Proactive Disruption Management: A Decision Support Tool for Day-Before Cancellations Resulting From Runway Capacity Reductions at Hub Airport

Stein Munting

Delft University of Technology

Proactive Disruption Management: A Decision Support Tool for Day-Before Cancellations Resulting From Runway Capacity Reductions at Hub Airport

by

Stein Munting

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Monday July 17, 2023.

Student number:	4837967	
Project duration:	September 27, 2022 – July 17, 2023	
Thesis committee:	Dr. ir. B.F. Santos,	TU Delft, supervisor & chair
	ir. P.C. Roling,	TU Delft examiner
	Dr. ir. Daan M. Pool,	TU Delft examiner
	ir. B. Denekamp,	External

This thesis is confidential and cannot be made public until July 3, 2025.

Cover Image: Amsterdam Schiphol Airport by Mark Wagtendonk

Acknowledgements

I would like to express my deepest gratitude to the following people who have supported me throughout my research:

First and foremost, I would like to thank my supervisor Dr. Bruno Santos for his expertise, valuable insights and constructive feedback. Your guidance has been crucial in determining the direction and quality of my thesis project.

Secondly, I want to thank my all colleagues at KLM for sharing their knowledge with me. In the first place, Bianca Denekamp, who has been my daily supervisor and was always available to answer my questions and have a critical view on my work. Secondly, Eric Müller, who had first given me the opportunity to do my internship at KLM and afterwards help me in setting up my thesis project. Lastly, I want to thank all the other interns I have met during my time at KLM for the coffee breaks, collaboration and shared experiences.

Finally, a special thank you goes to my family for their continuous support and encouragement throughout my academic journey.

Working on this large project in the field of airline disruption management strengthened my passion for this topic and confirmed my interest in pursuing a career in this field in the future. The project's challenges and complexities confirmed my interest in analytical problem-solving, strategic decision-making, and the dynamic nature of the aviation industry. This report concludes my time as a student. I am proud of the work and final results I delivered.

Stein Munting
Delft, June 2023

Contents

List of Figures

List of Abbreviations

AAR	Airport Arrival Rate
ACNL	Airport Coordination Netherlands
ADR	Airport Departure Rate
ASA	Airport Slot Allocation
ATC	Air Traffic Control
ATFM	Air Traffic Flow Management
CDM	Collaborative Decision Making
DDST	Deterministic Decision Support Tool
DRGH	Dynamic Revisable Ground Holding
FAA	Federal Aviation Administration
FMP	Flow Management Problem
GDP	Ground Delay Program
GHP	Ground Holding Problem
IACM	Integrated Airport Capacity Model
IFR	Instrument Flight Rules
IMC	Instrument Meteorological Conditions
ISA	International Standard Atmosphere
LIFR	Low Instrument Flight Rules
LVNL	Air Traffic Control Netherlands
METAR	Meteorological Aerodrome Reports
MGDP	Metroplex Ground Delay Problem
MILP	Mixed Integer Linear Program
MIP	Mixed Integer Programming
MMC	Marginal Meteorological Conditions
OC	Operations Control
OCC	Operations Control Center
OCT	Operations Control Team
PAAR	Planned Airport Acceptance Rate
PCPF	Pro-active Cancellation Probability Framework
PP	Passenger Promise

RCE	Runway Configuration Estimator
RCM	Runway Capacity Model
SAA	Sample Average Approximation
SAGHP	Single Airport Ground Holding Problem
TMI	Traffic Management Initiative
TMI	Traffic Management Initiatives
VFR	Visual Flight Rules
VMC	Visual Meteorological Conditions

Introduction

Adverse weather events, such as heavy precipitation, severe storms, fog, and low visibility, have been causing significant disruptions, leading to flight delays, cancellations, and subsequent effects on crew schedules and passenger connections. While operational disruption management techniques are commonly used to respond to disruptions on the day of operations, there has been limited exploration of proactive, tactical disruption management approaches. This study aims to address this gap by proposing a pro-active disruption management tool for day-before cancellations. The tool would enable airlines to evaluate meteorological conditions, anticipate runway capacity reductions, and make informed decisions about flight cancellations, rescheduling, and resource allocation. By implementing such proactive measures, airlines can optimize operations, reduce delays and cancellations, communicate efficiently with passengers, and achieve cost savings by minimizing ad-hoc cancellations. The tool has been validated in collaboration with a major European airline with a single hub airport.

The models designed to recover from disruptions in the literature are primarily focused on tactical planning. This means that after a disruption occurs, solutions will be found and decisions will be made. This is a reactive strategy for disruption management. Research on proactive approaches for specific expected disruptions is lacking in the literature. One of the proactive approaches that is already widely used in the airline industry is schedule robustness; this is an approach for disruption management, but only in the long term when specific details of the disruption are unknown.This is also known as strategic planning. Proactive approaches to disruptions between long-term strategic planning and tactical planning are particularly lacking in the literature. Based on this research gap, the research question can be formulated as:

How can an uncertain airport runway capacity reduction, at the hub-airport, be translated into actionable decisions for an airline, a day before operations?

The structure of this thesis report is as follows: The scientific paper that describes the framework is presented in Part I. Part II contains the relevant Literature Study that supports the research.

I

Scientific Paper

Proactive Disruption Management: A Decision Support Tool for Day-Before Cancellations Resulting From Runway Capacity Reductions at Hub Airport

Stein Munting,*

Delft University of Technology, Delft, The Netherlands

Abstract

In the dynamic and complex world of air transportation, airports can serve as vital hubs that connect numerous of destinations and facilitate the smooth flow of passengers and aircraft. However, the efficient operation of hub airports can be significantly disrupted by various factors, one of which is the reduction in runway capacity. Runway capacity reductions can be caused by several factors such as maintenance activities, adverse weather conditions, or airspace congestion, resulting in flight cancellations and delays that impact carefully planned schedules of airlines. By evaluating the expected disruption leading to a runway capacity reduction a day before operations, proactive measures could be taken with effective consideration of aircraft, passenger and crew schedules. To address this challenge, a stochastic, multiple scenario, decision support tool was created with the aim of providing the airline with required proactive cancellations for an expected disruption. A case study is performed on the framework in collaboration with a major European airline providing data. Validation has been performed by benchmarking the results of the framework with actual decisions and results of a deterministic tool including only a single scenario. Finally, a new disruption scenario was used to demonstrate the capabilities of the framework and show the sensitivity of some main input parameters. Benchmarking results of the framework with actual implemented cancellations and results of a deterministic tool show that the framework is capable of providing cheaper cancellation options, and overall, a cheaper cost of solving the disruption while complying with crew reserves and maintaining valuable passenger connections.

1 Introduction

The effective operation of airlines can be highly impacted by adverse weather conditions and industrial actions. Airlines are becoming increasingly concerned as a result of the record-breaking number of unfavorable weather events, such as heavy precipitation, severe storms, fog, and low visibility (Eurocontrol, 2019). In Europe, the operations and profitability of airlines are significantly impacted by these disruptions. Weather-related factors alone account for 20% of all air traffic delays and cost billions of Euros annually, not taking into account reactionary delays. In terms of en route air traffic flow management (ATFM) delays, weather-related delays come in second most frequently (Eurocontrol, 2020). Flight delays put pressure on the carefully planned schedules of aircraft and crew, introducing complexities to the operational efficiency of airlines. Flight delays due to weather conditions, or any other disruption, can affect the entire sequence of subsequent flights leading to a domino effect throughout the airline's network, thereby having an effect on the schedules of aircraft and crew. Additionally, these disruptions at hub airports cause problems for passengers missing connecting flights.

Disruptions can be solved by means of airline disruption management techniques which can be subdivided into tactical disruption management and operational disruption management each with a different focus. Operational disruption management techniques are used on the day of operations when disruptions have occurred, delivering immediate response and decision-making to minimize the impact of the disruption by means of recovery techniques. On the other hand, tactical disruption management is mostly employed before the day of operations, closely monitoring the (weather) situation and make timely adjustments, proactively adjusting to the situation, to ensure smooth operation of the day of operations (Ogunsina and Okolo, 2022). Literature clearly indicates that there has been limited exploration of the topic of tactical disruption management (Hassan et al., 2021). However, proactively preparing for a disruption by means of tactical disruption management techniques could greatly benefit not only passengers and crew but also the entire airline's network. Airlines can optimize crew scheduling, allocate resources effectively, and modify flight operations in advance, resulting in fewer delays, cancellations, and inconveniences, because of the implementation of tactical disruption management techniques. By implementing tactical disrup-

---

[*] Msc Student, Sustainable Air Transport, Faculty of Aerospace Engineering, Delft University of Technology

tion management techniques, airlines can proactively re-optimize crew scheduling, re-allocate resources in advance, modify flight operations, and prepare for the next day. This proactive approach minimizes delays, cancellations, and inefficiencies, as it allows to anticipate and address disruptions with greater flexibility, rather than having to react to disruption with limited schedule adjustments. Additionally, proactive planning enables airlines to communicate with passengers in an efficient manner, informing them of potential delays and providing them with alternate travel options to reduce disruptions to their travel plans.

