blocker position if this leads to a shorter route.

6.5.2 Swap Heuristic

The simplest (but very effective) heuristic involves using the swap-operator. As explained before, the aircraft are initially sequenced in the natural ordering. If a route of a new aircraft has a delay longer than the threshold value, this approach tests another sequence and uses the

better one. In the case of the swap heuristic, the route of the causer aircraft is taken out of the solution and the new aircraft is then routed and scheduled based on the QPPTW algorithm, before re-routing the causer aircraft. All of the other routes and schedules are fixed in order to maintain fairness and to aim for reduced communication requirements.

Tests were also performed to investigate the potential benefits of using the swap-operator and also allowing all of the other aircraft's routes to be changed. First, the final sequence found by the approach was used to run the QPPTW algorithm and quantify the benefit. Then, after swapping two aircraft in the sequence, the approach re-routed all of the intermediate aircraft and tested whether this lead to a reduced total taxi time compared to adding the new aircraft to the end of the old sequence.

6.5.3 Shift Heuristic

A shift-operator is used here instead of a swap-operator. In contrast to the previous heuristic, the new aircraft is added just before the causer aircraft in the aircraft's sequence. Obviously, all of the aircraft afterwards may have to be re-routed to find a feasible overall solution of the problem.

6.5.4 Best-shift Heuristic

Both of the above heuristics aim for a better overall solution by considering routing the new aircraft earlier than the causer aircraft. Hence, the concept of a causer aircraft is the main idea behind the improvements. This heuristic works in a different way and is based on the concept of Constrained Position Shifting (CPS) (Dear 1976; Dear and Sherif 1991). CPS allows the shifting of an aircraft by at most a predefined number of positions in the sequence. All of the insert positions which meet the CPS are explored for a new aircraft in our heuristic and the best is chosen. Again, all of the aircraft after the new position may have to be re-routed to guarantee a feasible solution.

6.5.5 Off-line Heuristic

To provide a baseline for all of the online heuristics discussed so far, the potential sequences were explored using an off-line approach. An initial sequence was used and swap- and shift-operators were randomly applied to delayed aircraft to find a better sequence, using a hill-climbing approach: the new sequence replaced the old sequence if the new sequence had a better overall quality.

6.6 Results and Discussions

This section starts with a table collating the key results, to ease comparison. The explanation of the results will follow. The results of the taxi time estimation which was presented in Section 6.3 are then discussed. An analysis of the results from the ground movement decision support system, which was described in Section 6.4, is provided next, followed by more results considering the different heuristics to improve the solution quality.

Table 6.2: Summary of the results

	Total taxi time [s]	Average taxi time per aircraft [s]
Actual taxi time	2489262.0	443.5
Fuzzy rule-based system		
Total taxi time estimation	2458400.4	438.0
Total taxi time estimation (unimpeded)	1685798.5	300.3
QPPTW algorithm with FCFS		
Using unimpeded taxi time estimates	1736020.9	309.3

The relevant results are summarised in Table 6.2. The first row of results, labelled "Actual taxi time", shows the actual total and average taxi times for dataset "ZRH 2011", including the queuing time at the runway.

The taxi time function, which was developed, was applied to each aircraft to estimate the taxi times and the results are shown in the next two rows, under the heading "Fuzzy rule-based system". In the first case, the function was applied assuming the actual traffic level and we note that the difference between the predicted and actual times is less than 2%. In the second case, the traffic-related components of the function were zeroed (as discussed in Section 6.6.1),

to estimate the taxi times if there had been no delays due to other aircraft, and the difference illustrates the amount of the taxi time which was a result of such delays.

The unimpeded taxi times were also used within the QPPTW algorithm, based on a first-come-first-served (FCFS) ordering of the aircraft. The total and mean resulting taxi times are shown in the table under the heading "QPPTW algorithm with FCFS". These results are analysed and explained further in the following two sections.

6.6.1 Analysis of Taxi Time Estimation

Once the pushback duration had been included in the Mamdani fuzzy rule-based system (see Section 6.3.1), the coefficient of determination R^2 of 94.15% showed that the FRBS was able to explain the variability of the taxi time data very well for the real world Zurich dataset.

The fitted FRBS model was then used to predict a taxi time for each aircraft in the dataset, with and without the factors which represented the effects of the delays due to other aircraft (see Section 6.3.1). The results can be seen in Table 6.2. The model predicts that 31.4% of the taxi time was related to delays due to other aircraft, including delays in queues behind other aircraft at the runway. There would be an average saving of 137.7s per aircraft if these delays could be eliminated. The influence of the interactions between the aircraft which lead to the waiting times is analysed in the next section.

6.6.2 Experimental Details Using the QPPTW Algorithm

The framework was programmed in Java as a single-threaded application and executed on a personal computer (Intel Core 2 Duo, 3GHz, 2GB RAM). In these experiments, all aircraft were allowed to use all of the taxiways and only intersecting and adjacent edges were considered to be in conflict and were, therefore, not allowed to be used by two aircraft simultaneously. The buffer time (Section 6.4.7) was set to zero. An analysis of different buffer times showed that the taxi time would have been enlarged by only a linear factor of the buffer time. Similar results were also found in Ravizza (2009).

Extensive analysis was performed using the QPPTW algorithm, with a FCFS consideration

sequence for aircraft, to solve the ground movement problem using the data from and layout of Zurich Airport (dataset "ZRH 2011"). The aircraft were routed sequentially using the taxi speed estimations from the Mamdani FRBS which was discussed in Sections 6.3 and 6.6.1. The resulting total taxi times can be found in Table 6.2, where the taxi times used were those which were estimated for unimpeded aircraft (ignoring the influence of factors related to other aircraft on the surface), the average taxi time (including re-routing and waiting delays) was 309.3s per aircraft.

The estimations of the unimpeded taxi times from the Mamdani FRBS prediction approach provide a lower bound for the taxi times, since they assume no re-routing delays or queuing behind other aircraft. The QPPTW algorithm is designed to predict the delays which are actually necessary due to the interactions between aircraft for the specific routings and timings which the algorithm assigns to aircraft. Comparison of the resulting taxi times from the QPPTW algorithm against the lower bound reveals an increase in the taxi time from 1685798.5 to 1736020.9 seconds, showing that the additional taxi times for the re-routing and waiting summed to 50222.4s over the entire week, an increase of around 3% in the total taxi time. The 3% increase over the lower bound (rather than optimal) times indicates that its use as a ground movement decision support system seems very promising for this problem.

It is also interesting to compare the approach described here against the actual performance of the airport on this particular week of operation. Data from Zurich Airport reports a total taxi time of 2489262.0s. Comparison with the results for the QPPTW algorithm with unimpeded taxi time estimation highlights potential maximum savings of about 30.3%, an average of 134.2s per aircraft. This only indicates an upper bound for the potential savings, since the real times will include slack time for the departures at the runway to ensure a high runway throughput.

The solution time to solve the entire week of operation with 5613 aircraft was 216887ms, an average solution time of 39ms per aircraft. This supports the potential use of the algorithm in an online decision support system. No infeasible solution occurred within any of the executions of the simulation. These findings are consistent with earlier work by Atkin et al. (2011b), using another dataset from Zurich Airport ("ZRH 2007") and taxi times which were generated from the linear regression approach in Chapter 4.

6.6.3 Analysis of Different Ordering Heuristics

The heuristics were first tested on the smaller datatset from Zurich Airport ("ZRH 2007") and the best heuristic was then used for improving the results on the larger dataset. The threshold value to accept a small delay was set very low, to 5 seconds.

Table 6.3: Analysis of ordering heuristics

	FCFS	swap	shift	best-shift	off-line
Difference from lower bound	4391s	2771s	2494s	2450s	2305s
Reduction of gap	0.0%	36.9%	43.2%	44.2%	47.5%
Approximation ratio	1.022	1.014	1.012	1.012	1.011
Solution time	11.6s	60.1s	13.2min	49.8h	-
Solution time per aircraft	17.1ms	88.5ms	1.2s	4.4min	-

All of the relevant results are summarised in Table 6.3. The columns categorise the ordering heuristics which were used: FCFS ordering, swap heuristic, shift heuristic, best-shift heuristic with a maximal position shift of 25 (because it is highly unlikely that a bigger limit would lead to significant improvements but would increase the computational time even more) and finally the solution from the off-line heuristic (starting from the best solution found by the other approaches).

As reported in our initial paper (Atkin et al. 2011b), the total taxi time for the FCFS ordering applied to the dataset "ZRH 2007" is 207723 seconds and a lower bound of the problem is 203332 seconds, implying that the optimality gap is at most 4391 seconds, with an approximation ratio of 1.022.

The results for the different approaches for improving the solution quality were ordered by their complexity. It can be seen that a reduction of 36.9% of the gap between the initial solution and the lower bound was found by applying the swap heuristic. Further improvements were found when using any of the other approaches, but these were surprisingly small. The solution times for ordering an entire sequence were in the opposite order. The swap heuristic needed more time due to the fact that the approach first had to check whether a route had any delay and then find the causer aircraft before trying the swapped sequence. The two shift heuristics needed much longer since all of the intermediate aircraft had to be re-routed, which would also imply more communication for the pilots and controllers.

The off-line approach used the sequence which resulted from the best-shift heuristic as the input. The presented solution was found after 3320 iterations of swap- or shift-operators, which corresponded to around 22 hours of calculation. Another additional 20000 iterations did not improve the solution any further and it is very likely that the approach had found a local optimum.

Results are not shown for the variations of the swap heuristic which were previously discussed since none had a better reduction in the taxi time than the other approaches in comparison to their solution time and the number of aircraft affected.

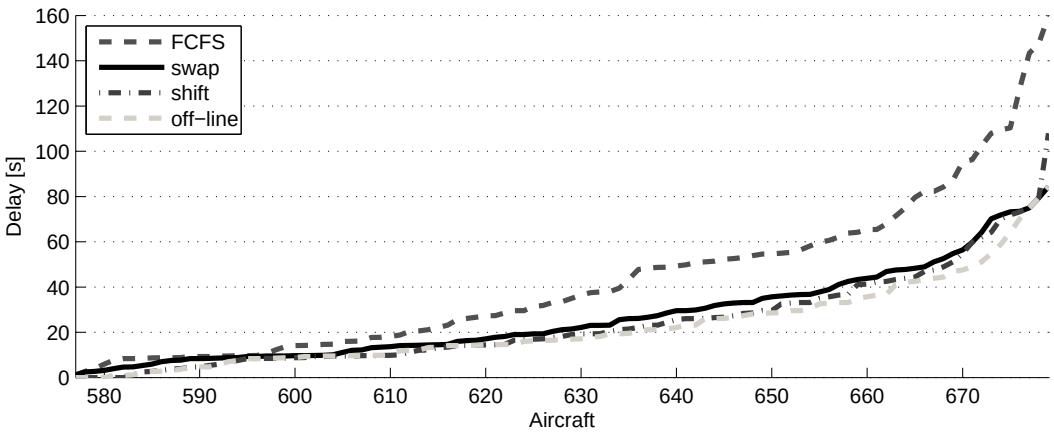

Figure 6.3: Sorted delay for each aircraft from the different heuristics

Figure 6.3 shows the sorted individual delays for the aircraft that resulted from the different heuristics. Since both shift heuristics lead to very similar lines, the best-shift heuristic is not presented. In all approaches, at least 577 aircraft were routed by the algorithm without any delays and are not included in the figure. It can be seen that the heuristics can greatly improve the solution and that the simple but effective swap heuristic reduced the longest delay from 160 seconds to 84 seconds.

6.6.4 Studies of a Swap Heuristic for the Larger Dataset

Table 6.4 provides a comparison of the routing and scheduling algorithm with and without the swap heuristic for the larger dataset "ZRH 2011". The different columns represent the different days in the dataset and the total for the entire week. The first three rows of the table report the

number of aircraft movements during each day and it can be seen that the airport has lighter traffic at the weekend (day 6 and day 7). Rows two and three differentiate between departures (DEP) and arrivals (ARR). The second block shows the results of the QPPTW algorithm with the FCFS order (without the swap heuristic) and the third block shows the results with the swap heuristic. The lower bound was computed using the estimated taxi times but with each aircraft routed in isolation, so no waiting times or detours were included. The following block shows the absolute gap between the lower bound and the results for the FCFS and the swap heuristic, respectively. The reduction in the gap is the relative improvement from using the swap heuristic compared with the FCFS ordering.

The results were similar for the different days and the total taxi times were approximately double for departures compared to arrivals, independent of the sequencing method. Obviously, the introduction of the swap heuristic increased the solution time per aircraft, however, the algorithm is still fast enough to be used in an online environment.

The swap heuristic based sequencing method was able to reduce the gap between the routing which was found and the lower bound by 30% on average over the entire week, with a bigger reduction rate for departing aircraft (33%) than arriving aircraft (25%).

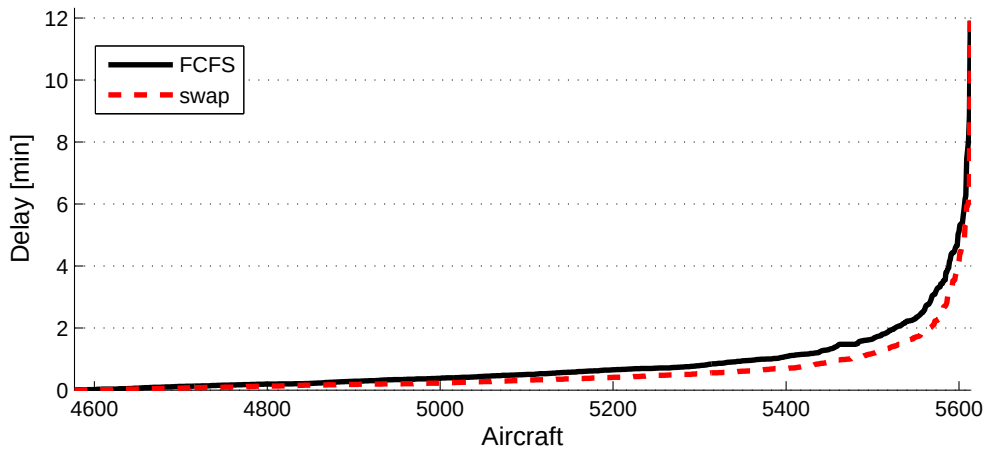

Figure 6.4: Sorted delay for each aircraft with and without swap heuristic

The sorted individual delays for the aircraft, resulting from the analysis with and without the swap heuristic are shown in Figure 6.4. In both cases, at least the first 4578 (out of 5613) aircraft had no delays in their planned schedules and are not included in the figure. The delays

Table 6.4: Analysis of routing and scheduling algorithm with and without swap heuristic

		Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	Total
# Aircraft	Total	818	806	781	839	825	757	787	5613
	DEP	407	405	392	416	421	379	387	2807
	ARR	411	401	389	423	404	378	400	2806
FCFS	Total taxi time [s]	251231.3	248751.7	244449.9	256497.1	256662.5	234632.2	243796.2	1736020.9
	Total taxi time DEP [s]	168731.8	168503.6	172718.4	172070.7	183031.6	163713.7	168059.2	1196829.0
	Total taxi time ARR [s]	82499.5	80248.1	71731.5	84426.4	73630.8	70918.5	75737.0	539191.9
	Solution time [ms]	33640	32562	28765	33078	30483	28859	29500	216887
	Solution time per aircraft [ms]	41	40	37	39	37	38	37	39
Swap heuristic	Total taxi time [s]	249355.0	246128.0	242097.1	253869.1	254924.1	233204.3	241216.3	1720793.8
	Total taxi time DEP [s]	167382.5	166201.3	171258.8	170221.0	181683.2	162652.0	166168.5	1185567.4
	Total taxi time ARR [s]	81972.5	79926.7	70838.3	83648.1	73240.8	70552.2	75047.7	535226.4
	Solution time [ms]	123919	109248	110685	117701	101373	89311	101248	753485
	Solution time per aircraft [ms]	151	136	142	140	123	118	129	134
Lower bound	Total taxi time [s]	243699.3	240061.6	237926.9	248121.3	250165.6	228864.3	236911.6	1685750.5
	Total taxi time DEP [s]	163890.8	161979.0	169068.7	166165.2	178146.0	159547.7	163469.9	1162267.3
	Total taxi time ARR [s]	79808.5	78082.7	68858.1	81956.1	72019.5	69316.6	73441.6	523483.2
FCFS gap	Total taxi time [s]	7532.1	8690.1	6523.0	8375.8	6496.9	5767.9	6884.7	50270.4
	Total taxi time DEP [s]	4841.0	6524.7	3649.7	5905.5	4885.6	4166.0	4589.3	34561.7
	Total taxi time ARR [s]	2691.0	2165.4	2873.4	2470.3	1611.3	1601.9	2295.4	15708.7
Swap heuristic gap	Total taxi time [s]	5655.7	6066.4	4170.2	5747.8	4758.5	4340.0	4304.7	35043.3
	Total taxi time DEP [s]	3491.7	4222.3	2190.1	4055.8	3537.2	3104.4	2698.6	23300.1
	Total taxi time ARR [s]	2164.0	1844.1	1980.2	1692.0	1221.3	1235.6	1606.1	11743.2
Reduction of gap	Total taxi time	25%	30%	36%	31%	27%	25%	37%	30%
	Total taxi time DEP	28%	35%	40%	31%	28%	25%	41%	33%
	Total taxi time ARR	20%	15%	31%	32%	24%	23%	30%	25%

from Figure 6.4 are summarised in Table 6.5 in a numerical way, showing the percentage of aircraft, which have more than a certain amount of delay. The swap heuristic was able to improve most of the percentages by almost a factor of 2. Again, these results are consistent with Section 6.6.3 based on the older dataset from the same airport.