To address the challenge of preparing for a disruption, this study focuses on the creation of a proactive disruption management tool for day-before cancellations. Airlines can evaluate meteorological conditions, anticipate runway capacity reductions, and make informed decisions about flight cancellations, rescheduling, and resource allocation by taking proactive measures. The proposed framework includes recovery of passengers, crew, and aircraft and takes into account important considerations for airlines such as airport slot management amongst others.

The remainder of the paper is structured as follows. Section 2 presents the problem context. An overview of comparable literature as well as the research gap is presented in Section 3. Section 4 describes the structure of the framework and how the framework has been set up including the optimization model. In Section 5, a case study for a major European Airline with a hub and spoke network is performed, with validation presented in this section as well. Lastly, conclusions and recommendations for future work are provided in Section 6.

2 Problem Context

This section aim to provide the reader with a problem description. As outlined in the introduction, this study focuses on the creation of a proactive disruption management tool for day-before cancellations. This solution arose from the following problem. A runway capacity reduction results in some serious network disturbances for all airlines operating at the airport. Airlines with a hub and spoke network with their hub at this airport are mostly disturbed by these capacity reductions. This study is therefor aimed at airlines with a hub and spoke network. Delayed flights cause passenger misconnections, crew misconnections and aircraft unavailability, to name a few. Airlines want to mitigate the effects of a capacity reduction which can be achieved by analyzing the effects airport wide. While the need to solve this problem is from an airline's perspective, it is important to bear in mind that this reduction affects all flights at the airport. This means that the scope of the problems expands to all flights operating at the airport. Lack of data of flights from other airline complicate the problem and will add uncertainty as only general data of those flights is known. Flight specific data such as cancellation cost, delay cost and amount of passengers is not known. For the airline, it is important to track the effects of the capacity reduction on passenger itineraries, crew schedules and aircraft availability amongst others. This implies an integrated approach for the model, taking into account all these components, as well as the uncertainty of flights from other airlines. Furthermore, the tool should cover the uncertainty of the runway capacity reduction, as the evaluation is done a day before operations and thus the exact reduction is not known yet. The solution would be a decision support tool with day-before cancellations as a solution to mitigate the effects of the capacity reduction minimizing costs for the airline. This approach will reduce ad-hoc cancellations and by using the model a day before operations, time pressure is less present and better decisions can be made.

3 Literature Review

A topic in literature related to this problem is known as the Ground Delay Problem (GDP). Several research initiatives over the past decades have focused on formulating decision support tools and optimized arrival strategies for the GDP. The first attempts to characterize the ground delay problem date back to 1987, when a study on the flow management problem in air traffic control by Odoni (1987) was published. At the period, the aviation sector was rapidly expanding, and airline traffic scaled up rapidly while airport capacity remained mainly constant. Odoni (1987) recognized the costs related to the congestion and created the Flow Management Problem (FMP) to solve congestions by focussing on strategic flow actions. The solution makes a trade-off between ground-holds and airborne delays for flights that are involved in the capacity-demand imbalance. Shortly after this paper was published, a paper (supervised by Odoni) was published which explores the types of models that can be used to solve the problem (Terrab, 1990). To begin, the problem can be solved deterministically. Deterministic models do not account for randomness or uncertainty. For the GDP, this means that the airport capacity is known and fixed. It is also possible for the problem to become static-stochastic. Randomness and uncertainty are considered in stochastic models. The capacities in static-stochastic models are random variables with a probability distribution that represents a forecast for these capacities. Because the model is static, the probability distribution is known before any aircraft departure or arrival. When a dynamic-stochastic model is adopted, the problem gets considerably more difficult. The capacities are also random variables in this model, however, the probability distribution of those capacities changes over time, therefore this probability distribution can vary after any arrival or departure of a flight.

From the simple deterministic models, several researches started to explore the possibility of using stochastic models to solve the problem with uncertainties. Richetta and Odoni (1993) were the first to publish a paper on this topic. A single stage stochastic model was proposed which included the uncertainty of runway capacity. This uncertainty was captured in several scenarios, which is common in stochastic models. Hoffman et al. (1999) also explored the area of stochastic modelling by creating a stochastic model not much different than Richetta and Odoni (1993). The main take-away from the paper is the fact that models solving the GDP should be turned into dynamic models to have more reliable results. According to Hoffman et al. (1999) the GDP is both stochastic and dynamic. Stochasticity follows from uncertain runway capacity and the model is dynamic because of changing information over time. In the meanwhile, Richetta and Odoni (1994) already explored the field of a dynamic model to solve the GDP. The results from the model are promising however the first signs of computational complexity arises. The static model is able to find a solution under 50 minutes, however the dynamic model takes almost 4 hours when taking into account three aircraft classes. An improved dynamic model was proposed by Mukherjee and Hansen (2007). This formulation included a scenario tree that accounted for all possible variations in weather outcomes. This scenario tree, however, has the potential to expand in size, proving the model computationally inefficient.

The Ground Delay Problem is a widely researched topic in literature, however it is linked to a very specific disruption, namely a reduction in arrival capacity which does not fully cover the problem to be solved. Also, GDP is a reactive disruption management technique where often flight cancellations are not included and flights will be ground holded or air holded as long as needed. Zooming out, a more general look into disruption management in the airline industry will be taken below. Because disruptions in the airline industry frequently happen and could result in major challenges in the network of the airline, disruption management is also extensively researched in literature. Disruption management is not only a hot topic in airlines but also in other industries such as railway industry (Ghaemi et al., 2017). All industries that are working with schedules have to deal with disruption management, some more complex than others. Carefully planned schedules can become infeasible due to disruptions. Factors causing disruptions in airlines are too many to list, but very generally they can be divided into *airline resource shortages* and *airport and airspace capacity shortages*. Solving these disruptions is known as disruption management (Ball et al., 2007). Lee et al. (2020) classify disruptions into three categories which also captures the propagation of the disruption in the network.

- **Contingent disruptions:** inefficiencies in airline and passenger operations such as late passengers, late crew, aircraft maintenance and late aircraft, to name a few. These disruptions are airline related.
- **Systematic disruptions:** imbalances in demand-capacity at airport or en-route.
- **Propagated disruptions:** disruptions spreading through the time-space network due to interrelations.

Lee et al. (2020) emphasized the significance of propagated disruptions through the network of an airline. Again as described in the paper of Ball et al. (2007), hiding under the flight schedule, there are three other layers of schedules; aircraft schedules, crew schedules and passenger schedules. Given the fact that these are all related, suggests that a disruption on a single flight leg can have significant 'downstream' effects which can lead to several delays in other flight legs.

The urge to solve the disruptions as fast and efficiently as possible comes from the fact that disruptions are very expensive for airlines. The cost of disruptions can be divided into hard and soft costs (Hassan et al., 2021). Hard costs are costs that are made and should be directly paid to the customer, such as compensation, hotel costs, vouchers and rebooking costs. In addition to this, there is also the concept of soft costs. There are referred to as potential future revenue which are missed due to the disruption. These costs differ per passenger class, flown distance and delay duration. Despite this, it is quite complex to estimate soft and hard costs and they are most of the time estimated in literature by using models of Cook (2015).

Within airline planning there are several stages in which different processes are performed. The time horizon runs from long term to short term. In this context long term could be from ten years to one year before operation of the flight. Short term is from one year until one day before operations. In this planning period, planning starts strategic and becomes tactical closer to the day of operations. During these planning periods, the main objective of the airline also changes. Within strategic planning, the airline defines its strategy and vision and tries to match the forecasted demand with their resources, with the main objective to match the demand with supply. Within the tactical phase, airline's main focus is to maximize the revenue by readjusting to the market. Close to the day of operations, the airline mainly wants to minimize the costs, this period starts about four weeks before operations. On the day of operations, obviously additional costs are minimized (Bouarfa et al., 2014).

Generally, disruption management is done in both the strategic and tactical phase of the airline planning process. In the strategic phase, airlines apply proactive measures to limit the amount of disruptions. Schedules are developed that are more flexible and robust to limit the impact of disruptions. proactively, this can be achieved by either using time flexibility or resource flexibility (Abdelghany and Abdelghany, 2018). As described by Aloulou et al. (2010), time flexibility can be implemented by using buffer time in the flight schedule to absorb (part of) the delay. Faster recovery can be achieved by implementing resource flexibility. This includes for example the use of a reserve aircraft which can substitute a disrupted aircraft. Proactively managing disruptions are techniques that are implemented before the actual disruptions happened which implies that uncertainty and robustness are terms that are closely related to proactive disruption management (Szabo et al., 2015). Given that the disruption's precise form and effects are unknown, uncertain parameters can be modeled into scenarios that can be used as inputs for stochastic models.

Disruption management in the tactical phase is done on the day of operations. Once a disruption occurs, Operations Control (OC) of an airline tries to find a solution with the least impact on the network. Actions to be taken include, delaying, cancelling, swapping aircraft, use a reserve aircraft, high speed flying or delaying maintenance operations. These action are also known as recovery techniques. Recovery, as the name implies, is known as a reactive approach. Analysing disruption management techniques, (Hassan et al., 2021) found that apart from a single article, literature has the main focus on reactive approaches on the day of operations itself.

It can be concluded from the presented literature that reactive approaches to disruption management are widely studied. Related to runway capacity reductions, the ground delay problem is the main recovery technique. Proactive approaches are only implemented in the strategic way without exactly knowing the kind of disruption. Concluding from this, it can be stated there exists a gap in the literature in the area of proactive disruption management in the tactical phase of the airline planning process when the approximate effects on the runway capacity reduction are known. In literature, this model will add to the existing research on stochastic models of disruptions following from airport wide capacity reductions with an extended planning time horizon of at least one day before operations. In the airline industry, this model will help to prepare for an expected disrupted day. By modelling the weather uncertainty of the coming day in several scenarios, a robust plan can be made. Proactively cancelling flights benefits passengers, crew, and network management personnel. Even more significant are the cost savings.

4 Modelling Framework

This paper presents a Proactive Cancellation Probability Framework (PCPF) which is build up by several steps and will be outlined in this section. An overview of these steps is presented in Figure 1. Once a disruption is expected, data of the day of the expected disrupted will be fed for pre-processing which consists of flight schedules, passengers itineraries and crew itineraries amongst others. Also, the expected disruption should be expressed in runway capacity reductions. Weather forecast can be translated into several scenarios where probabilities can be assigned to these scenarios. This step prepares the data into sets and parameters which will be the input for the optimization model. The optimization model prepares the schedule for a runway capacity reduction which is an integrated approach taking into account flights, passengers and crew to name a few. These results will be post-processed to save relevant KPIs. Next, variations of the disruption time will be created by means of a shifting disruption window which is part of the framework. Finally, the framework will output a list of flights which should be cancelled proactively a day before operations by collecting the probability a flight will be cancelled over all solutions.

4.1 Pre-processing

In this section, important pre-processing steps will be discussed. Pre-processing data is important for an optimization model and thus for the framework. There are several advantages however they are all linked to the main advantage which is reduce computational time. By pre-processing the problem size can be reduced by eliminating redundant constraints, variables, and/or inequalities. Also improved scalability can be achieved which is the result of reducing the computational complexity of the problem, leading to faster computation times. The complete pre-processing step is visualized in Figure 2. As can be seen, it consists of two main parts, a machine learning model and processing import data. In this section, first some relevant pre-processing steps will be discussed. Afterwards, the machine learning model for cancellation costs will be discussed. The main output from this step are the sets and parameters that will be used in the optimization model.

4.1.1 Flight Set

The problem will be modelled as a slot assignment problem where flights will be assigned to available runway slots. This means that for a specific time slot a limited amount of aircraft can be assigned such that the amount of aircraft does not exceed the available capacity. The available time slots are determined by the input capacity profile, the model is then free to assign flights to certain slots. In the model, scheduled arrival and departure time of flights are used. In the original schedule, arrival and departure times are rounded to 5 minutes but the in the model, the arrival and departure times will be rounded to the start time of the time slot they fall into. For example when using time slots of 20 minutes, a flight departing at 10:25 will fall into the timeslot 10:20-10:40 and the departure time used in the model is 10:20. Flights also have

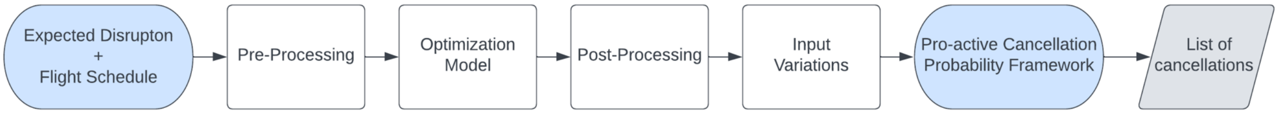

Figure 1: Overview of methodology

a set of available slots they can be assigned to because of a maximum allowed delay. So, in the pre-processing step, for each flight a set of available time slots will be generated with the latest time slot being the last available slot before exceeding the maximum delay. This greatly decreases the size of the flight set and limits the options for the model to assign flights to time slots which also decreases the amount of constraints reducing the computational time. The reference airline can set a limit on the delay which will have an effect on the delay options of a flight. The maximum delay for flights from the reference airline may differ from other flights than the reference airline. As the maximum delay for flights other than the reference airline is not fixed and not known, randomly an additional delay period is added or subtracted from the maximum delay to create some noise, introducing some randomness in the uncertain maximum delay. The importance of this will be in Section 4.4.

Within the flight set, it is also determined if a flight is allowed to be cancelled on the day before operations or on the day of operations. There are several factors that influence this. Flights other than the reference airline are all considered allowed to be cancelled as from the point of view of the airline it is not possible to determine if a flight might not be allowed to be cancelled. For flights of the reference airline there are a few exemptions. First of all, it is not desirable to cancel intercontinental flights due to high yield and high amount of connecting passengers. For this reason, the cancellation cost of these flights is multiplied by 10 to prevent cancellations of these flights. A second reason why a flight may not be cancelled is because of slot reasons.

Airport slots are of very high importance for an airline. An airport slot gives an airline the right to depart or arrive at an airport at a specific day of the week and time. Airport slots were introduced to control busy airports. Every airline could get a set of slots. However, it is also possible for an airline to lose their slot at an airport. This happens when the airline does not use the slot sufficiently, during each season (winter and summer) an airline must use the slot 80% of the time. A slot belongs to a flight and a day of the week. A flight to Paris on Monday has a different slot than the same flight on Tuesday. If a season consists of 20 weeks, an airline is allowed to cancel 4 flights of that specific slot. An airline will hold its slot as long as it is sufficiently used. Sufficiently flying the slots is of high importance due to the price of them, in 2016 Omar Air purchased a slot on Heathrow Airport from KLM for US$75 million (SimpleFlying, 2022). Due to the proven value of slots, it is required to include the slot status in the model, to remove the risk of cancelling a flight with only one flight left to cancel or cancelling a lot of flights early in the season causing flexibility issues at the end of the season. A method should be found which could assess the day-to-day value of a slot when the assessment of the slot status is done at the end of the season. It is assumed that the airline does not have a value for each slot and thus this was implemented in a simplified manner. In this example the winter period is used, which consist of 20 weeks (20 series of flights). As stated above, the airline must fly 80% of the slots in each season, for this winter period this means an airline may cancel 4 flights. For this reason, it is not wise to cancel already 3 flights in the beginning of the season as this reduces flexibility at the end of the season. To determine if a flight is allowed to be cancelled, Equation 1 is used. If below equation is true, a flight is not allowed to be cancelled by the model. This is a very simplified way of implementing slot series in the model, however very effective. In the model, this equation is used for every flight and a set of flights is created which are not allowed to be cancelled due to slots. This equation can be used for airlines which do not have their own rule set on cancelling flights taking into account slots. The most accurate approximation would be if the airline has its own cost model regarding slots, in this way, the associated value of a slot can be compared between flights. As opposed to this simplified equation, more complicated ways of determining the value of a slot could be used. The simplified way represents the slot value as a linear line where it is not allowed to drop below at any point in the season. This could be expanded to a way where penalties are awarded when this value drops below 20% at any point in the season but still assuring that at the end of the season 80% of the slots have been flown.

$$\frac{\text{Flights Allowed To Cancel} - 1}{\text{Flights Left In Series}} < \frac{1}{5} \tag{1}$$

To continue, there is a set of flights which cannot be cancelled on the day of operations as part of the rotation already has been completed. These are rotations of which the outbound flight from the hub airport departed on the day before operations. This means that on the day of operations it is not possible to cancel the complete rotation as the outbound flight already departed. So, the inbound flight of that rotation is not allowed to be cancelled in the second stage. All flights from the reference airline are associated to a rotation.

Lastly, each flight is also identified by a rotation id. A flight rotation consists of a set of flights. The minimum time needed to perform a rotation is the sum of the block times of the flight adding the minimum turnaround time for the aircraft type performing the rotation. Computing the minimum time to perform the rotation in this way means that this number is fixed for each rotation. Delays cannot be recaptured by a shorter flight time or applying cruise speed changes.