Table 6.5: Percentage of aircraft with more than a certain amount of delay

	Without swap heuristic	With swap heuristic
Have any delay	18.47%	18.31%
More than 1min	4.22%	2.51%
More than 2min	1.57%	0.86%
More than 3min	0.80%	0.48%
More than 4min	0.45%	0.27%
More than 5min	0.27%	0.12%

6.6.5 Scenarios with more Ground Traffic

New scenarios were generated based on the data from the summer of 2011, simulating more ground traffic at Zurich Airport. The analysis focused upon Monday as a representative day. Each movement of an aircraft was duplicated and the copy was shifted by 30 minutes to generate the scenario with 200% ground traffic. For the 300% scenario each movement was duplicated twice and one copy was shifted by 15 minutes and the other copy by 30 minutes. The scenarios for the settings with 120%, 140%, 160% and 180% were generated by randomly removing some of the duplicated aircraft movements from the 200% case and the scenarios between 200% and 300% were created by randomly removing movements from the second duplication. It has to be noted that within this analysis the focus was entirely upon analysing the ground movement problem with more ground traffic and, obviously, separations and deadlines were considered for neither taking-off nor landing (since the runway throughput would not be achievable), nor was it guaranteed that no overlaps occurred in the gate allocations. The aim is to consider only whether the algorithm can cope with increased traffic load, and if so to determine the size of the consequent delays which would be allocated to aircraft.

Table 6.6 shows the results of the analysis. Each column represents a scenario with the appropriate amount of ground traffic related to the actual setting. The table is structured similarly

Table 6.6: Analysis of the routing and scheduling algorithm with and without swap-heuristic with artificially more ground traffic

		100%	120%	140%	160%	180%	200%	220%	240%	260%	280%	300%
FCFS	Total taxi time [s]	251231	304523	354337	406332	463119	522593	581120	640659	736357	858409	929010
	Total taxi time DEP [s]	168732	204889	235528	267120	306219	350662	390517	432088	506596	607601	657527
	Total taxi time ARR [s]	82500	99634	118809	139212	156900	171931	190604	208571	229761	250808	271483
Swap-heuristic	Total taxi time [s]	249355	301591	349673	401258	456006	513862	570208	638429	715518	827778	887346
	Total taxi time DEP [s]	167383	202594	232713	264669	302068	345455	384747	435042	492151	585126	624224
	Total taxi time ARR [s]	81973	98997	116961	136589	153938	168407	185461	203387	223367	242652	263122
Lower bound	Total taxi time [s]	243699	293209	339086	385438	435692	487399	537638	586888	633717	682267	731098
	Total taxi time DEP [s]	163891	197666	226342	254709	289150	327782	362454	397385	427691	459642	491672
	Total taxi time ARR [s]	79808	95543	112744	130729	146542	159617	175184	189504	206026	222626	239425
FCFS gap	Total taxi time	3%	4%	4%	5%	6%	7%	8%	9%	16%	26%	27%
Swap-heuristic gap	Total taxi time	2%	3%	3%	4%	5%	5%	6%	9%	13%	21%	21%
Reduction of gap	Total taxi time	25%	26%	31%	24%	26%	25%	25%	4%	20%	17%	21%

to Table 6.4 to ease comparison. It can be seen that the lower bound increases linearly which is due to the construction of the problems. The numbers also show an approximately linear increase of the approach which was based on the FCFS consideration of aircraft until the ground traffic reached the 240% level. After that the gap between the QPPTW algorithm without the swap-heuristic increased from values between 3% to 9% before that to values between 16% and 27% after it. The swap-heuristic achieved an average of a 22% reduction in the gap between the lower bound and the QPPTW algorithm with FCFS ordering. This was relatively consistent for the scenarios with lower traffic and higher traffic levels. The only exception was the 240% scenario, where the reported reduction of the gap was only 4%. The implementation of the swap heuristic was, therefore, generally worthwhile.

6.6.6 Further Use for Simulations

The main purpose of this chapter is to enhance decision support systems which can be used in control towers. Nevertheless, a prototype of this approach could also be used for simulations of management or operational strategies. From an airport point of view several kinds of analysis would be possible. A taxiway layout could be analysed to highlight where the bottlenecks are and by how much the operations are restricted if a part of the network is blocked, such as for maintenance requirements. Airports often have a concept of where certain aircraft should be routed and variations of such concepts could also be tested by either restricting certain combinations of taxiway parts or by favouring certain combinations. Furthermore, a ground movement simulation could be integrated with runway sequencing or gate assignment to perform a broader analysis.

Airlines could also use simulations to better understand the situation at an airport. This could improve their own operations. For instance, they could be used to identify which times of the day are less likely to cause waiting times. Airlines could then adjust their schedules to improve the operational performance, assuming that the other carriers maintain their existing schedules. A good example of such a study is “Delta’s Operation Clockwork” (Petroccione 2007). De-peaking of their operations at Hartsfield-Jackson International Airport was able to save waiting times for aircraft equivalent to adding nineteen aircraft into the fleet, which

were then re-inserted into the system to provide more connections. A test phase confirmed the findings of the analysis, but the airline decided to revert to the old schedule afterwards due to reductions in revenue from reduced passenger demand.

Simulation has been widely used by research groups and software vendors to get insights into airport operations and to evaluate the impacts of uncertainties. Rosenberger et al. (2000, 2002) presented a stochastic model for airline operations within the SIMAIR project, with the primary purpose being to evaluate crew scheduling plans and recovery policies. Simulation tools for airport and airspace operations, such as SIMMOD from the Federal Aviation Administration (FAA), RAMS from Eurocontrol, DPM from Sabre and TAAM from the Preston Group, can model existing and planned operations very well, but may lack in the area of automatically improving operations which can be performed with optimisation systems.

6.6.7 Impact of Results

This section highlights the possible savings in fuel costs of the introduced algorithm by using the same approach as in the analysis by Brinton et al. (2011). In our analysis with the integration of taxi time estimation, the QPPTW algorithm and the swap heuristic, an average aircraft had 306.6 seconds of fuel burn instead of the 443.5 seconds which was reported in the historic data. The saving of 136.9 seconds per aircraft movement accumulates to around 637000 minutes per year, based on the 279000 movements as was reported at Zurich Airport in 2011. Brinton et al. (2011) based their calculations on a jet aircraft using 25 pounds of fuel per minute while taxiing, which fits the guidelines from ICAO for the settings of a “Single Aisle Jet”. With an assumed $4 US per gallon of fuel, the annual cost savings in fuel at Zurich Airport would be approximately $9.6 million. It should be noted that other sources question the actual fuel rate for taxiing, which is possibly overestimated by ICAO (Morris 2005; Kim and Rachami 2008).

6.7 Conclusions

This chapter described a more realistic and potentially more environmentally friendly ground movement decision support system, compared to previous approaches. The overall framework

is designed to combine the runway sequencing problem and ground movement problem, aiming for better global solutions, although only the ground movement element was considered in this chapter. This work extends the basic ground movement problem of minimising the travel times by including the concept of absorbing possible waiting times for departures at the gate/stand, to reduce the fuel burn and environmental impact. The sequential QPPTW algorithm which was described here is based on graph theoretical concepts and can include restrictions such as limitations on which taxiways aircraft can use, which taxiways block which and when, and any turning limitations at taxiway junctions. In addition, the algorithm provides the opportunity to add buffer times for blocking the reserved taxiways for longer than expected, to absorb small delays and schedule disturbances.

Experiments used data for an entire week of operations at Zurich Airport, the largest hub airport in Switzerland. This data was used to generate more accurate taxi time estimations for each aircraft, using a taxi time prediction function which was generated from an extensive statistical analysis and a fuzzy rule-based system, applied to the same dataset. These taxi time estimations were then utilised within the QPPTW algorithm to route and schedule the ground movement. The results are very promising and show potential maximum savings in total taxi time from using the decision support system described here, in conjunction with the taxi time prediction system, of about 30.3%, compared to the actual performance at the airport. Further research is necessary to determine the amount of buffer time and runway delay which should be utilised to account for any remaining taxi time uncertainty and avoid starving the runways.

The experimental results of the developed decision support approach show average solution times of only a few milliseconds per aircraft, and are, therefore, adequate for the implementation of such a system for real-time use at airports.

The potential benefits of applying different ordering heuristics for the sequential ground movement problem were also explored. The most promising approach was to use a simple but effective swap-operator. The quality of the solution was shown to be substantially improved with comparatively little additional computational time, making it still suitable for real-time use at airports. Very few changes are needed in the initial sequence, hence the communication between controllers and pilots is kept to a minimum.

We intend to investigate various extensions of this work in future, in addition to the combination of the ground movement problem with the runway sequencing problem. Firstly, the QPPTW algorithm enables the possible waiting times to be spread in different ways. In this chapter, they were allocated so as to maximise the stand hold time and to better adapt to schedule disturbances, but an alternative approach would be to develop smoother speed profiles for aircraft, using the engine in a more efficient and environmentally friendly way. Secondly, we would like to perform a similar analysis for different airport layouts, to better understand the effects of the layout upon the best solution approach, but it will be necessary to obtain more data and support from other airports in order to do so.

7

Trade-off Analysis between Taxi Time and Fuel Consumption in Airport Ground Movement

> A new idea comes suddenly and in a rather intuitive way. But intuition is nothing but the outcome of earlier intellectual experience.
>
> Albert Einstein

7.1 Introduction

Air transportation represents a growing sector and this trend is predicted to continue for the foreseeable future. However, there are increasing concerns, from a wide range of stakeholders, about the environmental impact of the sector. Aircraft ground movement is an operation which is particularly affected by these two conflicting trends. With the increase in aircraft movements, it is likely that hub airports, especially, will form bottlenecks for air transportation. The ground movement problem plays a key role in addressing the goal of reducing delays. It is important

to note that lower accelerations may sometimes be more fuel efficient, even though movement times may be increased. An ambitious goal stated in the report of the High Level Group on Aviation Research for the European Commission attempts to have emission-free aircraft movements when taxiing in the year 2050 (European Commission 2011).

The details of the ground movement problem vary depending upon the aims of the airport but it can be summarised as the goal of producing conflict-free routings for aircraft on the airport's surface, usually from gates/stands to runways and vice versa. A variety of different constraints and objective functions have been used in the literature (see Chapters 2 and 6). Previous research on ground movement often focused on minimising the total taxi time (Pesic et al. 2001; Marín 2006; Roling and Visser 2008; Atkin et al. 2011b) or minimising the makespan (García et al. 2005; Herrero et al. 2005). Multi-objective approaches have also been used. In addition to minimising the total taxi time, penalising deviations from a scheduled time of departure/arrival has been considered (Smeltink et al. 2004; Balakrishnan and Jung 2007; Deau et al. 2009). Gotteland et al. (2003) investigated penalising deviations from a departure time interval. Other research has employed a weighted linear objective function to simultaneously consider the total routing time, the delays for arrivals and departures, the number of arrivals and take-offs, the worst routing time and the number of controller interventions (Marín and Codina 2008). Although multi-objective approaches have been employed, we have not found any research focusing on the integration of objectives related to the environmental impact.

There is little coverage of the environmental considerations of the taxi operations within the current research literature. The main focus has been upon stand holding, which shifts waiting times for aircraft from the runway queue back to the gate, in order to reduce fuel burn (Burgain et al. 2009; Atkin et al. 2010a, 2011a). The assumption made by them was that by reducing the total taxi time, one can simultaneously improve the efficiency of airport operations and reduce the fuel consumption. However, as indicated in Chen and Stewart (2011), this may not be true for all cases or airports, since the detailed relationship between fuel consumption and the corresponding speed profile was not investigated in previous research. Atkin et al. (2010b) suggested the value of considering speed profiles when routing aircraft in order to avoid unnecessary fuel burn due to acceleration and deceleration. Lesire (2010) applied a

postprocessing stage in his routing approach to smoothen the speed profiles. Similar ideas have also been presented by Cheng and Sweriduk (2009). Finally, Chen and Stewart (2011) presented an approach to analyse the trade-off between taxi time and fuel consumption for a single trajectory of an unimpeded aircraft.

In this research, we analyse the trade-off between the total taxi time and the fuel consumption for the conflict-free routing problem for aircraft on an airport's surface. In contrast to the approach of Chen and Stewart (2011), the interactions between multiple aircraft are also considered. These interactions affect the speed profiles of the aircraft involved and massively increase the solution space of the routing approach. Hence, a sophisticated new procedure had to be developed to make such an analysis possible.

A related problem can be found in energy-efficient running time control for metro lines. For example, Binder and Albrecht (2012) recently presented a predictive dynamic control system to save energy for the Hamburg metro system. Furthermore, Bektas and Laporte (2011) introduced a new vehicle routing problem (VRP) variant, called Pollution-Routing Problem (PRP), which takes pollution into account.

This chapter is structured as follows: Section 7.2 presents the case which was analysed, then the newly developed multi-objective approach for analysing the trade-off between taxi time and fuel consumption is detailed in Section 7.3. The results of the application of the algorithm to the dataset are shown in Section 7.4; before the chapter ends with conclusions in Section 7.5.

7.2 Problem Details

Different approaches for fuel burn estimation are considered in this section, together with details about the settings which are used for maximal speeds and acceleration.

Data for an entire week's operations was utilised for this research (dataset "ZRH 2011", see Section 3.2.1 for more details). The information was used to represent the entire airport layout as a directed graph, where the edges represent the taxiways and the vertices represent the junctions or intermediate points (Figure 7.1 illustrates a part of this graph and Figure 3.2 the entire airport).

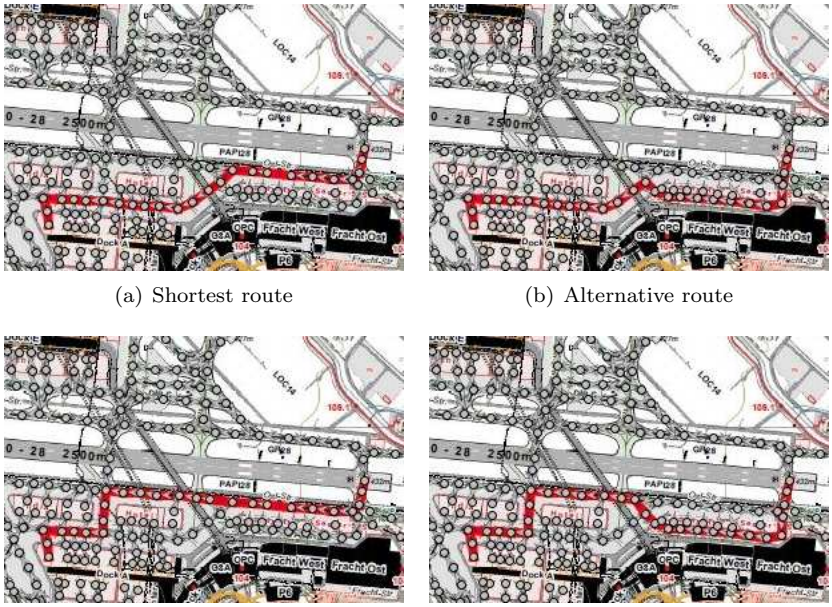

(a) Shortest route (b) Alternative route

(c) Alternative route (d) Alternative route

Figure 7.1: Different routes from pier A to runway 28 at ZRH

The explanation of the categorisation, which is used for aircraft, is discussed in Section 3.3.