4.1.2 Aircraft Set

The maintenance of aircraft is also included in this model. If an aircraft is in maintenance, it is not available to perform a flight. This is important for the constraints that model the subtype availability, as will be explained in Section 4.2. The maintenance schedule for the day of operations is used to create a set for each subtype which holds the times at which an aircraft needs to go to maintenance, and when it returns from maintenance. In this way, an aircraft can be subtracted from the available aircraft at a certain time and added when it returns from maintenance. Maintenance schedules will affect aircraft availability. It may be possible that due to a delay, an aircraft is not able to perform a rotation before it needs to go to maintenance. As this model does not allow to delay or cancel maintenance activities, this could result in a cancellation. Furthermore, for each aircraft a turn around time is defined. This turn around time is defined as the minimum time that is needed at the hub airport between two flights.

[redacted]

4.1.4 Crew Set

Crew data is imported for crew flying on continental flights. For intercontinental flights, no constraints are required because there is no connection between flights and mislocation of crew is not within the scope of the model as this model has the focus on a single day.

For each crew member, the flight schedule is known, which consists of a sequence of flights. In this way, sets can be created with flights which have a crew connection, including connection time. Also for crew connections, connections will be removed from the list with a connection time greater than the minimum connection time for crew plus the maximum delay of a single flight. Furthermore, duty times of crew are taken into account. Maximum duty times of the crew are set by the airline and depend on start time of duty and the amount of flights to be flown. For the last flight in the crew schedule, it is checked if a the maximum delay of that flight causes an exceedance of duty time. If this is the case, delay options are removed from the flight such that with the new maximum delay the crew member will always stay within its duty hours. Also, a set with flights that cause a mislocation of a crew member in case of cancellation is created. A mislocation of a crew member is caused by a crew member not being at the same airport as the departure airport of their next flight which can be caused by a cancellation of one of their previous flights.

4.1.5 Cancellation Cost Set

All flights in the flight set will need to have a cancellation cost and a delay cost. Delay cost need to be computed for every available time slots the flight may be assigned to. The cancellation cost needs to be determined for two moments in time, cancelling the flight before the day of operations and cancelling the flight on the day of operations. It is important that the cost of cancelling a flight a day before operations is distinguished from the cost of the day of operations as this determines the working of the model. If there would be no advantage of cancelling a flight a day before, there would be no need for this model and the airline would always wait to see which scenario is playing out and cancel flights needed for that specific scenario. In the remainder of the paper, the difference in cancellation cost between the day of operations and the day before operations will be noted by the cancellation factor. A factor > 1 indicates that it is more expensive to cancel a flight on the day of operations, a factor < 1 indicates it is cheaper to cancel on the day of operations. The input data will provide the cancellation cost of the flights of the reference airline. These costs are the cost of cancelling a flight on the day before operations as this is the moment the data is retrieved from the airline systems. There are two sets of flights which use a difference approach of retrieving the cancellation cost on the day before operations and on the day of operations, namely a set flights from the reference airline and a set of other flights operating at the airport. First addressing the flights of the reference airline.

Reference airline

The cancellation cost is defined as the sum of actual passenger related costs and the future value losses. The actual cost of cancelling the flight on the day of operations is not known, as this model is run a day before operations, however it can be estimated. This estimation is split in two parts, cost related to passengers and operational costs. Operational savings are fixed and known for a combination of aircraft type and

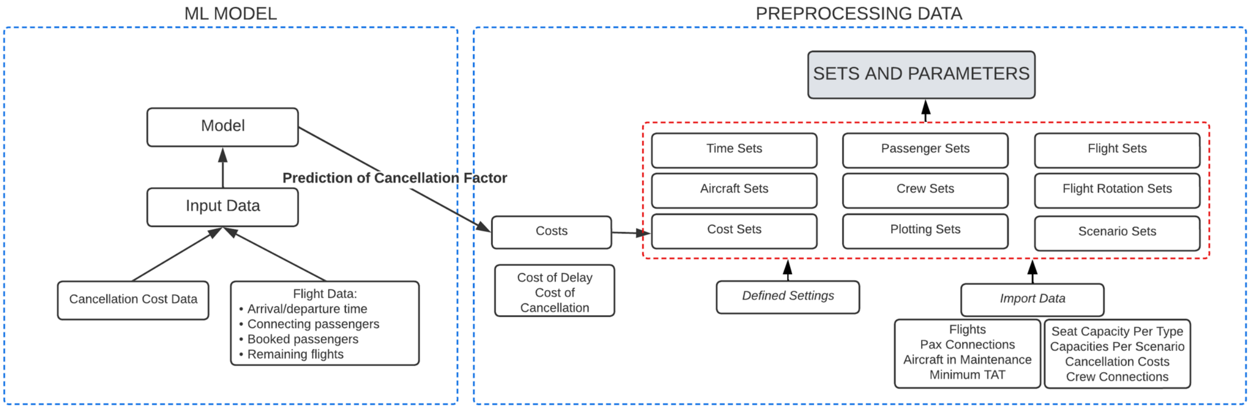

Figure 2: Overview of pre-processing step

destination. For a given aircraft type and destination, the cost savings of cancelling a day before operations is expressing in Euros for the complete rotation. The passenger related cost is not known and not fixed but it can it be estimated based on historical data by estimating the cancellation factor. This can be done by a machine learning model, as depicted in Figure 2. This model will be explained in Section 4.1.7. In the final cancellation cost, operational benefits are subtracted after the estimation of the machine learning model as this model only contains passenger related costs.

[redacted]

4.1.6 Delay Cost Set

Delay costs are determined for every flight's available time slots. The way delay costs are modelled is important for the working of the model. As stated before, this model is airline centric however it is modelling the full airport runway operations. This results in a set of flights from the reference airline for which delay costs are known, and a set of flights for which this is not known. Delay costs play an important role in the decision to delay a flight or not. Cheaper flights will be more likely to be delayed. For the model, it is important that flights of the reference airline and other flights have comparable delays, as this is also the case in real life as delays due to congestion are given by air traffic control. There are multiple ways of achieving a fair distribution of delays which can either be linked to the pre-processing step, or achieving this with constraints in the optimization model. Achieving a fair distribution of delays can be realized by assigning similar delay costs based on the amount of passengers to all flights. In this way, the model will not always benefit a flight from the reference airline over any other flight simply because the costs are lower. Resulting in a fair distribution of delays. Airline delay costs are build up from future value losses and passenger misconnection cost. For the reference airline, these complete costs are known. For other flights this is genuinely hard to estimate. Although, what can be estimated is the future value loss as this depends only on the amount of passengers on board which can be estimated. For the reference airline, the passenger misconnection costs will be included in another set.

A passenger whose flight is delayed is expected to spend less on the airline in the future. This is captured in future value loss. Future value loss curves have been estimated by Cook (2015). For flights of the reference airline, future value loss is calculated for local passengers ie. passengers who are not connecting at the hub airport. Connecting passengers are excluded because the hub airport is not the final destination of the passenger. If the passenger miss their connection because of a delay, the additional costs are accounted for in the passenger misconnection costs. For flights of other airlines, future value loss is calculated for all passengers onboard as the amount of connecting passengers is not known. The total passengers onboard is assumed to be the maximum capacity of the aircraft multiplied by an average load factor of 0.8.

The future value curve as implemented is shown in Figure 3. The curve is non-linear, however it is implemented as a piece-wise linear function to assure linearity of the model. When an airline defined its own future value curves based on more recent estimations, these could be implemented and used as well.

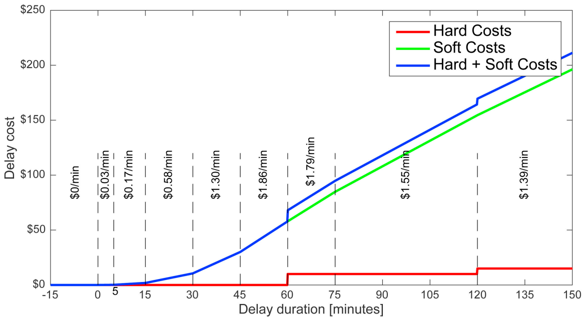

Figure 3: Soft costs (future value costs) as determined by Cook (2015)

4.1.7 ML Model

As explained in Section 4.1.5, the goal of this machine learning model is to predict the cancellation cost of cancelling a flight on the day of operations based on the cost of cancelling a day before operations and some other features. This difference in cost yields the cancellation cost factor. The available data and features available are labeled and therefor supervised learning is preferred. A Random Forest Regression (Coulston et al., 2016) was opted for this machine learning model due to its high accuracy, scalability and ease of use. Two separate models were set up, one for arriving and one for departing flights. The accuracy of the model determines to what extend the model is able to predict the required value. In this case the accuracy is a function of the absolute error. For this model, the accuracy was opted to be sufficient and more accurate than assigning a fixed cancellation cost factor to all flights for both the arriving flights dataset as the departing flights set. An analysis on the accuracy of the model is provided in Section 5.

The set of features which resulted in the highest accuracy have been chosen. The dataset is divided into a set of arriving flights, and a set of departing flights, which is done for the features as well. This resulted in the following set of features for arriving flights: *Arrival Time*, *Outbound Connecting Pax %*, *Aircraft Type*, *Weekday of Arrival*, *Departure Airport*, *Remaining Flights*. For departing flights the features are: *Departure Time*, *Inbound Connecting Pax %*, *Aircraft Type*, *Weekday of Departure*, *Arrival Airport*, *Remaining Flights*. The feature dataset is said to be dense, which means there are almost no missing values or zeros in the features data. The dataset consists of both numerical (eg. *Connecting Pax %*) and categorical features (eg. *Aircraft Type*). Outliers have been removed based on the assumption that the cancellation cost factors are normally distributed. Outliers are identified by values which fall outside the range of the mean ± two standard deviations ($\mu \pm 2\sigma$).

4.2 Optimization Model

Fed by the sets and parameters from the pre-processing step, the optimization model can be build. The optimization model is a Mixed Integer Linear Model (MILP) and the complete formulation described in this section. The results of the optimization model can be used for further analysis of the solution by post-processing.

A two-stage stochastic mixed integer linear programming (MILP) model was formulated for the integrated flight rescheduling. A number of different plausible scenarios were considered, and the flight rescheduling and slot assignment problem, which aimed to minimize cost based on the probabilities of these scenarios, was solved. This naturally led to a two-stage stochastic programming approach with first-stage decision variables being independent of the scenarios, and the second-stage decision variables being scenario-dependent. The first-stage variables represented decisions to cancel flights a day before operations, while the second-stage variables included cancellations on the day of operations and assigned delays.
The two-stage model is solved by means of a single objective function with both first and second stage decision variables. The link between the first and second stage is made in the constraints. To make this notation general following definitions will be used, *airline* is the name of the airline and *hub airport* is the hub airport of that airline where the expected disruption will occur.

4.2.1 Sets and Parameters

The sets and parameters, created by the pre-processing step, are listed in Table 1 and 2 respectively. Additionally, decision variables for the optimization model are listed in Table 3.

Table 1: Sets

Sets	
	Aircraft
TYPES	set of all aircraft types flown by *airline*
GROUPS	set of all flight groups from *airline*
CARRIERS	set of all carriers operating at airport
TYPES^{carr}	set of aircraft types operating at airport for carrier *carr* $\in$ CARRIERS
	Flights and Rotations
F	set of all flights departing and arriving on *hub airport*
F^{arr}	$\subset F$ set of all arriving flights
F^{dep}	$\subset F$ set of all departing flights
$F^{airline}$	$\subset F$ set of all flights of *airline*
F^{canx_1}	$\subset F$ set of flights which are not allowed to be cancelled in the first stage due various reasons
F^{canx_2}	$\subset F$ set of flights which are not allowed to be cancelled in the second stage due various reasons
R	set of rotations (outbound/inbound connected flights)
R^{type}	$\subset R$ set of rotations of *type* $\in$ TYPES
$R^{scenario}$	$\subset R$ set of rotations which are not allowed to be cancelled on day of operations
PAX_CONN	set of pair of flights $(i,j) \in F$ that has a passenger connection
	Crew
CREW	set of all crew members on flight $f \in F^{airline}$
F^{cr}	sequence of flights of crew member $cr \in$ CREW
CREW_CONN	set of pair of flights $(i,j) \in F$ that has a passenger connection
RANKS	set of available ranks of all crew members
	Timeslots and Capacities
T	set of runway time slots
T_f	set of available time slots for flight $f \in F$
$K_{t,[inb,outb]}$	set of runway slots discretized by $t \in T$. For every t, there is a corresponding inbound and outbound capacity $[inb, outb]$
S	set of scenarios which consists of multiple variations of K

Table 2: Parameters

Parameters	
	Aircraft
MA_{type}	number of aircraft of *type* $\in$ TYPES
TAT_{type}	minimum turnaround time for *type* $\in$ TYPES
$AIM_{type,t}$	amount of aircraft in maintenance for *type* $\in$ TYPES at $t \in T$
	Flights and Rotations
$C_{D_{i,t}}$	cost of assigning flight $i \in F$ to time slot $t \in T_i$
C_{C_i}	cost of cancelling flight $i \in F$ a day before operations
C_{C2_i}	cost of cancelling flight $i \in F$ on the day of operations
MRT_r	minimum time needed to complete rotation $r \in R$
ANA_{type}	amount of aircraft for *type* $\in$ TYPES not available to fly at airport
$CT_{i,j}$	scheduled connection time between flight $(i,j) \in F$
[redacted]	
	Crew
$MCTC$	minimum connection time for crew
MCM_{gr}	maximum amount of crew misconnections per flight group $gr \in$ GROUPS
$CF_{i,rank}$	amount of crew of rank *rank* $\in$ RANKS on flight $i \in F$
$CML_{i,rank}$	amount of crew members of *rank* $\in$ RANKS that will be mislocation when flight $i \in F^{airline}$ is cancelled
$CRA_{gr,rank}$	amount of crew reserves available for $gr \in$ GROUPS and *rank* $\in$ RANKS
	Passengers
$MCTP$	minimum connection time for passengers
$C_{CMP_{i,j}}$	cost of a passenger misconnection between flight $(i,j) \in$ PAX_CONN
	Timeslots and Capacities
$cap_{t,[inb/outb],s}$	runway capacity for $t \in T$ configuration $[inb/outb]$ and $s \in S$
p_s	probability of occurrence for scenario $s \in S$
[redacted]	
T_H	discretization per hour
M	a big number

Table 3: Decision Variables

Decision Variables		
Z_i	$\in \{0,1\}$	whether flight $i \in F$ is cancelled a day before operations
$Z_{2_{i,s}}$	$\in \{0,1\}$	whether flight $i \in F$ is cancelled in scenario $s \in S$
$X_{2_{i,t,s}}$	$\in \{0,1\}$	whether flight $i \in F$ is assigned to time slot $t \in T_i$ in scenario $s \in S$
Derived Decision Variables		
$TA_{type,t,s}$	≥ 0	amount of aircraft of *type* $\in$ TYPES available at start of time period $t \in T$ in scenario $s \in S$
$PCM_{i,j,s}$	$\in \{0,1\}$	whether there is a miss connection of passengers between flight $(i,j) \in$ PAX_CONN in scenario $s \in S$
$CCM_{i,j,s}$	$\in \{0,1\}$	whether there is a miss connection of crew between flight $(i,j) \in$ CREW_CONN in scenario $s \in S$
CL_s	$\in \{0,1\}$	whether the amount of cancellations exceeds a set limit in scenario $s \in S$

4.2.2 Mathematical Model

This section outlines the mathematical model starting with the objective function which is followed by the constraints. Afterwards, each constraint will be discussed.

OBJECTIVE FUNCTION

$$\text{Minimize:} \quad \sum_{i \in F} Z_i \cdot C_{C_i} + \tag{2a}$$

$$\sum_{s \in S} p_s \cdot (\sum_{i \in F} Z_{2_{i,s}} \cdot C_{C2_i} + \sum_{i \in F} \sum_{t \in T_i} X_{2_{i,t,s}} \cdot C_{D_{i,t}}$$

$$+ \sum_{(i,j) \in \text{PAX\_CONN}} PCM_{i,j,s} \cdot C_{CMP_{i,j}})$$

subject to:

SLOT ASSIGNMENT

$$\sum_{t \in T_i} X_{2_{i,t,s}} + Z_{2_{i,s}} + Z_i = 1 \qquad \forall i \in F \quad \forall s \in S \tag{2b}$$

RUNWAY CAPACITY

$$\sum_{i \in F^{[inb/outb]}} \sum_{t \in T_i : t = time} X_{2_{i,t,s}} \leq cap_{t,[inb/outb],s} \tag{2c}$$

$$\forall time \in T \qquad \forall s \in S$$

ROTATION CANCELLATION

$$Z_i = Z_j \qquad \forall r = (i,j) \in R \tag{2d}$$

$$Z_{2_{i,s}} = Z_{2_{j,s}} \qquad \forall r = (i,j) \in R, \quad \forall s \in S \tag{2e}$$

$$Z_{2_{i,s}} = 0 \qquad \forall r = (i,j) \in R^{scenario}, \quad \forall s \in S \tag{2f}$$