7.2.1 Fuel Consumption, Taxi Speed and Acceleration

As is common practice, the International Civil Aviation Organization engine emissions database (ICAO 2008) has been used for estimating the fuel consumption of aircraft. It states that the engine power setting for taxi/ground idle is 7% of full rated power but does not distinguish between the different phases of taxiing. This setting was also used by the FAA (2005a,b) and Simaiakis and Balakrishnan (2010). Morris (2005) showed that levels of around 5% to 6% are more realistic for most engine types and Kim and Rachami (2008) also stated that values below 7% are more likely. A newer approach by Nikoleris et al. (2011) and Jung et al. (2011) used a set of four different values for different taxi operation phases: 4% for idle thrust, 5% for taxiing at a constant speed or brake thrust, 7% for perpendicular turn thrust and 9% for breakaway thrust. In their study about air quality and public health impacts of UK airports, Stettler et al. (2011) used a setting of 4-7% (a uniform random distribution with a mean of 5.5%) for taxiing (for maintaining a constant speed, decelerating, or holding) and a setting of 7-17% (a

triangular distribution with mode of 10%) for taxiway acceleration. Based on the results by Wey et al. (2006), they stated that the fuel flow of the engines is approximately proportional to the engine thrust setting. Khadilkar and Balakrishnan (2011, 2012) presented an approach to estimate fuel burn using linear regression. They concluded that the total taxi time is the main factor, although the number of acceleration events was also a significant factor. Our analysis has approached the problem using a physics-based model which is introduced later, in Section 7.3.1. We do not consider single engine taxiing in this chapter (Deonandan and Balakrishnan 2009).

Different researchers have worked with different taxi speed settings. Rappaport et al. (2009) showed, using quantitative analysis, that the average speed on straight taxiways (29.4 km/h) was higher than the average speed during turns (23.2 km/h) at Detroit Metropolitan Wayne County Airport (DTW) in Michigan, USA. Cassell and Evers (1998) reported that 95% of aircraft taxi at less than 30 knots (around 56 km/h) and the average speed was found to be 10 knots (around 19 km/h) during turns. This setting was also used in the work by Chen and Stewart (2011), where the maximal speed during taxiing was set to 30 knots (around 56 km/h) and the speed during turns to 10 knots (around 19 km/h). The same setting has been applied in this research, where a turn is considered to be when an aircraft has to make a change of direction of more than 30 degrees on a part of a taxiway. The maximal acceleration and deceleration is here set to $0.1 \cdot g$, to ensure passenger comfort, as in the latter reference, where $g = 9.81$ m/s^2 is the acceleration due to gravity.

It is assumed in this analysis that the airport has no significant taxiway slopes. It is also assumed that there is no heavy wind, which affects the fuel burn of aircraft and that no drag or lift considerations are needed in the model for estimating fuel consumption.

7.3 Methodology

The focus of this research is entirely on the ground movement part of the airport operations of aircraft. In addressing this, the pushback/landing time of aircraft are as specified by the dataset and are assumed to be fixed.

This section first introduces the objective functions, before an overview of the developed integrated procedure is given. Two key elements of the procedure are presented in separate sections afterwards.

7.3.1 Objective Functions

This research analyses the trade-off between taxi time and fuel consumption in airport ground movement. The first objective function aims to minimise the total taxi time (including waiting times during taxiing) combined with moving possible waiting times to the gate where the engines are not running. The second objective function aims to minimise fuel burn. As in the research by Chen and Stewart (2011), a fuel consumption index is used. This penalises high acceleration rates during taxiing and uses a physics-based model. In essence, the formula for the force of acceleration is assumed to be given by $F_a = m \cdot a_p$, (Newton's second law of motion) where a_p is the acceleration of an aircraft during a phase p and m is its weight. The rolling resistance F_r is then also taken into account (see Section 3.3 for the formula and values). The fuel consumption index is defined as the sum of the force of acceleration plus the rolling resistance, multiplied by the time for which it was applied. If the sum of the force of acceleration plus the rolling resistance is negative, due to deceleration in a phase, then the sum is set to zero, since aircraft need fuel to accelerate or taxi with constant speed but cannot recover fuel while decelerating. The trade-off between the two objective functions for an example taxi route is shown in Figure 7.2.

7.3.2 Integrated Procedure

A routing approach was developed, based upon the algorithm presented in Chapter 6, utilising the trade-off information gleaned from the algorithm proposed in Chen and Stewart (2011). It is a sequential, vertex-based, label-setting algorithm working on a graph representing the airport's surface. Since two conflicting objective functions are considered, the approach has to be enhanced by using an adapted version of the algorithm, in a sophisticated integrated procedure.

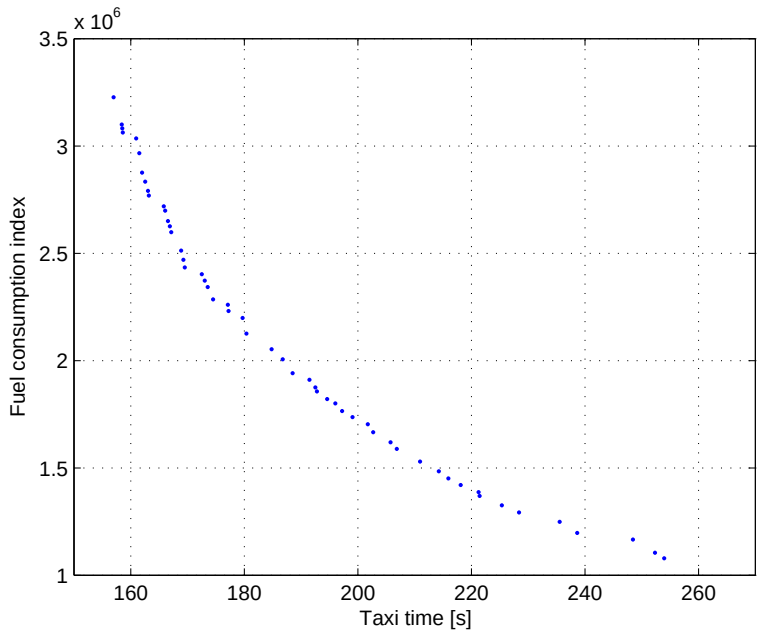

Figure 7.2: Pareto-front of unimpeded taxi trajectories

The general idea of this procedure was proposed by Climaco and Martins (1982), whose aim was to develop a shortest path algorithm for finding the Pareto-front of optimal paths for two criteria. The objective functions which they used were minimising the total time and minimising the cost of the path, where each edge had two values assigned to it. Their method generates a sequence of k shortest paths with respect to the first objective function, until the path with the minimal value with respect to the second objective function is obtained, leading to a Pareto-front of all optimal paths.

Our problem differs from the problem which Climaco and Martins (1982) were facing in two main points. Firstly, not all edges are available at all times since other aircraft are also travelling around the airport and will block some parts of the taxiways at certain times. Secondly, the second objective function cannot be evaluated with a simple Dijkstra's algorithm for finding the shortest path (Dijkstra 1959) in this situation, but needs a more elaborate method due to its non-additivity.

In summary, Algorithm 3 generates sequentially several feasible routes for each aircraft and picks the one with the desired trade-off between taxi time and fuel consumption. Considering all

Algorithm 3: Integrated procedure for trade-off analysis

```
1  Sort all flights by pushback/landing time
2  foreach objective function discretisation i ← 1 to l do
3      foreach aircraft a do
4          Find the best k routes w.r.t. minimal taxi times using the k-QPPTW algorithm
5          foreach route k of aircraft a do
6              Approximate the Pareto-front of both objectives, using the population adaptive
               immune algorithm (PATT-PAIA)
7          Generate the combined Pareto-front for the source-destination pair of aircraft a
8          Discretise this Pareto-front into l roughly equally spaced points
9          Select the ith point and reserve the relevant route for aircraft a
10     Save the accumulated values for all aircraft for both objective functions for the global
       Pareto-front
11 Output: Global discretised Pareto-front
```

aircraft with the same desired trade-off, one possible schedule is found which is then represented as a point in the global discretised Pareto-front (see Figure 7.4). As an input, the details of the aircraft are needed together with the layout of the airport. The output can be used to better understand the mentioned trade-off.

Algorithm 3 shows the proposed integrated procedure at a glance. The approximation of the global Pareto-front is generated in a discretised way due to the complexity of the problem. The parameter l defines the number of generated points on the global Pareto-front approximation. In each iteration of the outer loop (lines 2-10), the objective values are generated for both objectives, starting with the most time-efficient solution then incrementally changing to the most fuel-efficient solution. For each outer loop, the entire set of aircraft has to be scheduled. The algorithm routes and schedules the flights sequentially and is based on an initial sequencing (line 1) by pushback/landing times of all aircraft. Different (adaptive) sequencing methods could be used, as was done by Ravizza and Atkin (2011), but this was not investigated here.

The first subroutine (line 4) finds the best k routes for aircraft a related to the total taxi time. In doing so, reservations of already routed aircraft have to be taken into account. The

k-Quickest Path Problem with Time Windows (k-QPPTW) was developed for this purpose and is explained in more detail in Section 7.3.3. A possible set of generated routes can be seen in Figure 7.1.

The second subroutine (line 6) analyses each of the k routes independently. A population adaptive immune algorithm (PATT-PAIA, see Section 7.3.4) approximates the Pareto-front of different speed trajectories for aircraft a on a particular route, complying with the unblocked time-windows for each edge and the detailed speed behaviours of this aircraft. A more detailed description of this subroutine is given in Section 7.3.4 and an example of the output can be seen in Figure 7.2. It should be noted that also other multi-objective evolutionary algorithms could be used and that the decision to use the proposed algorithm was due to the fact, that we had access to an implementation which was already tailored to this particular problem.

The subroutine in line 7 combines the k different Pareto-fronts and selects, with the same dominance rules as in the PATT-PAIA, the global Pareto-front for a given source-destination pair of aircraft a (see Figure 7.3). The resulting Pareto-front is discretised into l points, as equally spaced as possible (line 8). The approach aims to split the border of the Pareto-front between the most time-efficient and most fuel-efficient solutions into equally spaced segments and always selects the closest non-dominated point to each of the ends of these segments. Line 9 selects the ith point (according to the outer loop of the algorithm) out of the l ordered representative points. In addition, the detailed route associated with this point is fixed for this aircraft and the scenario is changed in such a way that upcoming aircraft cannot use the same parts of the taxiways at the same time.

The inner loop (lines 3-9) is repeated until all of the aircraft from the dataset have been routed and the total taxi time and the total fuel consumption can be accumulated to generate a single point in the global Pareto-front (line 10). Obviously, before repeating the outer loop (line 2) all of the changes which have been made to the reservations of the aircraft have to be reversed, since the scenario is then evaluated for a different objective function discretisation.

Since the subroutine on line 6 is comparatively time consuming, the procedure could be parallelised for this stage and executed on a cluster of processors.

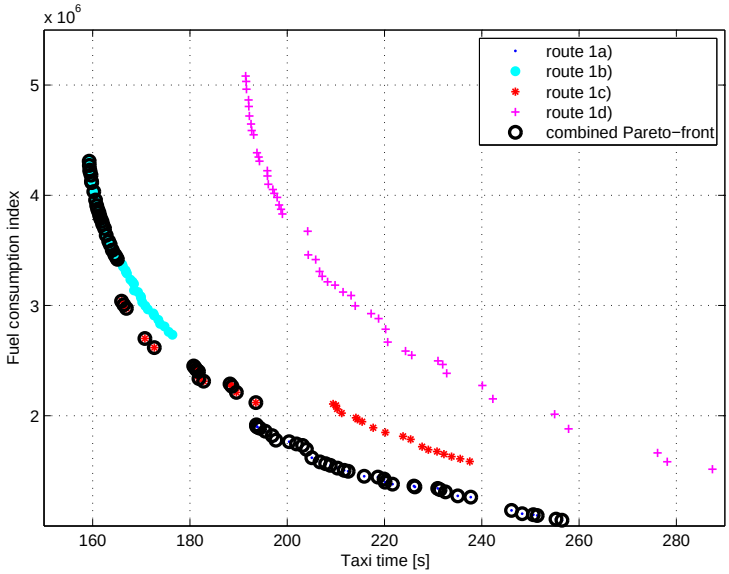

Figure 7.3: Combined Pareto-front from four different routes which are shown in Figure 7.1 having associated time-windows

7.3.3 Sequential K-QPPTW

Schüpbach and Zenklusen (2011) recently showed that a simplified version of the conflict-free routing problem for a group of n vehicles is NP-hard, even when the underlying graph is a path, using a reduction from the 3-partitioning problem. Thus the approach, which was discussed in Chapter 6 and Atkin et al. (2011b), was used based upon a sequential routing of the aircraft. This vertex-based, label-setting algorithm works on a graph representing the airport's surface, does not need any time discretisation, respects reservations of parts of the taxiways for previously routed aircraft and can compute a route very quickly.

The Quickest Path Problem with Time Windows (QPPTW) algorithm was extended to form a k-QPPTW algorithm by adapting it to not only generate the "best" (in this setting, fastest) route, but a set of the k best solutions. This extension is based upon the ideas of Yen (1971) and Lawler (1972). Yen (1971) introduced an algorithm to find the k shortest loopless paths in a network, where the computational upper bound of the algorithm only increases linearly with the value k. The main idea behind the approach is that the $(j+1)$th path can only deviate from the root of one of the best j paths in one vertex. Hence, it is only necessary to look for

all shortest deviations from the best j paths and then select the deviation which has the best objective value for the entire path.

The QPPTW algorithm was similarly adapted and generates the best route in the conventional way. It then iterates until it has found the best k routes. In each iteration j, it generates all deviations from the $(j-1)$th best routes which are different from routes which have already been found. The jth best route is then the best one of all of these routes which has not already been identified as one of the $(j-1)$ best routes. To speed up the entire algorithm and to minimise storage space, only subroutes need to be stored along with the information about which route they are deviating from, instead of storing the entire route.

7.3.4 Planning Aircraft Taxiing Trajectories via a Population Adaptive Immune Algorithm (PATT-PAIA)

Chen and Stewart (2011) proposed an immune inspired multi-objective search algorithm which utilised a physics-based aircraft dynamic model to search for different taxiing trajectories for a given route. Each of these trajectories represents a different trade-off between taxi time and fuel consumption. This algorithm has been extended in this research to incorporate time-window constraints (collaborative work between the University of Lincoln and the University of Nottingham and Jun Chen made the adjustments to his MATLAB code). In the following, the PATT-PAIA is briefly discussed. Interested readers are referred to Chen and Mahfouf (2006) and Chen and Stewart (2011) for more details.

Algorithm 4 shows the proposed PATT-PAIA with time-window constraints at a glance. As discussed in Section 7.2, the entire airport layout is represented as a directed graph. Time-windows, corresponding to edges between the vertices, represent when a part of a taxiway is not used by any other aircraft. Unlike the vertices shown in Figure 7.1, PATT-PAIA only considers junctions and divides the entire taxi route of an aircraft into segments. Each of these segments may contain several intermediate vertices. There are two types of segments, namely straight segments and turning segments. The maximum speed for a straight segment is 30 knots (around 56 km/h) and the speed during turns is fixed to 10 knots (around 19 km/h). For a straight segment, there are four consecutive transitional phases for an aircraft:

Algorithm 4: PATT-PAIA

1 Approximate the most time-efficient speed profile for each segment satisfying all given time-windows
2 Randomly generate additional initial solutions in its neighberhood
3 Objective evaluation and non-dominated sorting
4 **foreach** of the *Gen* iterations **do**
5 Fitness evaluation: evaluate the fitness of the candidate solutions
6 Selection and cloning: selection of good solutions based on the fitness values; selected solutions are cloned
7 Mutation: variation of clones
8 Constraint handling: check if the mutated solutions meet all time-windows and feasible bounds
9 Objective evaluation and reselection: reselect good solutions from the combined solutions based on non-dominated sorting
10 **Output:** Approximation of Pareto-front of taxi trajectories

a) acceleration phase, b) constant speed phase, c) deceleration phase and d) fast deceleration phase. By adjusting the acceleration and deceleration rates and the switching points between the phases, one can obtain different speed profiles and their corresponding fuel consumption indices using the aircraft categorisation (see Section 3.3) and the second objective function (see Section 7.3.1). For a turning segment, no optimisation is needed, since the speed is fixed.

The algorithm is a tailored adaptive immune algorithm with the following main features: fitness evaluation, selection and cloning, mutation, constraint handling and objective evaluation. As an input, the path of an aircraft is given and time-window constraints.