MINIMUM TURNAROUND

$$\sum_{t \in T_j} (t \cdot X_{2j,t,s}) - \sum_{t \in T_i} (t \cdot X_{2i,t,s}) \tag{2g}$$

$$\geq MRT_{i,j} - (Z_i + Z_j + Z_{i,s} + Z_{j,s}) \cdot MRT_{i,j}$$

$$\forall (i,j) \in R \quad \forall s \in S$$

SUBTYPE AVAILABILITY PART I

$$TA_{type,time[0],s} = MA_{type} - ANA_{type} - AIM_{type,time[0]} \tag{2h}$$

$$+ \sum_{(i,j) \in R_{type}} Z_i \forall type \in \text{TYPES} \qquad \forall s \in S$$

SUBTYPE AVAILABILITY PART II

$$TA_{type,time,s} = TA_{type,time-1,s} - \sum_{t \in T_i : t = time, i \in F^{dep}_{type}} X_{2i,t,s} \tag{2i}$$

$$+ \sum_{t \in T_i : t + TAT_{type} = time, i \in F^{arr}_{type}} X_{2i,t,s} - AIM_{type,time}$$

$$\forall time \in T[1:] \quad \forall type \in \text{TYPES} \quad \forall s \in S$$

NOT ALLOWED TO CANCEL

$$Z_i = 0 \quad \forall i \in F^{canx_1} \tag{2j}$$

$$Z_{2_{i,s}} = 0 \quad \forall i \in F^{canx_2} \quad \forall s \in S \tag{2k}$$

PASSENGER CONNECTIONS

$$\sum_{t \in T_j} (t \cdot X_{2j,t,s}) - \sum_{t \in T_i} (t \cdot X_{2i,t,s}) \geq \tag{2l}$$

$$MCTP - M \cdot (PCM_{i,j,s} + (Z_i + Z_j + Z_{i,s} + Z_{j,s}))$$

$$\forall (i,j) \in \text{PAX\_CONN} \quad \forall s \in S$$

CREW CONNECTIONS

$$\sum_{t \in T_j} (t \cdot X_{2j,t,s}) - \sum_{t \in T_i} (t \cdot X_{2i,t,s}) \geq \tag{2m}$$

$$MCTC - M \cdot (CCM_{i,j,s} + (Z_i + Z_j + Z_{i,s} + Z_{j,s}))$$

$$\forall (i,j) \in \text{CREW\_CONN} \quad \forall s \in S$$

CREW MISCONNECTIONS

$$\sum_{(i,j) \in \text{CREW\_CONN}^{gr}} CCM_{i,j,s} \leq MCM_{gr} \tag{2n}$$

$$\forall gr \in GROUPS \quad \forall s \in S$$

CREW MISLOCATIONS

$$\sum_{i \in F^{gr}} (Z_i + Z_{2_{i,s}}) \cdot CML_{i,rank} \leq CRA_{gr,rank} \tag{2o}$$

$$\forall gr \in \text{GROUPS} \quad \forall rank \in \text{RANKS} \quad \forall s \in S$$

[redacted]

The objective function is a minimization function outlined in Equation 2a. This function minimizes the cost over all scenarios and consists of two parts. The first part are the cost of flights that are cancelled the day before operations and thus for all scenarios. The second part is the sum of cost for a scenario multiplied by the probability of that scenario occurring. For each scenario, costs are the sum of the following components; cancellation cost, delay cost and passenger misconnection cost.

Equation 2b describes the slot assignment constraint. Every arriving and departing flight should either be assigned to a slot or the flight should be cancelled. In this constraint the decision variable of stage one (Z_i) is added which means that if it is decided to cancel flight i in stage one, this flight is cancelled in all scenarios. This constraint is the link between stage one and two. The runway capacity constraint is described by Equation 2c. The amount of flights arriving or departing in a time slot should be less or equal to the capacity of the corresponding time slot. Equation 2d - 2f describe constraints related to cancellation of flight rotations. If the outbound of a rotation is cancelled, the inbound will be cancelled as well. Equation 2f states that flights that already departed cannot be cancelled on the day of operations. For a rotation, the time taken to complete the rotation should be larger than the minimum time needed as scheduled. The minimum time needed for the rotation is the sum of the block time of the outbound flight, the block time of the inbound flight and the minimum turnaround time of the aircraft type at an outstation. If one of the two flights is cancelled, this constraint will not be active anymore. This is captured in Equation 2g.

Equation 2h and 2i deal with the subtype availability. Rotations are not assigned to aircraft registrations, instead throughout the day it is checked if there are enough aircraft available for the amount of flights for that subtype. For each type a maximum number of available aircraft is defined. At each time period, the total amount of flights assigned to a subtype should be less than the maximum available aircraft. Modelling this constraint on subtype level instead of registration level allows for swapping aircraft to another registration of the same subtype however, it does not allow swaps between different subtypes. This constraint is modelled as an inventory constraint. This type of constraint is often used in supply chain management and production planning models. Beginning with a starting value of available aircraft, an aircraft will be subtracted from the available set if departed. An aircraft will be added at the moment it is ready for departure again which is after arrival and minimum turnaround time at the hub airport. First, determining the starting value, this is the number of aircraft which are available at the first time period at 00:00. The amount of aircraft of each type available at 00:00 are the maximum amount of aircraft available of a subtype minus the amount of aircraft of that type which are scheduled to fly or scheduled on a night stop, also subtracting aircraft in maintenance. In case a rotation, either a nightstop, or rotation in flight during the night, is cancelled in stage one by Z_i, an extra available aircraft is added. This is expressed in Equation 2h. This constraint serves as a starting point, now this can be used to loop through all time periods to make sure that the available types at all time periods will be greater or equal to 0, this is described in Equation 2i.

For certain flights, cancellation is undesirable due to various reasons, one of which could be slot constraints, as previously discussed. In F_{canx_2} nightstop flights are included, as these flights cannot be cancelled on the day of operations, as part of the flight (outbound) has already been operated. For these flights, the cancellation decision variable will be set to zero in Equation 2j and 2k. Passenger and crew connection constraints are captured in Equation 2l and 2m. The constraints are very similar and a misconnection is triggered when the time between two flights with a connection is less than the minimum connection time for either crew or passengers. For each flight group, the amount of misconnection of crew is limited by Equation 2n, which depends on the crew reserves available. Apart from misconnections, crew can also be mislocated which is captured in Equation 2o. A mislocated crew member is a crew member which is not at the airport of which the next flight in the flight schedule will depart. For every flight $i \in F^{airline}$ it has been determined if by cancelling this flight a crew member of $rank \in$ RANKS will be mislocated. The amount of mislocations is limited by the amount of crew reserves available. [redacted]

[redacted]

4.3 Post-processing

In the post-processing, the results of the optimization model can be analyzed. This is done based on a set of KPIs which are listed in Table 4. In this way, the solution of different runs of the optimization model can be compared on flight, crew and passenger level as well as model performance. All KPIs are self-explanatory except for the last one. Passenger Promises (PP) is a passenger centric KPI. This KPI is build up from three other KPIs with promises from the airline to the customer. For every passenger it is determined if the passenger arrives within 15 minutes from the scheduled arrival time, the passenger is able to complete its whole journey and the passenger has all their bags at the destination airport. If all three conditions are met, the promises to the passenger are met. PP describes the percentage of passenger to which the promises have been fulfilled. Typically, PP presents a low value in case of disruptions as flights are delayed, maybe even cancelled and due to short connection times, bags might not make it to the next flight. This KPI is presented as a percentage of all passengers.

KPI	Unit
Solution time	[s]
Objective Function Value	[MU]
Cost of Cancellations	[MU]
Passengers on cancelled flights	[#]
Passenger misconnection cost	[MU]
Passenger misconnections	[#]
Crew misconnections	[#]
Crew mislocations	[#]
Average delay reference airline	[minutes]
Average delay non reference airline	[minutes]
Passenger Promise	[%]

Table 4: KPIs saved in post-processing step

4.4 Input Variations

This section explains the implementation of the optimization model in the Proactive Cancellation Probability Framework by means of input variations for the optimization model. The main reason for this implementation follows from the fact that the optimization model itself is quite rigid and fixed. Input capacities are hourly without variation, and scheduled arrival and departure times are used. Apart from the rigidity, there is also a part of the model with random elements. These elements have been discussed before and are all related to flights other than the reference airline. To summarize in short, the cancellation cost of these flights are sampled from the cancellation cost of flights of the reference airline, delay options are removed or added for some flights [redacted] This means that running the optimization model multiple times will result in slightly different solutions. For this reason, it was decided to run the optimization model multiple times to see the effect of the random elements, furthermore, variation in the input capacities will be tested as well. The set up of the framework is shown in Figure 4 with the goal of having a more reliable output by adding variation as well as including the uncertainty of other flights than the reference airline with the random parameters.

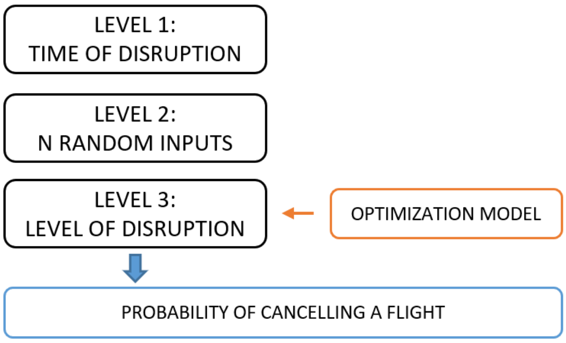

Figure 4: Multi level implementation of optimization model in the Proactive Cancellation Probability Framework

As can be seen in Figure 4, the framework consists of three levels. In level 1, variation of the disruption window is added. Imagine an expected disruption between 6:00 and 10:00. In the input, capacities will be lowered between these hours. However, this is still an expected disruption so start and end time might differ a bit which cannot be captured by the input as this is hourly. To account for this, a shifting disruption window was created as visualized in Figure 5. The green block in the middle represents the input. The blocks around have a slightly different disruption window. To be more precise, the start and end are either one period (20 minutes) earlier or later. The disruption window in the orange blocks have the same length as the input window, the ones in blue differ in length compared to the input. The shifting disruption window applies for both inbound and outbound disruptions in all scenarios. Figure 5 shows the input disruption window and 8 disruption window variations which means that the optimization model will be run 9 times with the slightly adjusted disruption window.

5:40 – 9:40	5:40 – 10:00	5:40 – 10:20
6:00 – 9:40	6:00 – 10:00	6:00 – 10:20
6:20 – 9:40	6:20 – 10:00	6:20 – 10:20

Figure 5: Shifting disruption window by 1 period (20 minutes) from the input window of 6:00-10:00

In level 2, N extra random inputs can be generated. As explained earlier, the random inputs follow from the fact that other flights than the reference airline are also taken into account. [redacted] Other random parameters for those flights are the maximum delay and the cancellation cost. Each time the optimization model runs, these random parameters will have new values, also in the different runs of level 1. In level 2, extra random runs can be added until a balanced solution is found. The value for N will be determined during validation. Ideally, N is as high as possible which will result in the most accurate result however, adding an extra random input, also means that that Level 1 will be run an extra time adding 9 extra runs of the optimization model. For this reason a trade-off between computational time and accuracy will be made. Level 3 denotes the level of disruption which is the input for the optimization model. The level of disruption is captured in several scenarios with corresponding probabilities of occurring.

The framework will have a total of *9 x N x #input scenarios* solutions. From this set of solutions, a robust set of flights that need to be cancelled will be extracted. The set of flights consists of two parts. Part one are flights that need to be cancelled proactively and thus a day before operations which are the same for all scenarios. Part two are flights that need to be cancelled on the day of operations and are scenario specific. The flights to be cancelled will be chosen from a list which is ordered by the probability that flight is cancelled in all solutions. This probability will be computed in the following way:

$$P(\text{Flight } i \text{ cancelled}) = \sum_{solutions} P(\text{Level 1}) \cdot P(\text{Level 2}) \cdot P(\text{Level 3}) \cdot Z_i$$

The probability of Level 1 is assumed to be 1/9 for each alternate disruption window as no difference in terms of probability can be determined. Also for Level 2, the random inputs are equally probable which corresponds to a probability of $1/N$. Probabilities of Level 3 are defined in the input scenarios and are most of the times not equally distributed. Applying above equation for all flights in all solutions results in a list of flights which are most probable to be cancelled. Note that for flights which are cancelled in the first stage, $P(\text{Level 3}) = 1$, flights which are cancelled in the second stage have a probability dependent on the scenario they are cancelled in. A probability that a flight will be cancelled of 1 means that the flight is cancelled in all solutions generated by the framework. For flights in the first stage this means that in all solutions this flight was cancelled in the first stage. For cancelled flights in the seconds stage, this means that in all solutions this flight was cancelled in scenario s. Now, in order to find a final set of flights that need to be cancelled that is robust for all solutions, this list of flights with associated probabilities will be used. Each solution suggests a certain amount of flights to be cancelled on the day before operations, ie. in the first stage. The average amount of cancelled flights over all solutions determines the amount of flights in the final set of cancellations. In this final list, flight rotations are taken into account which means that only full rotations are cancelled which may result in an extra cancellation above the average.

4.4.1 Output

A list of proactive cancellations required as this is the main outcome of the framework. Of course, it could be possible that there is also an additional set of flights which need to be cancelled on the day of operations itself, for this set also a list of probabilities will be created. Like the proactive cancellations, this set will also be shortened to the average amount of cancellations per scenario. In this way, once the airline knows which scenario is playing out, they have a list of flights, ordered by probability, to be aware off and start cancelling flights starting from the top of the list to make space in the schedule and limit the effects of the capacity reduction.

5 Case Study

This section presents the implementation of the framework in the operations of a major European airline with a hub and spoke network. At the hub airport, the airline operates more than 50% of the flights. It can be imagined in case of a runway capacity reduction, the hub and spoke network of this airline will be severely disrupted because of delays. This case study investigates the effects of using the framework in different weather scenarios resulting in runway capacity reductions. First, with the use of an historic disruption, the framework will be validated.

5.1 Case-Study Assumptions

For the case study a few assumptions have been made such that the framework can be used for the major European Airline. The model has been adjusted to these assumptions. These are listed below:

- No cruise speed changes to recapture time during flight (High Speed Flying)
- Future Value Losses are calculated using models of the airline
- Simple slot model is used
- Turnaround time is constant per aircraft type
- Crew reserves are defined per flight group

5.2 Framework Validation

Validation will be performed by simulating a historical disruption and comparing cancellation actions by the Operations Control of the airline and the results of the deterministic tool with the results of the framework. A day was found in which wind was expected to come from an unfavourable direction for the runway lay-out of the hub airport. This means that in peak hours only one runway could be used instead of two, subsequently runway capacity could half in peak hours. This is not a heavy disruption, but a good test for the model as delays are inevitable triggering several constraints regarding passenger connections and crew connections as well as aircraft availability. Once cancellations follow, the crew mislocations will also play a role. In order to perform the simulation with the same knowledge as available on the day before the expected disruption, the weather forecast of the day before the expected disruption was extracted from the local meteorological institute. Using this information, an expert in the industry with proven experience was asked to translate the probability forecast into several scenarios with corresponding probabilities. It was opted to ask an expert in the industry to translate the forecast into scenarios as this is not a trivial task and experience plays a great role in this. The weather forecast has been translated into four scenarios as shown in Figure 6. The first three scenarios imply a capacity reduction for both inbound and outbound with different starting and end times. The 4th scenario is the single scenario that was constructed by the Operations Control Team (OCT) which is used in a deterministic tool, this will be discussed later. In the last scenario, no capacity reductions are implied which means no disruption is expected by the expert.

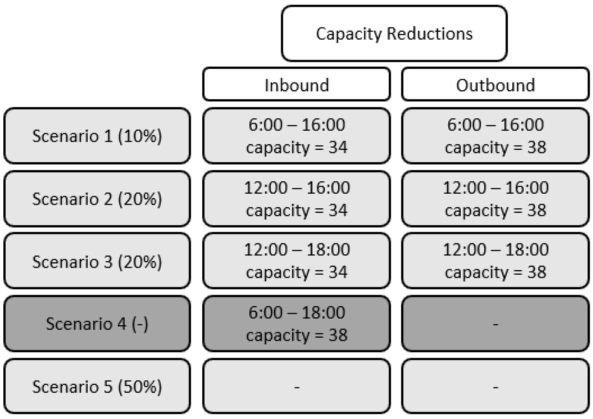

Figure 6: Capacity reductions used for validation of framework

The set of scenarios, as visualized in Figure 6, was implemented in the framework and with the general settings a solution was found. The settings of the model are listed in Table 5. The value for N as used in Level 2 of the framework will be chosen based on the results of the validation. In order to get a solution, a high number of N was initially used ($N = 10$). Scenario 1, 2, 3 and 5 have been determined by the expert and scenario 4 will be used for other validation work. The probability for scenario 4 is set to 0% so it will be neglected by the framework. According to the framework, it was not required to cancel any flights the day before operations. This result can be explained by the fact that the input consists of a scenario without disruption with a probability of 50%. In this way, it is cheaper to wait for the day of operations and see which scenario is playing out and consequently cancel flights if needed. The risk would be too high to cancel already flight proactively. For scenario 1, 2 and 3 there were some cancellations which could be implemented on the day of operations. Despite the framework not being in line with what was actually implemented, the result was in line with the expert's opinion, no proactive cancellations are required.