To obtain a good approximation of the Pareto-front of both objectives, PATT-PAIA is devised as follows. First, an initial population pool is randomly generated around a feasible solution (trajectory) which fulfils all of the time-window constraints (lines 1 and 2). This feasible solution is generated using a heuristic to find the most time-efficient trajectory which takes into account a more realistic speed profile. Non-dominated sorting will then be utilised to distinguish between dominated and non-dominated solutions (line 3). The loop (lines 4-9) iterates *Gen* times to improve the current Pareto-front and aims to have the Pareto-front equally spread. One of the good solutions will be randomly selected in order to calculate its distance from the rest of the

solutions, which defines the fitness of each solution (line 5). Based upon the fitness values, good solutions will be selected to be cloned with a higher probability (line 6). The cloned solutions will be mutated with small variation steps, to locally search the neighbourhood (line 7). On the other hand, bad solutions will only be cloned once and will be subjected to greater mutation in order to explore more of the search space. Constraint handling is used each time, immediately after new solutions are generated, to check whether the mutated solutions are still within the feasible bounds, as discussed by Chen and Stewart (2011), and also to calculate the arrival time at each vertex to see whether the current solution still complies with the given time-windows (line 8). The mutated solutions and the previous solutions will be combined and passed to a reselection stage, so that the best solutions survive into the next generation (line 9). The output of the algorithm is an approximation of the Pareto-front for a given route of an aircraft within the given time-windows (line 10).

The feasible solution which is generated using the heuristic in line 1 not only speeds up the search of the PATT-PAIA, but it also guarantees at least one feasible solution at the end. A possible improvement to spread the solutions more equally between the most time-efficient route and the most fuel-efficient route, would be to also generate the most fuel-efficient trajectory.

7.4 Results and Discussion

The following section is divided into three parts. It starts with the visualisation of the global Pareto-front, continues with an analysis across an entire week's operations and finishes with a sensitivity analysis related to a different objective function.

7.4.1 Global Pareto-front

Analysis showed that, by using the three best routes ($k = 3$, in Algorithm 3), the procedure can find very good approximations to the global Pareto-front. Similarly, the number of iterations which were needed by the PATT-PAIA to find a good approximation was also tested, and this value was then fixed to $Gen = 40$. The execution time of Algorithm 3 on a personal computer (Intel Core 2 Duo, 3GHz, 2GB RAM) with these settings was around 100 minutes

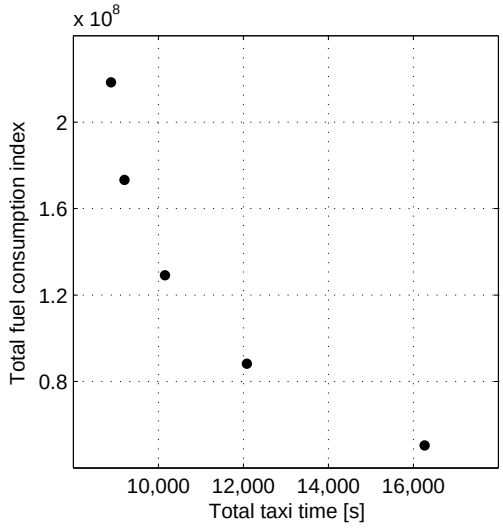

Figure 7.4: Global discretised approximation of the Pareto-front for 57 aircraft

for one data point on the global Pareto-front (inner loop of the algorithm), for a dataset of 57 aircraft. Due to the long execution time, the focus of this research was restricted to analysing the busiest time of the day, which was the hour between 11am and noon. Figure 7.4 shows the global Pareto-front with five discretised values for the "Monday" dataset. The point at the top indicates the analysis where each of the 57 aircraft was taxiing as time-efficiently as possible, whereas the furthest point to the right indicates the analysis where each aircraft was taxiing as fuel-efficiently as possible.

7.4.2 Analysis over a Week's Operations

An analysis was performed to see how consistent the results were over a week's operations of the hour between 11am and noon. For this purpose, only the two extreme cases were studied instead of analysing the entire global Pareto-front. Table 7.1 shows the results for the analysed week. The values are reported as the average values per aircraft, since the number of aircraft in the dataset varied between 46 and 63 over the week. The first column ("Mon") restates the extreme values from Figure 7.4 (averaged per aircraft). The last row of each block highlights the growth from the best solution to the other extreme value. The values seem to be relatively consistent and the dataset for Monday seemed to be a good representation of an average day.

Table 7.1: Analysis over a week's operations of the hour between 11am and noon with the focus upon the extreme values

	Mon	Tue	Wed	Thu	Fri	Sat	Sun
∅ **Taxi time**							
Time-efficient [s]	156	157	128	174	152	165	154
Fuel-efficient [s]	285	293	214	320	292	316	295
Growth	83%	87%	67%	84%	92%	91%	91%
∅ **Fuel cons. index**							
Time-efficient [$\times 10^3$]	3832	3291	3492	4002	3173	3754	3718
Fuel-efficient [$\times 10^3$]	884	762	742	922	700	849	823
Growth	334%	332%	371%	334%	353%	342%	352%
Number of aircraft	57	58	46	58	56	63	52

7.4.3 A Different Objective Function

A further experiment was run to see how sensitive the algorithm was to the fuel-related objective function. For this purpose, the setting from Stettler et al. (2011) was used as a replacement for the second objective function. As stated in Section 7.2.1, two different settings were used, one for acceleration and one for taxiing with constant speed, deceleration or holding. A stepwise function was utilised to measure the fuel used (in kg) during taxiing, based on the fuel flow coefficient. The parameters were set so that an aircraft burns 10% of the maximal fuel flow during acceleration and 5.5% when it is not accelerating (the PATT-PAIA still models segments with four transitional phases). With such a setting, the PATT-PAIA is encouraged to always accelerate with the maximal acceleration rate and mainly controls the length of the acceleration phases. Table 7.2 shows the results for the "Monday" dataset and is structured in the same way as Table 7.1, with the only difference being that the new fuel-related objective function is used instead.

Table 7.2: Analysis with the focus upon the extreme values where the fuel-related objective function was replaced in reference to the research by Stettler et al. (2011)

	Different objective function
∅ **Taxi time**	
Time-efficient [s]	155.5
Fuel-efficient [s]	156.7
Growth	0.8%
∅ **Fuel flow**	
Time-efficient [kg]	23.8
Fuel-efficient [kg]	23.5
Growth	1.2%

It can be seen that the trade-off analysis very much depends upon the fuel-based objective function which is used. With this approach, based on the research by Stettler et al. (2011), there seems to be hardly any difference between optimising an aircraft's trajectory for time-efficiency or fuel-efficiency. The actual values cannot be directly compared to those in Table 7.1, but, due to the fact that fuel flow of the engines is approximately proportional to the engine thrust setting (Wey et al. 2006, p. 7), the growth between the most extreme solutions indicates the potential for using a trade-off analysis.

Calculations highlight why the trade-off is sensitive. Using the example of the Airbus A320, the rolling resistance is 11.48 kN, which is around 5.1% of the total rated output. Hence, the two fuel-related objective functions behave very similarly in phases of constant speed, deceleration or during a hold. However, during acceleration with maximal acceleration rate, the physics-based objective function adds the rolling resistance and the force of acceleration which is 88.0 kN in total. This is around 39.3% of the total rated output, which is considerably more than the 10% of the other function. More research is needed to better understand in detail the fuel burn during taxiing, before the question can be answered as to whether there is actually a trade-off between time-efficient and fuel-efficient taxiing, and to quantify any potential trade-off.

7.5 Conclusions

A new model was developed to analyse the trade-off between two different, potentially conflicting, objective functions for the ground movement problem at airports. A sophisticated combination of two algorithms has enabled the development of a framework to run simulations for different datasets, and to perform sensitivity analysis. The first utilised algorithm finds the best possible routes for an aircraft at an airport and the second algorithm finds an approximation of the Pareto-front for different speed profiles for each of these routes, in relation to the given objective functions.

Historic data from Zurich Airport was utilised for this analysis. The objective functions consisted of the taxi time, which is a commonly used measure, and a physics-based function related to the force needed from the aircraft engines during taxiing. The results show that the integrated

procedure is able to tackle this hard problem in a comparatively efficient way. Furthermore, results seem to be consistent over several days. Sensitivity analysis has highlighted that the potential trade-off between the two objectives depends very much upon the actual modelling of the fuel-based objective function, which appears not to be fully understood at the moment.

Future research is needed for better understanding the details of the fuel usage during taxiing and the standard practices of pilots during taxiing. Such insights could then be used to show more clearly the influence of operational and environmental targets during the taxiing process.

8

Conclusions

> It is easier to measure something than to understand what you have measured.
>
> Anon.

8.1 Discussion

The central aim of this thesis is the support of ground movement operations of aircraft at airports from an operations research, statistics and data mining point of view. Two different areas were the focus of this research. On the one hand, algorithms were introduced and analysed for better predicting taxi times of arrivals and departures and a routing and scheduling algorithm is presented which needs very little execution time and facilitates stand holding. On the other hand, this thesis makes a contribution to the understanding of airport operations. In particularly, Chapters 4 and 5 identify the factors which are important for estimating taxi times and which regression approaches can find the most accurate predictions. Chapter 7 tackles a new research direction of understanding the trade-off between minimising taxi time and minimising fuel consumption for airport ground movement. This research had a clear focus of being practical for use at an airport and for being able to be integrated with other approaches to finally end up with a fully integrated decision support system for airside airport operations including

gate assignment, stand holding, ground movement and runway sequencing.

The analysis is mainly based on European hub airports, but Appendix D shows that some of the results can be similarly valid for other airports. Some parts of this thesis could be used already by any airport at which they have access to historic data, especially to improve their taxi time predictions. The routing and scheduling approach from Chapter 6 needs considerably more IT infrastructure in place. We believe that such a system could be extended in such a way that it does not only support controllers in the tower with suggesting routes, but that a system in the cockpit of each aircraft could eventually act as an autopilot on the ground.

The following sections highlight the key results of this PhD thesis and disclose some future research areas where we see potential.

8.2 Key Results

The key results are summarised here in the same order as the chapters within this thesis.

Ground movement can be supported. This thesis and research literature in general has repeatedly demonstrated that prototypes of decision support system are able to solve the daily operations which ground controllers are facing. In many instances, the results obtained can decrease the total taxi time and reduce the carbon footprint.

Integration is the future. Ground movement decision support systems not only have to be integrated with any collaborative decision making tool at an airport, but they should simultaneously optimise the integrated airside airport operations. In a first stage, such integration contains links with the stand holding problem, the gate assignment problem and the runway sequencing problem, so that all operations of an aircraft are considered between the landing phase and the take-off phase. In a second stage, en-route optimisation should also be integrated. Obviously, this starts to link the operations of various airports and greatly enlarges the complexity of the problem. For this reason, the focus of this study was on fast solution approaches, especially for the routing and scheduling algorithm, which can also deal with many real-world constraints.

Other needs in the area. Section 2.7 highlighted some of the challenges which are considered

in this thesis and some where we think it is important that further research should continue. One area with increasing relevance is the consideration of environmental issues during all phases of a flight, including taxiing on the ground. Benefits in this area are considered throughout the entire thesis, but especially in Chapter 7. Furthermore, it is important that academic models are dealing with the actual problems at an airport and not simplified scenarios. Therefore, any approach has to be able to integrate all of the relevant constraints at an airport, to handle uncertainty in an appropriate way, and to be robust in general.

Advanced taxi time prediction leads to improvements. Chapters 4 and 5 highlighted the benefits of improving taxi time prediction. Analysis showed that Zurich Airport could benefit greatly by applying more advanced models to predict taxi times, of the form which were introduced within this research. Moreover, two different areas are suggested which could benefit from such approaches. Firstly, many stakeholders at an airport need accurate predictions of all the operations to better plan their tasks and improved taxi time prediction could help them. Secondly, ground movement decision support systems need predictions for unimpeded taxi times as such systems consider interaction between aircraft explicitly. In contrast to many other publications, not only departures but also arrivals were included in the studies.

Better understanding of taxi times. Chapter 4 analysed which factors are significant when predicting taxi times. The models were tested considerably on different airports. A key requirement for improved taxi time prediction is the integration of the surface layout at an airport. This was often neglected in other approaches and is especially relevant for European airports where the runway queues are less dominant than at US airports. The important factors turned out to be the total distance, the turning angle, the differentiation between arriving and departing aircraft and the number of other aircraft on the ground. Moreover, the operational mode (which runways were in use) can have an effect on the taxi times.

Better model for taxi time prediction. It is not only important to understand the significant factors for prediction taxi times, but it is also fundamental to analyse which regression approach is best for doing so. Chapter 5 showed an extensive analysis of different regression models and concluded that TSK fuzzy rule-based systems provided the best results for the analysed airports.

Routing and scheduling can be done within milliseconds. A graph-based decision support system was introduced and tested which can route and schedule aircraft in a fraction of a second. Even though the approach can add buffer times within the model to make the schedule more robust, an instantaneous approach can deal with unexpected situations and is also able to be used in a more complex integrated framework.

Solution time does not jeopardise a realistic setting. Even though the mentioned approach is very fast, it can handle realistic constraints. First of all, the routing and scheduling approach utilises the information found from unimpeded taxi time prediction. Moreover, the approach can incorporate different constraints and can be tailored to the circumstances of the particular airport.

The swap-heuristic can improve the performance. By using a simple but effective heuristic, the quality of the solution can be further improved with comparatively little additional computational time. Very few changes in the initial sequence are needed. Hence, the communication between controllers and pilots is kept to a minimum.

Massive expected savings. The results are very promising and show potential maximum savings in total taxi time from using the introduced decision support system of about 30.3%, compared to the actual performance at the tested airport.

Framework to analyse trade-off between taxi time and fuel consumption. It is still an open question whether there is a trade-off between taxi time and fuel consumption for the ground movement problem. Chapter 7 presented a sophisticated combination of two algorithms to analyse this question. This new approach is able to tackle this hard problem in a comparatively efficient way. Sensitivity analysis has highlighted that the potential trade-off between the two objectives depends very much upon the actual modelling of the fuel-based objective function, which appears not to be well understood at the moment.

8.3 Future Work

A consequence of research is that by answering challenging questions, new questions will arise. The aim of this section is not to restate the individual future work sections of each chapter,

but to focus upon an overall direction where we see potential for more research.

Extraordinary occurrences. An area which was less discussed within this thesis is how to deal with extraordinary occurrences. Heavy wind and snow can affect the airside airport operations a lot. Changes in wind can often result in the fact that the airport has to change the operational mode. It is important that a decision support system enables such changes and can resolve the operation in real-time. During weather conditions with colder temperatures and snow, aircraft have to be de-iced. Obviously, this can affect the ground movement process, since often there are remote de-icing stations, so the routing of an aircraft has to be adjusted and the taxi time prediction has to be tuned to consider such operations.

Better understanding of fuel consumption. This thesis introduced a sophisticated framework to analyse the trade-off between taxi time and fuel consumption. Nevertheless, more knowledge about the fuel-based objective function is needed to conclusively answer the question of whether there is trade-off between the two objective functions and to quantify any trade-off.

Gate-waiting. Gate-waiting is the term used to describe the situation when an arriving aircraft has to wait until a gate becomes available, e.g. when the gate is currently occupied or blocked by another aircraft. Analysis suggested that this was not a major problem for the routing and scheduling approach at the tested airport. However, this can be an issue at other airports and it can be important to take this into account when modelling the operations of such an airport.

Runway crossing. As this thesis focuses upon the ground movement problem and is only partly integrated with the runway sequencing problem, runways are assumed to be free for crossing within this research. The routing and scheduling approach can be adjusted to also consider the information about when a runway is free to cross, but the effects of this should be further analysed in future work.

Integration. It was discussed many times within this thesis how important it is to integrate different airside airport operations in a future step, such as ground movement, stand holding, gate-waiting, gate assignment and runway sequencing for arrivals and departures. The presented approaches have focused on the linking part of this integration - the ground movement problem. As it was always intended to use these approaches later within an integrated model, it was

important to develop flexible and fast algorithms. As such, we think that we have succeeded in providing a solid basis for further integration within this area.

Improvements of algorithms. We are by no means stating that the presented algorithms could not be further improved. We think that this thesis has presented significant new work and that it has pushed the boundaries of the state-of-the-art in different areas related to ground movement at airports. Moreover, it has tried to bridge the gap between the academic world and the needs of industry. The study about the taxi time prediction especially can vary from airport to airport and it is possible that other factors should be considered at other airports and that the structure of the problem changes in such a way that the most suitable regression approach will be different. In case of the framework for the trade-off analysis, the second stage of the framework related to the detailed taxi trajectory could possibly be tailored to any new fuel related objective functions.

Extension to more airports. We tried to generalise the findings by applying them to different airports and we have been able to test some of the approaches at four different hub airports. Regardless, other airports potentially have other bottlenecks in their systems and it could be very interesting to apply similar approaches to other settings and to see whether they work or what has to be adjusted to deal with the given airport setting.

Testing at airports and fine-tuning. This research tried to tackle problems which are relevant to the air transportation sector and we hope that some of the ideas and concepts will be used in operations. Discussions with experts from the field guided us to incorporate the important constraints from the actual problems at an airport. Nevertheless, some issues can only be identified by testing the concepts at an airport directly. Furthermore, a fine-tuning phase would be needed to fully benefit from the presented approaches.

References

Aarts E, Korst J, Michiels W (2005) Search Methodologies, Springer, chap Simulated Annealing, pp 187–210, ISBN: 0387234608

Abdelghany A, Abdelghany K, Narasimhan R (2006) Scheduling baggage-handling facilities in congested airports. Journal of Air Transport Management 12(2):76–81, doi: 10.1016/j.jairtraman.2005.10.001

ACI EUROPE (2010) An outlook for Europe's airports - Facing the challenges of the 21st century. Tech. rep., Airports Council International Europe

Allen DM (1971) Mean square error of prediction as a criterion for selecting variables. Technometrics 13(3):469–475

Anagnostakis I, Clarke JP (2003) Runway operations planning: A two-stage solution methodology. Hawaii International Conference on System Sciences 3:79a, doi: 10.1109/HICSS.2003.1174196

Andersson K, Carr F, Feron E, Hall WD (2000) Analysis and modeling of ground operations at hub airports. In: Proceedings of the 3rd USA/Europe Air Traffic Management R&D seminar, Napoli, Italy

Artiouchine K, Baptiste P, Dürr C (2008) Runway sequencing with holding patterns. European Journal of Operational Research 189(3):1254–1266, doi: 10.1016/j.ejor.2006.06.076

Asco A, Atkin JAD, Burke EK (2011) The airport baggage sorting station allocation problem.

In: Proceedings of the 5th Multidisciplinary International Scheduling Conference (MISTA), Phoenix, AZ, USA, pp 419–444

Atkin JAD (2008) On-line decision support for take-off runway scheduling at London Heathrow airport. Phd thesis, The University of Nottingham, UK

Atkin JAD, Burke EK, Greenwood JS, Reeson D (2006) The effect of the planning horizon and freezing time on take-off sequencing. In: Proceedings of the 2nd International Conference on Research in Air Transportation (ICRAT), Belgrade, Serbia and Montenegro

Atkin JAD, Burke EK, Greenwood JS, Reeson D (2007) Hybrid metaheuristics to aid runway scheduling at London Heathrow Airport. Transportation Science 41(1):90–106, doi: 10.1287/trsc.1060.0163

Atkin JAD, Burke EK, Greenwood JS, Reeson D (2008a) A metaheuristic approach to aircraft departure scheduling at London Heathrow airport. Computer-aided Systems in Public Transport, Lecture Notes in Economics and Mathematical Systems 600:235–252, doi: 10.1007/978-3-540-73312-6_12

Atkin JAD, Burke EK, Greenwood JS, Reeson D (2008b) On-line decision support for take-off runway scheduling with uncertain taxi times at London Heathrow Airport. Journal of Scheduling 11(5):323–346, doi: 10.1007/s10951-008-0065-9

Atkin JAD, Burke EK, Greenwood JS, Reeson D (2009) An examination of take-off scheduling constraints at London Heathrow airport. Public Transport 1(3):169–187, doi: 10.1007/s12469-009-0011-z

Atkin JAD, Burke EK, Greenwood JS (2010a) TSAT allocation at London Heathrow: The relationship between slot compliance, throughput and equity. Public Transport 2(3):173–198, doi: 10.1007/s12469-010-0029-2

Atkin JAD, Burke EK, Ravizza S (2010b) The airport ground movement problem: Past and current research and future directions. In: Proceedings of the 4th International Conference on Research in Air Transportation (ICRAT), Budapest, Hungary, pp 131–138

Atkin JAD, Burke EK, Greenwood JS (2011a) A comparison of two methods for reducing take-off delay at London Heathrow Airport. Journal of Scheduling 14(5):409–421, doi: 10.1007/s10951-011-0228-y

Atkin JAD, Burke EK, Ravizza S (2011b) A more realistic approach for airport ground movement optimisation with stand holding. In: Proceedings of the 5th Multidisciplinary International Scheduling Conference (MISTA), Phoenix, AZ, USA, pp 146–160

Atkin JAD, Burke EK, Ravizza S (2011c) A statistical approach for taxi time estimation at London Heathrow Airport. In: Proceedings of the 10th Workshop on Models and Algorithms for Planning and Scheduling Problems (MAPSP), Nymburk, Czech Republic, pp 61–63

Atkin JAD, Burke EK, De Maere G, Greenwood JS (2012) Addressing the pushback time allocation problem at Heathrow airport. Transportation Science (advance online publication), doi: 10.1287/trsc.1120.0446

Atkins S, Brinton C, Jung Y (2008) Implication of variability in airport surface operations on 4-d trajectory planning. In: Proceedings of AIAA Aviation Technology, Integration, and Operations (ATIO) Conference, Anchorage, AK, USA

Azeker M, Noyan N (2012) Stochastic optimization models for the airport gate assignment problem. Transportation Research Part E: Logistics and Transportation Review 48(2):438–459, doi: 10.1016/j.tre.2011.10.008

Baik H, Trani AA (2008) Framework of a time-based simulation model for the analysis of airfield operations. Journal of Transportation Engineering 134(10):397–413, doi: 10.1061/(ASCE)0733-947X(2008)134:10(397)

Baik H, Sherali HD, Trani AA (2002) Time-dependent network assignment strategy for taxiway routing at airports. Transportation Research Record: Journal of the Transportation Research Board 1788:70–75, doi: 10.3141/1788-09

Balakrishna P (2009) Scalable approximate dynamic programming models with applications in air transportation. PhD thesis, System Engineering & Operations Research Department, George Mason University, USA, pp. 3.D.3-1–3.D.3-12

Balakrishna P, Ganesan R, Sherry L (2008a) Airport taxi-out prediction using approximate dynamic programming: Intelligence-based paradigm. Transportation Research Record: Journal of the Transportation Research Board 2052:54–61

Balakrishna P, Ganesan R, Sherry L, Levy B (2008b) Estimating taxi-out times with a reinforcement learning algorithm. In: Proceedings of the IEEE/AIAA 27th Digital Avionics Systems Conference (DASC), doi: 10.1109/DASC.2008.4702812

Balakrishna P, Ganesan R, Sherry L (2009) Application of reinforcement learning algorithms for predicting taxi-out times. In: Proceedings of the 8th ATM R&D Seminars, Napa, USA

Balakrishna P, Ganesan R, Sherry L (2010) Accuracy of reinforcement learning algorithms for predicting aircraft taxi-out times: A case-study of Tampa Bay departures. Transportation Research Part C: Emerging Technologies 18(6):950–962, doi: 10.1016/j.trc.2010.03.003

Balakrishnan H, Chandran BG (2010) Algorithms for scheduling runway operations under constrained position shifting. Operations Research 58(6):1650–1665, doi: 10.1287/opre.1100.0869

Balakrishnan H, Jung Y (2007) A framework for coordinated surface operations planning at Dallas-Fort Worth International Airport. In: Proceedings of the AIAA Guidance, Navigation, and Control Conference, Hilton Head, SC, USA

Ball M, Barnhart C, Nemhauser G, Odoni A (2006) Air transportation: Irregular operations and control. In: Laport G, Barnhart C (eds) Handbooks in Operations Research & Management Science: Transportation, Handbooks in Operations Research and Management Science

Barnhart C, Belobaba P, Odoni AR (2003) Applications of operations research in the air transport industry. Transportation Science 37(4):368–391

Barth T, Pisinger D (2012) Scheduling of outbound luggage handling at airports. In: Klatte D, Lüthi HJ, Schmedders K (eds) Operations Research Proceedings 2011, Springer Berlin Heidelberg, pp 251–256, doi: 10.1007/978-3-642-29210-1_40

Beasley J, Krishnamoorthy M, Sharaiha Y, Abramson D (2004) Displacement problem and dynamically scheduling aircraft landings. Journal of the Operational Research Society 55(1):54–64, doi: 10.1057/palgrave.jors.2601650

Bektas T, Laporte G (2011) The pollution-routing problem. Transportation Research Part B: Methodological 45(8):1232–1250, doi: 10.1016/j.trb.2011.02.004

Bennell JA, Mesgarpour M, Potts CN (2011) Airport runway scheduling. 4OR: A Quarterly Journal of Operations Research 9:115–138, doi: 10.1007/s10288-011-0172-x

Bhadra D, Knorr D, Levy B (2011) Benefits of virtual queuing at congested airports using ASDE-X: A case study of JFK airport. In: Proceedings of the 9th USA/Europe Air Traffic Management Research and Development Seminar, Berlin, Germany

Bianco L, Dell'Olmo P, Giordani S (1999) Minimizing total completion time subject to release dates and sequence-dependent processing times. Annals of Operations Research 86:393–415

Bianco L, DellOlmo P, Giordani S (2006) Scheduling models for air traffic control in terminal areas. Journal of Scheduling 9:223–253, doi: 10.1007/s10951-006-6779-7

Binder A, Albrecht T (2012) Predictive energy-efficient running time control for metro lines. In: Proceedings of the Conference on Advanced Systems for Public Transport (CASPT12), Santiago, Chile

Böhme D, Brucherseifer R, Christoffels L (2007) Coordinated arrival departure management. In: Proceedings of the 7th USA/Europe Air Traffic Management Research and Development Seminar, Barcelona, Spain

Bolender MA (2000) Scheduling and control strategies for the departure problem in air traffic control. PhD thesis, University of Cincinnati, USA

Bonakdar L, Etemad-Shahidi A (2011) Predicting wave run-up on rubble-mound structures using M5 model tree. Ocean Engineering 38(1):111–118, doi: 10.1016/j.oceaneng.2010.09.015

Box GE, Jenkins GM, Reinsel GC (2008) Time series analysis: Forecasting and control, 4th edn. Wiley series in probability and statistics, Wiley, ISBN: 978-0470272848

Box GEP, Cox DR (1964) An analysis of transformations. Journal of the Royal Statistical Society Series B (Methodological) 26(2):211–252

Boysen N, Fliedner M (2011) Scheduling aircraft landings to balance workload of ground staff. Computers & Industrial Engineering 60(2):206–217, doi: 10.1016/j.cie.2010.11.002

Brinton C, Krozel J, Capozzi B, Atkins S (2002) Improved taxi prediction algorithms for the surface management system. In: Proceedings of AIAA Guidance, Navigation, and Control Conference and Exhibit, Monterey, CA, USA

Brinton C, Provan C, Lent S, Prevost T, Passmore S (2011) Collaborative departure queue management: An example of airport collaborative decision making in the United States. In: Proceedings of the 9th USA/Europe Air Traffic Management Research and Development Seminar, Berlin, Germany

Brucker P (2007) The job-shop problem: Old and new challenges. In: Proceedings of the 3th Multidisciplinary International Scheduling Conference (MISTA), Paris, France, pp. 15–22

Burgain P, Feron E, Clarke JP (2009) Collaborative virtual queue: Benefit analysis of a collaborative decision making concept applied to congested airport departure operations. Air Traffic Control Quarterly 17(2):195–222

Busacker T, Fricke M (2002) A contribution to proactively planning and managing airport ground traffic using a stochastic modelling approach. In: Proceedings of the 23rd Congress of International Council of the Aeronautical Sciences (ICAS), Toronto, Canada, pp. 771.1–771.7

Caimi G, Chudak F, Fuchsberger M, Laumanns M, Zenklusen R (2011) A new resource-constrained multicommodity flow model for conflict-free train routing and scheduling. Transportation Science 45(2):212–227, doi: 10.1287/trsc.1100.0349

Carpenter M, Stroiney S (2012) Managing gate and ramp operations to reduce delay, fuel burn, and costs. In: Proceeding of the Integrated Communications, Navigation and Surveillance Conference (ICNS), pp. H3-1–H3-22, doi: 10.1109/ICNSurv.2012.6218407

Carr F, Evans A, Clarke JP, Feron E (2002) Modeling and control of airport queueing dynamics under severe flow restrictions. In: Proceedings of the 2002 American Control Conference, vol 2, pp 1314–1319, doi: 10.1109/ACC.2002.1023202

Cassell R, Evers C (1998) Development of airport surface surveillance performance requirements. In: Proceedings of the AIAA/IEEE/SAE 17th Digital Avionics Systems Conference (DASC), pp. F32-1–F32-7, vol 2, doi: 10.1109/DASC.1998.739821

Chandran B, Balakrishnan H (2007) A dynamic programming algorithm for robust runway scheduling. In: Proceedings of the American Control Conference, New York, USA, pp 1161–1166, doi: 10.1109/ACC.2007.4282922

Chatfield C (2003) The analysis of time series: An introduction, 6th edn. CRC Texts in Statistical Science, Chapman & Hall, CRC Press, ISBN: 9780203491683

Chen J (2009) Biological inspired optimisation algorithms for transparent knowledge extraction allied to engineering materials process. PhD thesis, Department of Automatic Control and Systems Engineering, The University of Sheffield, UK

Chen J, Mahfouf M (2006) A population adaptive based immune algorithm for solving multi-objective optimization problems. In: Artificial Immune Systems, Lecture Notes in Computer Science, vol 4163, Springer Berlin, pp 280–293

Chen J, Mahfouf M (2012) Improving transparency in approximate fuzzy modeling using multi-objective immune-inspired optimisation. International Journal of Computational Intelligence Systems 5(2):322–342, doi: 10.1080/18756891.2012.685311

Chen J, Stewart P (2011) Planning aircraft taxiing trajectories via a multi-ojective immune optimisation. In: Proceedings of the 7th International Conference on Natural Computation (ICNC), Shanghai, China, vol 4, pp 2235–2240, doi: 10.1109/ICNC.2011.6022587

Chen J, Ravizza S, Atkin JAD, Stewart P (2011) On the utilisation of fuzzy rule-based systems for taxi time estimations at airports. In: Proceedings of the 11th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS), Saarbrücken, Germany, OpenAccess Series in Informatics (OASIcs), vol 20, pp 134–145, doi: 10.4230/OASIcs.ATMOS.2011.134

Cheng V (2007) Research progress on an automation concept for surface operation with

time-based trajectories. In: Proceedings of the Integrated Communications, Navigation and Surveillance Conference (ICNS), pp 1–13, doi: 10.1109/ICNSURV.2007.384166

Cheng VHL (2003) Airport surface operation collaborative automation concept. In: Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit, Austin, TX, USA

Cheng VHL, Foyle DC (2002) Automation tools for enhancing ground-operation situation awareness and flow efficiency. In: Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit, Monterey, CA, USA

Cheng VHL, Sweriduk GD (2009) Trajectory design for aircraft taxi automation to benefit trajectory-based operations. In: Proccedings of the 7th Asian Control Conference, Hong Kong, China, pp. 99–104

Clare G, Richards A (2009) Airport ground operations optimizer. In: Proceedings of the EUROCONTROL 8th Innovative Research Workshop & Exhibition (INO)

Clare G, Richards A, Sharma S (2009) Receding horizon, iterative optimization of taxiway routing and runway scheduling. In: Proceedings of the AIAA Guidance, Navigation and Control Conference, Chicago, IL, USA

Clare GL, Richards AG (2011) Optimization of taxiway routing and runway scheduling. IEEE Transactions on Intelligent Transportation Systems 12(4):1000–1013, doi: 10.1109/TITS.2011.2131650

Clewlow R, Simaiakis I, Balakrishnan H (2010) Impact of arrivals on departure taxi operations at airports. In: Proceedings of the AIAA Guidance, Navigation, and Control Conference, doi: 10.2514/6.2010-7698

Climaco JCN, Martins EQV (1982) A bicriterion shortest path algorithm. European Journal of Operational Research 11(4):399–404, doi: 10.1016/0377-2217(82)90205-3

Confessore G, Liotta G, Grieco R (2005) A simulation-based architecture for supporting strategic and tactical decisions in the apron of Rome-Fiumicino Airport. In: Proceedings of the 37th Conference on Winter Simulation (WSC), pp 1596–1605

Cooper WW, Cormier RH, Foster JG, Mills MJ, Mohleji S (2002) Use of the departure enhanced planning and runway/taxiway assignment system (departs) for optimal departure scheduling. In: Proceedings of the 2st Digital Avionics Systems Conference (DASC), Irvine, CA, USA, doi: 10.1109/DASC.2002.1067898

Cordeau JF, Toth P, Vigo D (1998) A survey of optimization models for train routing and scheduling. Transportation Science 32(4):380–404

Cormen TH, Leiserson CE, Rivest RL, Stein C (2001) Introduction to Algorithms, 2nd edn. MIT Press and McGraw-Hill, Cambridge, MA, USA

D'Ariano A, Pacciarelli D, Pranzo M (2007) A branch and bound algorithm for scheduling trains in a railway network. European Journal of Operational Research 183(2):643–657, doi: 10.1016/j.ejor.2006.10.034

Dear RG (1976) The dynamic scheduling of aircraft in the near terminal area. Tech. rep., Flight Transportaiton Laboratory, M.I.T., Cambridge, MA, USA

Dear RG, Sherif YS (1991) An algorithm for computer assisted sequencing and scheduling of terminal area operations. Transportation Research Part A: General 25(2-3):129–139, doi: 10.1016/0191-2607(91)90132-A

Deau R, Gotteland JB, Durand N (2008) Runways sequences and ground traffic optimisation. In: Proceedings of the 3rd International Conference on Research in Air Transportation (ICRAT), Fairfax, VA, USA

Deau R, Gotteland JB, Durand N (2009) Airport surface management and runways scheduling. In: Proceedings of the 8th USA/Europe Air Traffic Management Research and Development Seminar, Napa, CA, USA

Deb K, Anand A, Joshi D (2002) A computationally efficient evolutionary algorithm for real-parameter optimization. Evolutionary computation 10(4):371–395, doi: 10.1162/106365602760972767

Demšar J (2006) Statistical comparisons of classifiers over multiple data sets. Journal of Machine Learning Research 7:1–30

Deonandan I, Balakrishnan H (2009) Evaluation of strategies for reducing taxi-out emissions at airports. In: Proceedings of the 10th AIAA Aviation Technology, Integration, and Operations Conference (ATIO), doi: 10.2514/6.2010-9370

Desrochers M, Soumis F (1988) A generalized permanent labelling algorithm for the shortest path problem with time windows. INFOR 26(3):191–212

Diepen G, van den Akker J, Hoogeveen J (2009) Integrated gate and bus assignment at Amsterdam Airport Schiphol. Lecture Notes in Computer Science 5868:338–353, doi: 10.1007/978-3-642-05465-5_14

Diepen G, vanden Akker J, Hoogeveen J, Smeltink J (2012) Finding a robust assignment of flights to gates at Amsterdam Airport Schiphol. Journal of Scheduling 15:703–715

Dijkstra EW (1959) A note on two problems in connexion with graphs. Numerische Mathematik 1:269–271, doi: 10.1007/BF01386390

Ding H, Lim A, Rodrigues B, Zhu Y (2005) The over-constrained airport gate assignment problem. Computers & Operations Research 32(7):1867–1880, doi: 10.1016/j.cor.2003.12.003

Dorndorf U, Drexl A, Nikulin Y, Pesch E (2007) Flight gate scheduling: State-of-the-art and recent developments. Omega 35(3):326–334, doi: 10.1016/j.omega.2005.07.001

Dorndorf U, Jaehn F, Pesch E (2008) Modelling robust flight-gate scheduling as a clique partitioning problem. Transportation Science 42(3):292–301, doi: 10.1287/trsc.1070.0211

Dorndorf U, Jaehn F, Pesch E (2012) Flight gate scheduling with respect to a reference schedule. Annals of Operations Research 194(1):177–187, doi: 10.1007/s10479-010-0809-8

Drexl A, Nikulin Y (2008) Multicriteria airport gate assignment and Pareto simulated annealing. IIE Transactions 40(4):385–397, doi: 10.1080/07408170701416673

Durand N, Alliot JM (1998) Genetic crossover operator for partially separable functions. In: Proceedings of the 3rd annual Genetic Programming Conference, USA

Durbin J, Watson GS (1950) Testing for serial correlation in least squares regression, I. Biometrika 37:409–428

Durbin J, Watson GS (1951) Testing for serial correlation in least squares regression, II. Biometrika 38:159–179

Ernst AT, Krishnamoorthy M, Storer RH (1999) Heuristic and exact algorithms for scheduling aircraft landings. Networks 34(3):229–241

Esser J, Schreckenberg M (1997) Microscopic simulation of urban traffic based on cellular automata. International Journal of Modern Physics C 08(05):1025–1036, doi: 10.1142/S0129183197000904

Eurocontrol (2012) Airport CDM implementation manual. Tech. Rep. version 4, Eurocontrol

European Commission (2011) Flightpath 2050, Europe's vision for aviation. Tech. rep., Report of the High Level Group on Aviation Research, doi: 10.2777/50266

Even S, Itai A, Shamir A (1976) On the complexity of timetable and multicommodity flow problems. SIAM Journal on Computing 5(4):691–703, doi: 10.1137/0205048

FAA (2005a) SAGE: Technical manual. Tech. Rep. FAA-EE-2005-01, FAA, Office of Environment and Energy

FAA (2005b) SAGE: Validation assessment, model assumptions and uncertainties. Tech. Rep. FAA-EE-2005-03, FAA, Office of Environment and Energy

Fox J (2002) Time-series regression and generalized least squares: Appendix to an R and S-PLUS companion to applied regression

Ganesan R, Balakrishna P, Sherry L (2010) Improving quality of prediction in highly dynamic environments using approximate dynamic programming. Quality and Reliability Engineering International 26(7):717–732, doi: 10.1002/qre.1127

García J, Berlanga A, Molina JM, Casar JR (2005) Optimization of airport ground operations integrating genetic and dynamic flow management algorithms. AI Communications 18(2):143–164

Gawrilow E, Köhler E, Möhring R, Stenzel B (2008) Dynamic routing of automated guided vehicles in real-time. In: Mathematics - key technology for the future, Springer, pp 165–178

Gendreau M, Potvin JY (2005) Search Methodologies, Springer, chap Tabu Search, pp 165–186, ISBN: 0387234608

Goldberg DE (1989) Genetic algorithms in search, optimization, and machine learning. The University of Michigan, Addison-Wesley, ISBN: 0-201-15767-5

Gong C (2009) Kinematic airport surface trajectory model development. In: Proceedings of the 9th AIAA Aviation Technology, Integration and Operations Conference

Gotteland JB, Durand N (2003) Genetic algorithms applied to airport ground traffic optimization. In: Proceedings of the Congress on Evolutionary Computation (CEC), Canberra, Australia, vol 1, pp 544–551, doi: 10.1109/CEC.2003.1299623

Gotteland JB, Durand N, Alliot JM, Page E (2001) Aircraft ground traffic optimization. In: Proceedings of the 4th USA/Europe Air Traffic Management Research and Development Seminar, Santa Fe, NM, USA

Gotteland JB, Durand N, Alliot JM (2003) Handling CFMU slots in busy airports. In: Proceedings of the 5th USA/Europe Air Traffic Management Research and Development Seminar, Budapest, Hungary

Gröflin H, Klinkert A (2009) A new neighborhood and tabu search for the blocking job shop. Discrete Applied Mathematics 157(17):3643–3655, doi: 10.1016/j.dam.2009.02.020

Gupta G, Malik W, Jung Y (2010a) Incorporating active runway crossings in airport departure scheduling. In: Proceedings of the AIAA Guidance, Navigation, and Control Conference, Toronto, Canada, doi: 0.2514/6.2010-7695

Gupta P, Subramanian H, Pant RS (2010b) Generation of optimized routes and schedules for surface movement of aircraft on taxiways. In: Proceedings of the 10th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference, Forth Worth, TX, USA

Hall M, Frank E, Holmes G, Pfahringer B, Reutemann P, Witten IH (2009) The WEKA data mining software: an update. ACM SIGKDD Explorations Newsletter 11(1):10–18, doi: 10.1145/1656274.1656278

Hall NG, Sriskandarajah C (1996) A survey of machine scheduling problems with blocking and no-wait in process. Operations Research 44(3):510–525, doi: 10.1287/opre.44.3.510

Han J, Kamber M, Pei J (2011) Data Mining: Concepts and Techniques, 3rd edn. The Morgan Kaufmann Series in Data Management Systems, Elsevier Science, ISBN: 9780123814791

Herrero JG, Berlanga A, Molina JM, Casar JR (2005) Methods for operations planning in airport decision support systems. Applied Intelligence 22(3):183–206, doi: 10.1007/s10791-005-6618-z

Horst R, Tuy H (1990) Global Optimization: Deterministic Approaches. Springer, ISBN: 978-3-662-02598-7

ICAO (2008) International standards and recommended practices, Annex 16, Environmental protection: Aircraft engine emissions. Tech. rep., International Civil Aviation Organization

Idris HR (2001) Observation and analysis of departure operations at Boston Logan International Airport. PhD thesis, Massachusetts Institute of Technology, USA

Idris HR, Clarke JP, Bhuva R, Kang L (2002) Queuing model for taxi-out time estimation. Air Traffic Control Quarterly 10(1):1–22

Jaehn F (2010) Solving the flight gate assignment problem using dynamic programming. Zeitschrift für Betriebswirtschaft 80:1027–1039, doi: 10.1007/s11573-010-0396-9

Jang JS (1993) ANFIS: adaptive-network-based fuzzy inference system. IEEE Transactions on Systems, Man and Cybernetics 23(3):665–685, doi: 10.1109/21.256541

Jordan R, Ishutkina M, Reynolds T (2010) A statistical learning approach to the modeling of aircraft taxi time. In: Proceedings of the IEEE/AIAA 29th Digital Avionics Systems Conference (DASC), pp. 1.B.1-1–1.B.1-10, doi: 10.1109/DASC.2010.5655532

Jung Y, Hoang T, Montoya J, Gupta G, Malik W, Tobias L (2011) Performance evaluation of a surface traffic management tool for Dallas/Fort Worth International Airport. In: Proceedings of the 9th USA/Europe Air Traffic Management Research and Development Seminar, Berlin, Germany

Keith G, Richards A (2008) Optimization of taxiway routing and runway scheduling. In: Proceedings of the AIAA Guidance, Navigation and Control Conference, Honolulu, HI, USA, doi: 0.2514/6.2008-6827

Khadilkar H, Balakrishnan H (2011) Estimation of aircraft taxi-out fuel burn using flight data recorder archives. In: Proceedings of the AIAA Guidance, Navigation, and Control Conference, doi: 10.2514/6.2011-6383

Khadilkar H, Balakrishnan H (2012) Estimation of aircraft taxi fuel burn using flight data recorder archives. Transportation Research Part D: Transport and Environment 17(7):532–537, doi: 10.1016/j.trd.2012.06.005

Kim B, Rachami J (2008) Aircraft emissions modeling under low power conditions. In: Proceedings of the A&WMA's 101st Annual Conference and Exhibition, Portland, OR, USA

Kim S, Feron E (2011) Robust gate assignment. In: Proceedings of the AIAA Guidance, Navigation, and Control Conference, pp 2991–3002

Kim SH, Feron E, Clarke JP (2009) Assigning gates by resolving physical conflicts. In: Proceedings of the AIAA Guidance, Navigation, and Control Conference, Chicago, IL, USA, doi: 10.2514/6.2009-5648

Kim SH, Feron E, Clarke JP (2010) Airport gate assignment that minimizes passenger flow in terminals and aircraft congestion on ramps. In: Proceedings of the AIAA Guidance, Navigation, and Control Conference, Toronto, Canada, vol 2, pp 1226–1238

Kistler MS, Gupta G (2009) Relationship between airport efficiency and surface traffic. In: Proceedings of the 9th AIAA Aviation Technology, Integration, and Operations Conference, Hilton Head, SC, USA

Kosko B (1994) Fuzzy systems as universal approximators. IEEE Transactions on Computers 43(11):1329–1333, doi: 10.1109/12.324566

Lawler EL (1972) A procedure for computing the k best solutions to discrete optimization problems and its application to the shortest path problem. Management Science 18(7):401–405

Leese RA, Craig A, Ketzscer R, Noble SD, Parrott K, Preater J, Wilson RE, Wood DA (2001) The sequencing of aircraft departures. Tech. rep., Study report from the 40th European Study Group with Industry

Lesire C (2009) Automatic planning of ground traffic. In: Proceedings of the AIAA Aerospace Sciences Meeting (ASM), Orlando, USA

Lesire C (2010) Iterative planning of airport ground movements. In: Proceedings of the 4th International Conference on Research in Air Transportation (ICRAT), Budapest, Hungary, pp 147–154

Li L, Clarke JP, Chien HHC, Melconian T (2009) A probabilistic decision-making model for runway configuration planning under stochastic wind conditions. In: Proceedings of the Digital Avionics Systems Conference (DASC), pp. 3.A.5-1–3.A.5-10, doi: 10.1109/DASC.2009.5347528

Maharjan B, Matis TI (2011) An optimization model for gate reassignment in response to flight delays. Journal of Air Transport Management 17(4):256–261, doi: 10.1016/j.jairtraman.2011.02.012

Mamdani E (1974) Application of fuzzy algorithms for control of simple dynamic plant. Proceedings of the Institution of Electrical Engineers 121(12):1585–1588, doi: 10.1049/piee.1974.0328

Marín Á (2006) Airport management: Taxi planning. Annals of Operations Research 143(1):191–202, doi: 10.1007/s10479-006-7381-2

Marín Á, Codina E (2008) Network design: Taxi planning. Annals of Operations Research 157(1):135–151, doi: 10.1007/s10479-007-0194-0

Mascis A, Pacciarelli D (2002) Job-shop scheduling with blocking and no-wait constraints. European Journal of Operational Research 143(3):498–517, doi: 10.1016/S0377-2217(01)00338-1

Montgomery DC, Peck EA, Vining GG (2001) Introduction to Linear Regression Analysis, 3rd edn. John Wiley & Sons, Inc., ISBN: 978-0470125069

Mori R (2010) Modeling of aircraft surface traffic flow at congested airport using cellular automata. In: Proceedings of the 4th International Conference on Research in Air Transportation (ICRAT), Budapest, Hungary, pp 139–145

Morris KM (2005) Results from a number of surveys of power settings used during taxi operations. Tech. rep., ENV/KMM/1126/14.8, British Airways

Nadeau C, Bengio Y (2003) Inference for the generalization error. Machine Learning 52:239–281, doi: 10.1023/A:1024068626366

Nagel K, Schreckenberg M (1992) A cellular automaton model for freeway traffic. J Phys I France 2(12):2221–2229, doi: 10.1051/jp1:1992277

Nikoleris T, Gupta G, Kistler M (2011) Detailed estimation of fuel consumption and emissions during aircraft taxi operations at Dallas/Fort Worth International Airport. Transportation Research Part D: Transport and Environment 16(4):302–308, doi: 10.1016/j.trd.2011.01.007

Nishi T, Ando M, Konishi M (2005) Distributed route planning for multiple mobile robots using an augmented lagrangian decomposition and coordination technique. IEEE Transactions on Robotics 21(6):1191–1200, doi: 10.1109/TRO.2005.853489

Nishi T, Shimatani K, Inuiguchi M (2009) Decomposition of Petri nets and Lagrangian relaxation for solving routing problems for AGVs. International Journal of Production Research 47(14):3957–3977, doi: 10.1080/00207540701846244

Nishi T, Hiranaka Y, Grossmann IE (2011) A bilevel decomposition algorithm for simultaneous production scheduling and conflict-free routing for automated guided vehicles. Computers & Operations Research 38(5):876–888, doi: 10.1016/j.cor.2010.08.012

Ortiz MC, Sarabia LA, Herrero A (2006) Robust regression techniques: A useful alternative for the detection of outlier data in chemical analysis. Talanta 70(3):499–512, doi: 10.1016/j.talanta.2005.12.058

Osborne J (2010) Improving your data transformations: Applying the box-cox transformation. Practical Assessment, Research & Evaluation 15(12):1–9

Pesic B, Durand N, Alliot JM (2001) Aircraft ground traffic optimisation using a genetic algorithm. In: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO), San Francisco, CA, USA, pp 1397–1404

Petroccione L (2007) Delta's operation clockwork, transforming the fundamentals of an airline

Pina P, Pablo JD (2005) Benefits obtained from the estimation and distribution of realistic taxi times. In: Proceedings of the 6th ATM R&D Seminar, Baltimore, USA

Pina P, de Pablo J, Mas M (2005) Linking existing on ground, arrival and departure operations. In: Proceedings of the 6th ATM R&D Seminar, Baltimore, USA

Pinedo ML (2012) Scheduling: Theory, Algorithms, and Systems, 4th edn. Springer, ISBN: 978-1-4614-2361-4

Psaraftis HN (1980) A dynamic programming approach for sequencing groups of identical jobs. Operations Research 28(6):1347–1359, doi: 10.1287/opre.28.6.1347

Quinlan JR (1992) Learning with continuous classes. In: Proceedings of the 5th Australian Joint Conference on Artificial Intelligence, pp 343–348

Rappaport DB, Yu P, Griffin K, Daviau C (2009) Quantitative analysis of uncertainty in airport surface operations. In: Proceedings of the 9th AIAA Aviation Technology, Integration, and Operations Conference, Hilton Head, SC, USA, doi: 10.2514/6.2009-6987

Rathinam S, Montoya J, Jung Y (2008) An optimization model for reducing aircraft taxi times at the Dallas Fort Worth International Airport. In: Proceedings of the 26th International Congress of the Aeronautical Sciences (ICAS)

Ravizza S (2009) Control of automated guided vehicles (AGVs). Master's thesis, ETH Zurich, Switzerland

Ravizza S, Atkin JAD (2011) Exploration of the ordering for a sequential airport ground movement algorithm. Tech. Rep. 1543, University of Nottingham, UK

Ravizza S, Atkin JAD, Maathuis MH, Burke EK (2012a) A statistical approach for improving taxi time estimations at airports. Journal of the Operational Research Society (advance online publication), doi: 10.1057/jors.2012.123

Ravizza S, Chen J, Atkin JAD, Burke EK, Stewart P (2012b) The trade-off between taxi time and fuel consumption in airport ground movement. In: Proceedings of the Conference on Advanced Systems for Public Transport (CASPT12), Santiago, Chile

Richli R (2009) Approaches in simultaneous AGV routing. Master's thesis, ETH Zurich, Switzerland

Roling PC, Visser HG (2008) Optimal airport surface traffic planning using mixed-integer linear programming. International Journal of Aerospace Engineering 2008(1):1–11, doi: 10.1155/2008/732828

Rosenberger JM, Schaefer AJ, Goldsman D, Johnson EL, Kleywegt AJ, Nemhauser GL (2000) Air transportation simulation: SimAir: a stochastic model of airline operations. In: Proceedings of the 32nd conference on winter simulation, San Diego, CA, USA, pp 1118–1122

Rosenberger JM, Schaefer AJ, Goldsman D, Johnson EL, Kleywegt AJ, Nemhauser GL (2002) A stochastic model of airline operations. Transportation Science 36(4):357–377, doi: 10.1287/trsc.36.4.357.551

Rousseeuw P, Leroy A (1987) Robust Regression and Outlier Detection. Wiley series in probability and mathematical statistics: Applied probability and statistics, Wiley, ISBN: 9780471852339

Salehipour A, Modarres M, Naeni LM (2013) An efficient hybrid meta-heuristic for aircraft landing problem. Computers & Operations Research 40(1):207–213, doi: 10.1016/j.cor.2012.06.004

Sancho NGF (1994) Shortest path problems with time windows on nodes and arcs. Journal of Mathematical Analysis and Applications 186(3):643–648, doi: 10.1006/jmaa.1994.1324

Sastry K, Goldberg D, Graham K (2005) Search Methodologies, Springer, chap Genetic Algorithms, pp 97–125, ISBN: 0387234608

Schüpbach K, Zenklusen R (2011) Approximation algorithms for conflict-free vehicle routing. In: Proceedings of the 19th Annual European Symposium on Algorithms (ESA), pp. 640–651, doi: 10.1007/978-3-642-23719-5_54

Schüpbach K, Zenklusen R (2012) An adaptive routing approach for personal rapid transit. Mathematical Methods of Operations Research, doi: 10.1007/s00186-012-0403-8

Sedgewick R, Vitter JS (1986) Shortest paths in Euclidean graphs. Algorithmica 1(1):31–48, doi: 10.1007/BF01840435

SESAR (2006) Milestone deliverable D1 - Air transport framework: The current situation. Tech. Rep. Edition 3, Eurocontrol

SESAR (2009) European air traffic management master plan. Tech. Rep. Edition 1, Eurocontrol

Shapiro SS, Wilk MB (1965) An analysis of variance test for normality (complete samples). Biometrika 52(3-4):591–611, doi: 10.1093/biomet/52.3-4.591

Simaiakis I, Balakrishnan H (2009) Queuing models of airport departure processes for emissions reduction. In: Proceedings of the AIAA Guidance, Navigation, and Control Conference, Chicago, IL, USA, doi: 10.2514/6.2009-5650

Simaiakis I, Balakrishnan H (2010) Impact of congestion on taxi times, fuel burn, and emissions at major airports. Transportation Research Record: Journal of the Transportation Research Board 2184:22–30, doi: 10.3141/2184-03

Simaiakis I, Khadilkar H, Balakrishnan H, Reynolds TG, Hansman RJ, Reilly B, Urlass S (2011) Demonstration of reduced airport congestion through pushback rate control. In: Proceedings of the 9th USA/Europe Air Traffic Management Research and Development Seminar, Berlin, Germany

Smeltink JW, Soomer MJ, de Waal PR, van der Mei RD (2004) An optimisation model for airport taxi scheduling. In: Proceedings of the INFORMS Annual Meeting, Denver, CO, USA

Smola AJ (1998) Learning with kernels. PhD thesis, TU Berlin, Germany

Smola AJ, Schoelkopf B (2004) A tutorial on support vector regression. Statistics and Computing 14:199–222, doi: 10.1023/B:STCO.0000035301.49549.88

Soomer M, Franx G (2008) Scheduling aircraft landings using airlines' preferences. European Journal of Operational Research 190(1):277–291, doi: 10.1016/j.ejor.2007.06.017

Soomer MJ, Koole GM (2008) Fairness in the aircraft landing problem. In: Proceedings of the Anna Valicek Competition 2008

Srivastava A (2011) Improving departure taxi time predictions using ASDE-X surveillance data. In: Proceedings of the IEEE/AIAA 30th Digital Avionics Systems Conference (DASC), pp. 2B5-1–2B5-14

Stenzel B (2008) Online disjoint vehicle routing with application to AGV routing. Phd thesis, TU Berlin, Germany

Stettler MEJ, Eastham S, Barrett SRH (2011) Air quality and public health impacts of UK airports. Part I: Emissions. Atmospheric Environment 45(31):5415–5424, doi: 10.1016/j.atmosenv.2011.07.012

Takagi T, Sugeno M (1985) Fuzzy identification of systems and its applications to modeling and control. IEEE Transactions on Systems, Man and Cybernetics 15(1):116–132

Tanaka Y, Nishi T, Inuiguchi M (2010) Dynamic optimization of simultaneous dispatching and conflict-free routing for automated guided vehicles - Petri net decomposition approach. Journal of Advanced Mechanical Design Systems and Manufacturing 4(3):701–715, doi: 10.1299/jamdsm.4.701

Tang CH, Yan S, Hou YZ (2009) A gate reassignment framework for real time flight delays. 4OR: A Quarterly Journal of Operations Research 8(3):299–318, doi: 10.1007/s10288-009-0112-1

Tavakkoli-Moghaddam R, Yaghoubi-Panah M, Radmehr F (2012) Scheduling the sequence of aircraft landings for a single runway using a fuzzy programming approach. Journal of Air Transport Management 25(0):15–18, doi: 10.1016/j.jairtraman.2012.03.004

Tu Y, Ball MO, Jank WS (2008) Estimating flight departure delay distributions - a statistical approach with long-term trend and short-term pattern. Journal of the American Statistical Association 103(481):112–125, doi: 10.1198/016214507000000257

Venables WN, Ripley BD (2002) Modern Applied Statistics with S, 4th edn. Springer, ISBN: 0-387-95457-0

Vis IF (2006) Survey of research in the design and control of automated guided vehicle systems. European Journal of Operational Research 170(3):677–709, doi: 10.1016/j.ejor.2004.09.020

Wang J (2011) Methodology for analysis, modeling, and simulation of airport gate-waiting delays. PhD thesis, George Mason University, USA

Wang J, Shortle J, Sherry L (2009a) Analysis of New York LaGuardia Airport gate-waiting delays. In: Proceedings of the Transportation Research Board 89th Annual Meeting

Wang J, Shortle J, Wang J, Sherry L (2009b) Analysis of gate-waiting delays at major US airports. In: Proceedings of the 9th AIAA Aviation Technology, Integration, and Operations Conference (ATIO) and Aircraft Noise and Emissions Reduction Symposium (ANERS), doi: 10.2514/6.2009-7085

Wang LX, Mendel JM (1992) Fuzzy basis functions, universal approximation, and orthogonal least-squares learning. IEEE Transactions on Neural Networks 3:807–814, doi: 10.1109/72.159070

Weisberg S (2005) Applied Linear Regression, 3rd edn. Wiley/Interscience, ISBN: 0-471-66379-4

Wey CC, Anderson BE, Hudgins C, Wey C, Li-Jones X, Winstead E, Thornhill LK, Lobo P, Hagen D, Whitefield P, Yelvington PE, Herndon SC, Onasch TB, Miake-Lye RC, Wormhoudt J, Knighton WB, Howard R, Bryant D, Corporan E, Moses C, Holve D, Dodds W (2006) Aircraft Particle Emissions eXperiment (APEX). Tech. rep., NASA

Witten IH, Frank E, Hall MA (2011) Data mining: practical machine learning tools and techniques, 3rd edn. The Morgan Kaufmann Series in Data Management Systems, Elsevier Science, ISBN: 9780123748560

Wu C, Caves R (2002) Research review of air traffic management. Transportation reviews 22(1):115–132

Wu Y, Zhang B, Lu J, Du KL (2011) Fuzzy logic and neuro-fuzzy systems: A systematicintroduction. International Journal of Artificial Intelligence and Expert Systems 2(2):47–80

Yen JY (1971) Finding the k shortest loopless paths in a network. Management Science 17(11):712–716

Yin K, Tian C, Wang BX, Quadrifoglio L (2012) Analysis of taxiway aircraft traffic at George Bush Intercontinental Airport, Houston, Texas. Transportation Research Record: Journal of the Transportation Research Board 2266:85–94, doi: 10.3141/2266-10

Zadeh L (1965) Fuzzy sets. Information and Control 8(3):338–353, doi: 10.1016/S0019-9958(65)90241-X

Zhang Y, Chauhan A, Chen X (2010) Modeling and predicting taxi out times. In: Proceedings of the 4th International Conference on Research in Air Transportation (ICRAT), Budapest, Hungary, pp 31–35

Appendix

A Autoregressive AR(1) and AR(2) Models

Table A.1 shows additional results from Section 4.3.4.

Table A.1: Coefficients for Zurich Airport with and without autoregressive AR(1) and AR(2) models (ϕ was equal to 0.249 in the autoregressive AR(1) model and the ϕ values for the autoregressive AR(2) model were 0.221 and 0.105)

	Without AR(1) model		With AR(1) model		With AR(2) model	
	Coefficient	Std. Dev.	Coefficient	Std. Dev.	Coefficient	Std. Dev.
(Constant)	-2.601	0.250	-2.529	0.237	-2.528	0.238
ARR	0.260	0.018	0.253	0.017	0.255	0.017
Distance	-7e-05	2e-05	-5e-05	2e-05	-6e-05	2e-05
$\log_{10}$(Distance)	1.161	0.091	1.120	0.086	1.120	0.086
$\log_{10}$(Angle)	-0.143	0.039	-0.133	0.037	-0.133	0.037
$Q_{DEP,\#DEP}$	-0.046	0.004	-0.049	0.004	-0.049	0.004
$Q_{DEP,\#ARR}$	-0.032	0.004	-0.034	0.004	-0.034	0.004
$Q_{ARR,\#DEP}$	-0.067	0.006	-0.071	0.006	-0.069	0.006
$Q_{ARR,\#ARR}$	-0.081	0.007	-0.087	0.007	-0.086	0.007
$N_{DEP,\#DEP}$	0.025	0.005	0.028	0.005	0.027	0.005
$N_{DEP,\#ARR}$	0.029	0.007	0.033	0.007	0.034	0.007
$N_{ARR,\#DEP}$	0.019	0.004	0.020	0.004	0.019	0.004
$N_{ARR,\#ARR}$	0.054	0.008	0.062	0.008	0.060	0.008
$O_{Morning}$	-0.075	0.019	-0.079	0.024	-0.080	0.026
$O_{Evening}$	0.049	0.013	0.052	0.016	0.051	0.018

Table A.2: Coefficients for Stockholm-Arlanda Airport with and without autoregressive AR(1) and AR(2) models (ϕ was equal to 0.242 in the autoregressive AR(1) model and the ϕ values for the autoregressive AR(2) model were 0.205 and 0.146)

	Without AR(1) model		With AR(1) model		With AR(2) model	
	Coefficient	Std. Dev.	Coefficient	Std. Dev.	Coefficient	Std. Dev.
(Constant)	-2.349	0.091	-2.420	0.088	-2.419	0.088
ARR	0.211	0.015	0.213	0.014	0.212	0.014
$\log_{10}$(Distance)	0.922	0.029	0.944	0.028	0.943	0.028
Distance2	-5e-05	8e-06	-5e-05	8e-06	-5e-05	8e-06
$Q_{DEP,\#DEP}$	-0.044	0.005	-0.045	0.005	-0.044	0.005
$Q_{DEP,\#ARR}$	-0.034	0.004	-0.033	0.004	-0.033	0.004
$Q_{ARR,\#DEP}$	-0.052	0.006	-0.055	0.006	-0.054	0.005
$Q_{ARR,\#ARR}$	-0.066	0.011	-0.069	0.010	-0.064	0.010
$N_{DEP,\#DEP}$	0.031	0.005	0.032	0.005	0.031	0.005
$N_{DEP,\#ARR}$	0.036	0.006	0.036	0.006	0.035	0.006
$N_{ARR,\#DEP}$	0.029	0.004	0.030	0.004	0.028	0.004
$N_{ARR,\#ARR}$	0.049	0.011	0.050	0.011	0.047	0.011

B Multiple Linear Regression Models for Heathrow

Table B.1: Coefficients for London Heathrow Airport

	Departure on 27R	Departure on 27L	Arrival on 27R	Arrival on 27L
(Constant)	-0.960	-0.966	-0.039	-0.383
$\log_{10}$(Distance)	0.943	0.976	0.771	0.885
$Q_{DEP,\#DEP}$	-0.018	-0.015	-	-
$Q_{DEP,\#ARR}$	-0.008	-0.013	-	-
$Q_{ARR,\#DEP}$	-	-	-0.034	-0.038
$Q_{ARR,\#ARR}$	-	-	-0.020	-0.018
$N_{DEP,\#DEP}$	0.009	0.004	-	-
$N_{DEP,\#ARR}$	0.017	0.013	-	-
$N_{ARR,\#DEP}$	-	-	0.009	0.009
$N_{ARR,\#ARR}$	-	-	0.015	0.012

C Further Insights from Prediction Models

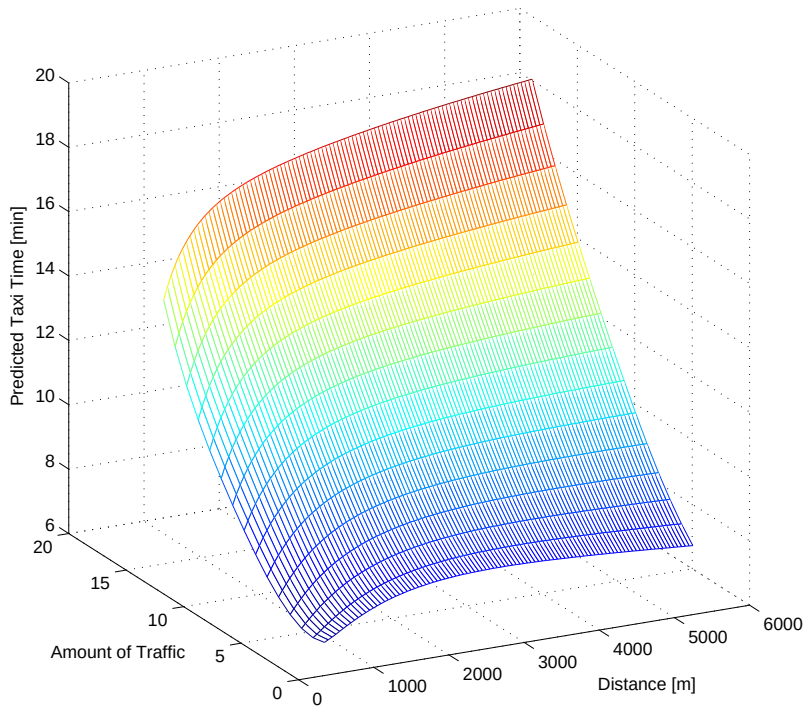

Figure C.1: Analysis of predicted taxi-out times at Stockholm-Arlanda Airport with the TSK FRBS

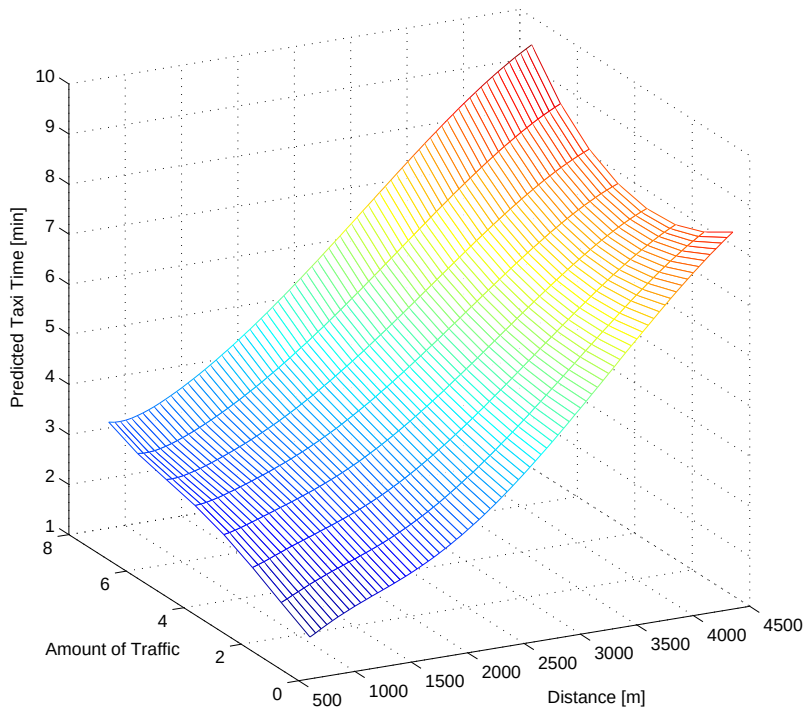

Figure C.2: Analysis of predicted taxi-in times at Zurich Airport with the TSK FRBS (daytime operational mode with fixed average turning angle)

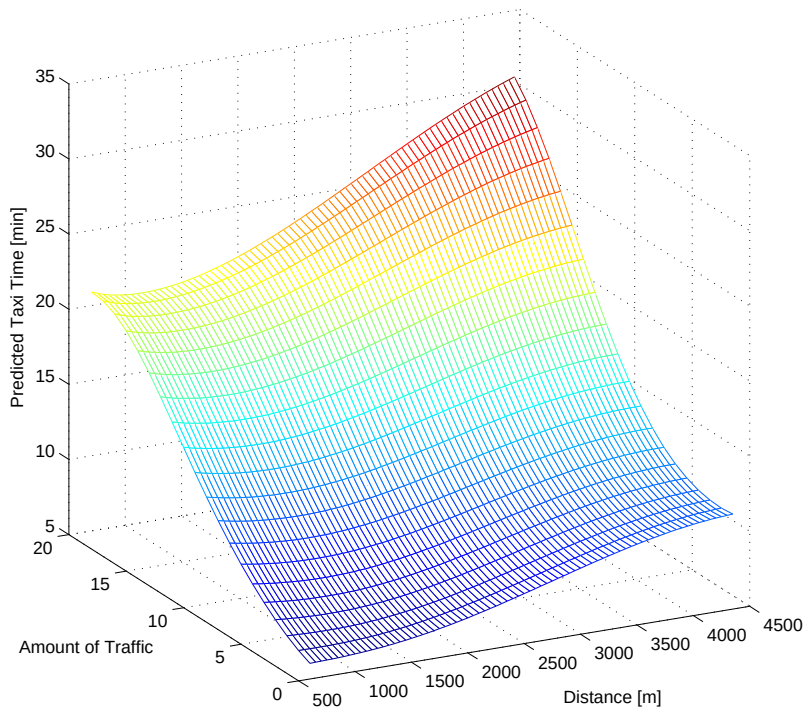

Figure C.3: Analysis of predicted taxi-out times at Zurich Airport with the TSK FRBS (daytime operational mode with fixed average turning angle)

D Taxi Time Prediction for Atlanta Airport

This appendix analyses some of the approaches from this PhD thesis, which were initially focused on European hub airports, extending them to the setting of the world's busiest airport - Hartsfield-Jackson Atlanta International Airport (ATL). In particularly, this appendix shows results for the application of the multiple linear regression approach from Chapter 4 to ATL data and the analysis of Chapter 5 with different regression models.

With around 92 million passengers and around 924000 flights in 2011, ATL is ranked as the busiest airport in the world measured by passengers and by the number of flights. The airport is also the primary hub of Delta Air Lines. We had access to detailed high-fidelity aircraft surveillance data from the 1st of May 2011 and 1948 aircraft movements were reported (993 arrivals and 955 departures). Hartsfield-Jackson Atlanta International Airport can operate with 5 runways as indicated in Figure D.1. All runways are parallel and labelled as 8L/26R, 8R/26L, 9L/27R, 9R/27L and 10/28. All flights were landing or taking-off from west to east in the utilised dataset, where the inner runways (runways 8R and 9L) were used for departures, the outer runways (runways 8L and 9R) for arrivals and runway 10 for both, but with considerably fewer movements.

The airport was modelled, as described in Chapter 3, with 200 gates, 739 nodes and 944 edges and can be seen in Figure D.1.

Multiple linear regression was used to analyse the taxi times of aircraft. Unlike the analysis in Chapter 4, no model was found which satisfied all of the needed statistical assumptions. The best found model is presented in Table D.1 and results in an R^2_{Adj} value of 0.888. This model has as the dependent variable the taxi times of the aircraft (in minutes). Figure D.2 shows the scatterplot for the linear fit of the regression model. These results are less meaningful than the results presented in Chapter 4, since the normality assumption is not valid and also the residual plots are showing abnormalities. Some of them can be seen in Figure D.2, where it is visible that some observed taxi times are considerably larger than the predicted taxi times, but there are no comparable observations with major underestimations.

Table D.1 shows the coefficients of the model as in Tables 4.1 and 4.2. All of the coefficients have

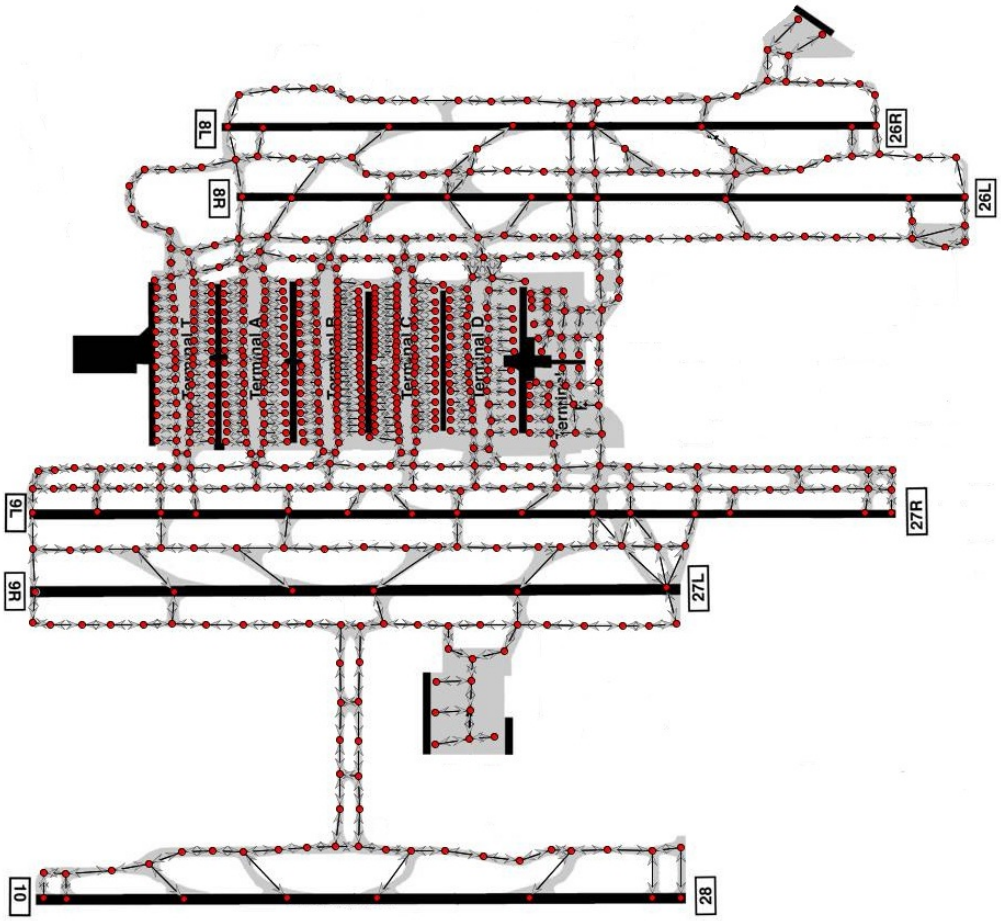

Figure D.1: Layout of Hartsfield-Jackson Atlanta International Airport modelled as a graph with vertices and edges

a different algebraic sign compared to the other tables, since this table reports the coefficients for the model related to taxi times and not the logarithmic transformation of the average taxi speed. Again, all the Q and N values appear to be highly significant and the arrivals generally have shorter taxi times than departures. The distance measure is less important, as was argued in Chapter 4, and is a difference between North American and European airports. However, the last column of Table D.1 should be handled with special care due to some violation of the statistical assumptions.

Table D.1: Coefficients for Hartsfield-Jackson Atlanta International Airport, Sig. indicates if the p-value is < 0.05 (*), < 0.01 (**) or < 0.001 (***)

	Coefficient $\hat{\beta}_i$	Std. Dev.	Standardised Coefficient	Sig.
(Constant)	6.927	0.201		***
$Q_{DEP,\#DEP}$	0.456	0.012	0.888	***
$Q_{ARR,\#DEP}$	0.430	0.015	0.645	***
$Q_{ARR,\#ARR}$	0.326	0.017	0.473	***
$Q_{DEP,\#ARR}$	0.232	0.012	0.401	***
Distance	1E-04	6E-05	0.020	*
ARR	-1.820	0.252	-0.152	***
$N_{ARR,\#ARR}$	-0.129	0.016	-0.166	***
$N_{DEP,\#ARR}$	-0.171	0.014	-0.227	***
$N_{ARR,\#DEP}$	-0.208	0.011	-0.329	***
$N_{DEP,\#DEP}$	-0.206	0.012	-0.360	***

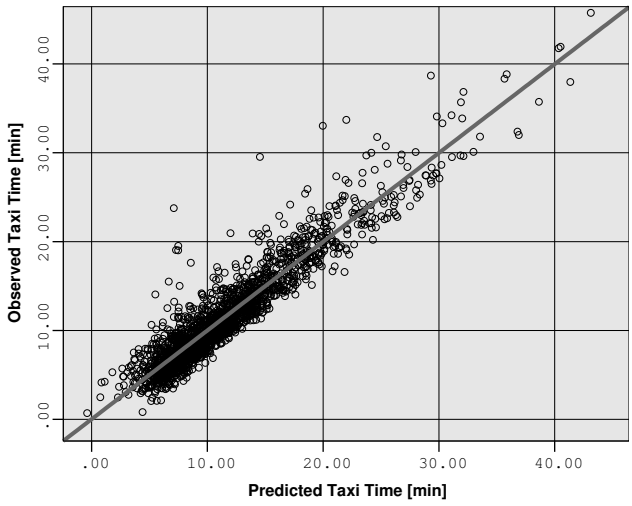

Figure D.2: Scatterplot showing the linear fit of the regression model in Table D.1 for Hartsfield-Jackson Atlanta International Airport

A second analysis was performed to predict taxi times at Hartsfield-Jackson Atlanta Interna-

tional Airport and again highlights the benefits of using different regression approaches. The procedure was identical to the predictions for Zurich Airport and Stockholm-Arlanda Airport. The only differences to the settings from Chapter 5 are the utilised dataset, the explanatory variables (as shown in Table D.1) and that the TSK FRBS used 6 rules instead of 4 (as for Stockholm-Arlanda Airport) or 8 (as for Zurich Airport).

Figure D.3 shows a visual comparison of the 6 different regression approaches as was done in Figure 5.6. Support vector regression performs badly and least median square regression seems to have a bad performance for the underestimated taxi times. It is hard to visually identify bigger differences from the best three approaches: M5 model trees, Mamdani FRBS and TSK FRBS.

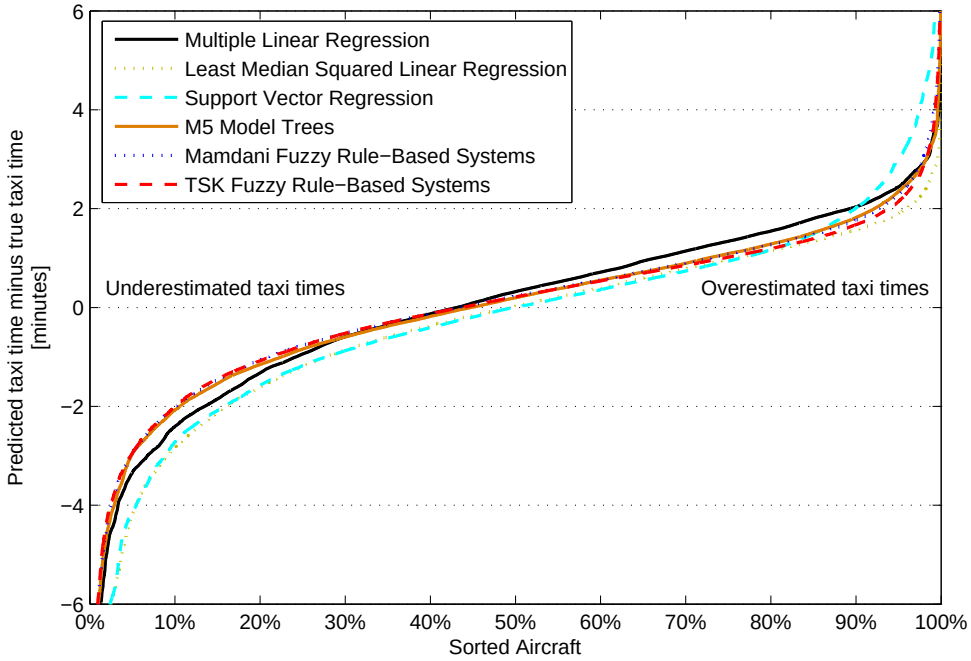

Figure D.3: Taxi time prediction accuracy at Hartsfield-Jackson Atlanta International Airport

The comparison of the performance measures is shown in Table D.2 in the same way as was done in Tables 5.2 and 5.3. The table clearly indicates that TSK FRBS with 6 rules leads to the best performance measures. These results are based on 10-fold cross-validation with 15 repetitions. The corrected resample t-test suggested that the TSK FRBS was significantly

better than least median square regression in all performance measures and was significantly better than support vector regression in all performance measures apart from the ± 10 minutes accuracy. Linear regression was outperformed by TSK FRBS in the ± 1, 2 and 3 minutes accuracy, the mean-absolute error and the relative-absolute error. In addition, TSK FRBS was significantly better than M5 model trees in relation to the ± 1 and 2 minutes accuracy and the mean-absolute error, but was not able to statistically outperform the Mamdani FRBS in any performance measure. Finally, it should be highlighted that the ± 1 minute accuracy can be improved by 26% when using a TSK FRBS instead of the baseline experiment with the multiple linear regression approach.

Table D.2: Comparisons of performance measures for Hartsfield-Jackson Atlanta International Airport

Performance Measure	LinReg	LMS	SMOreg	M5P	Mamdani	TSK
Root mean-squared error	1.98	2.16	2.32	1.81	1.80	**1.78**
Mean-absolute error	1.46	1.44	1.56	1.29	1.26	**1.23**
Root relative-squared error	33.20	36.17	38.85	30.31	30.36	**30.07**
Relative-absolute error	32.69	32.32	34.93	28.88	28.39	**27.58**
Accuracy within ± 1 min	42.29%	48.02%	48.72%	50.08%	51.97%	**53.35%**
Accuracy within ± 2 min	75.78%	79.87%	74.06%	81.58%	82.55%	**83.84%**
Accuracy within ± 3 min	91.56%	90.49%	86.17%	93.39%	93.01%	**93.47%**
Accuracy within ± 5 min	97.99%	96.20%	95.35%	98.26%	98.15%	**98.29%**
Accuracy within ± 10 min	99.64%	99.48%	99.64%	99.74%	99.75%	**99.80%**