Parameter	Value	Unit
[redacted]		

Table 5: Settings for Model during Validation

Based on Scenario 4 as outlined in Figure 6, the OCT took some proactive actions to prepare for this expected disruption. A day before, in total 20 flights (10 rotations) were cancelled to allow some more space in the schedule and keep delays below a threshold. The results of the framework (PCPF) will be compared to these cancellations. Comparison will be made on multiple levels. First comparison will be made on the amount of cancelled flights and other parameters on flight level, then the arrival/departure time of cancelled flights, following comparing the flight numbers of cancelled flights. Finally the cost of solving the disruption will be compared.

Table 6 outlines the results on the validation. The table can be split up in two parts. Actions which have been performed or should be performed a day before operations and actions on the day itself. The results of the amount of cancellation a day before (24H) are computed using *expected* runway capacities, the results on the day of operations (DAY) have been created using actual declared capacities and the results are more detailed also taking into account passengers and cost related to passengers. Crew is not taken into account here because data on crew mislocations and crew misconnections in lacking as well as costs related to this.

Starting at the first row, 20 flights have actually be cancelled by the OCT a day before operations impacting ±1700 passengers. On the day itself no extra cancellations have been made. The average delay related to weather reasons and reactionary delays was 50 minutes [redacted] scenario 4 as depicted in Figure 6 as input for the DDST which suggested to cancel 54 flights impacting ±4800 passengers which would result in a simulated average delay of 27.4 minutes. Due to uncertainty of the disruption, the OCT did not cancel all these flights, but only a subset of 20. Using the deterministic tool on the day of operations with actual declared capacities for the day of operations (DDST (sc. 5)), the tool suggests to cancel 22 flights and simulates a total of 8 flights being cancelled by other airlines. This results in ±1800 impacted passengers of the reference airline. These 30 cancellations cause an average delay

	24H		DAY									
	Canx	Canx Pax	Canx reference airline	Canx other airlines	Canx Pax	Avg. Delay	Pax Misconn.	PP	Canx Cost	Pax Misconn. Cost	Delay Cost	Total Cost
Actual	20	±1700	0	0	0	48.3	[redacted]	[redacted]	100%	100%	100%	100%
DDST (sc. 4)	54	±4800	0	0	0	27.4	[redacted]	[redacted]	147%		-42%	-10%
DDST (sc. 5)	0	0	22	8	±1800	29.2	[redacted]	[redacted]	-14%		-25%	-38%
PCPF	0	0	24	4	±2100	28.1	[redacted]	[redacted]	-24%	-76%	-58%	-54%

Table 6: Benchmarking the actual implemented cancellations, cancellations from the deterministic tool (DDST) and the cancellations resulting from the cancellation framework (PCPF), comparing actions to be performed in the day before operations and on the day itself and the effect on costs

of about 29 minutes. The DDST however is unable to record passenger misconnections and for this reason, the PP value is unknown as well as one components deals with passenger being able to make their complete journey. Lastly, as stated in the previous paragraphs, the framework did not imply any proactive cancellation but rather wait for the day itself. Using the actual declared capacities as input for the framework results in 24 + 4 cancelled flights of respectively the reference airline and other airlines, impacting ±2100 passengers. The average delay is comparable with the previously discussed result of the DDST taking into account actual declared capacities. Notable is the clear difference in PP between actual observed and what would be achieved by implementing the results of the framework which is mainly caused by the actual high average delay. Also this difference in delay explains the difference in misconnecting passengers.

Next, the cancelled flights will be examined in more details, this will first be done by comparing at what moment during the day the flights have been cancelled. This is outlined in Figure 7. In this figure, three sets of cancelled flights are compared, which are the actual cancelled flights, and the cancelled flights according to DDST and PCPF using the actual declared capacities. The cancelled flights of DDST (sc. 4) are not compared as these amount of cancellations are clearly not comparable and the actual cancelled flights are a subset of these cancellations. Please note that in this plot, two types of cancellations are compared. The actual flights cancelled by the airline have been cancelled proactively however the flights of the DDST and PCPF are cancelled on the day itself. Despite this, this figure aims to show that the cancellations strategy used by the airline is much less effective then advised by the decision support tools.

The start and end time of the disruption is shown in the plot by respectively the red and green vertical lines. The actual declared capacity only entails an inbound capacity reduction. It can be seen that most of the departing flights are being cancelled in the same hourly brackets especially in the beginning of the disruption, also the peak of arriving flights in the bracket 9:00-10:00 is visible in both model, which could suggest common flights being cancelled. Actual cancelled flights do not follow these above mentioned patterns but cancellations are spread throughout the disruption as can be seen in the plot of departing flights. Referring back to Table 6, it can be seen, especially in the average delay and amount of misconnections that cancelling flights in this way proves to be much less effective than the way cancellations are spread in DDST and PCPF resulting in a lower value for average delay and much less passenger misconnections.

The two models suggest cancellations mainly within the same hourly brackets. Now, comparing the flights that are suggested to be cancelled by the models. The models suggest to cancel 22 and 24 flights of the reference airline for DDST and PCPF respectively. Amongst these flights, there are 16 flights in common, which is over 70%. 100% common flights will be hard to achieve due to the uncertainty of cancellations of other airlines first of all in the amount of cancellations but also in the determination of which flights are cancelled and in which hourly bracket they fall. The method of determining this also differs between the DDST and PCPF.

Lastly, the costs of solving the disruption are compared. These are also outlined in Table 6. The total costs are divided into cancellation costs, delay costs and passenger misconnection costs. There are more costs related to disruptions however only these costs are available for comparison. The costs which resulted from a solution from either the DDST or PCPF have been benchmarked against the actual costs. For the DDST, no passenger misconnection data is available which means the cost for this category is missing as well. When comparing the total cost, it can be seen that the cost would be lower if the solutions of the models would have been implemented. The solutions of DDST (sc. 5) and PCPF show a very comparable number of cancellations, also the cancellation cost is very comparable. Bigger differences can be spotted in the delay costs. Delay costs have been calculated as was done in the PCPF, by future value loss. Despite the average delay being comparable, in the DDST apparently more flights have been delayed which evidently result in more future value losses. An explanation for this can be found by the fact that the DDST uses a different way of modelling for flight delays. The passenger misconnection costs cannot be compared, however it is expected that these costs will be comparable to the PCPF passenger misconnection costs.

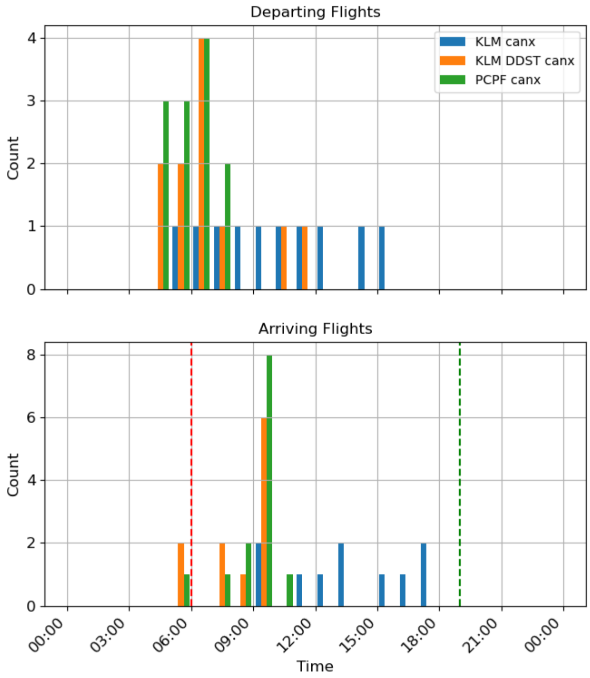

Figure 7: Comparing hourly division of flight cancellations between actual cancelled, DDST and PCPF

5.3 Framework Tuning

With the use of the validation of the results, an important parameter in the framework can be decided on. The framework is build up in three levels, where Level 2 is the amount of random inputs (N) on which the number can be decided on. This number should be a trade-off between computation time and reliable results. An extra random input means that the optimization model will run for each of the 9 shifted disruption window. Using the input which was also used for the comparison of DDST (sc. 5) and PCPF in Table 6, the framework was run 10 times (with $N = 1$) and the average amount of proactive cancellations over the 9 shifted disruption windows were saved. Based on these ten averages, a moving average line can be drawn. For example, choosing a value of $N = 3$, the average of *three* runs will determine the amount of proactive cancellations. The average amount of cancellations for all 10 runs is shown in Figure 8(a). Note that the amount of cancellations for each run is calculated by taking the average amount of cancellations from the solutions of the framework, in this way, this number is not always an integer.

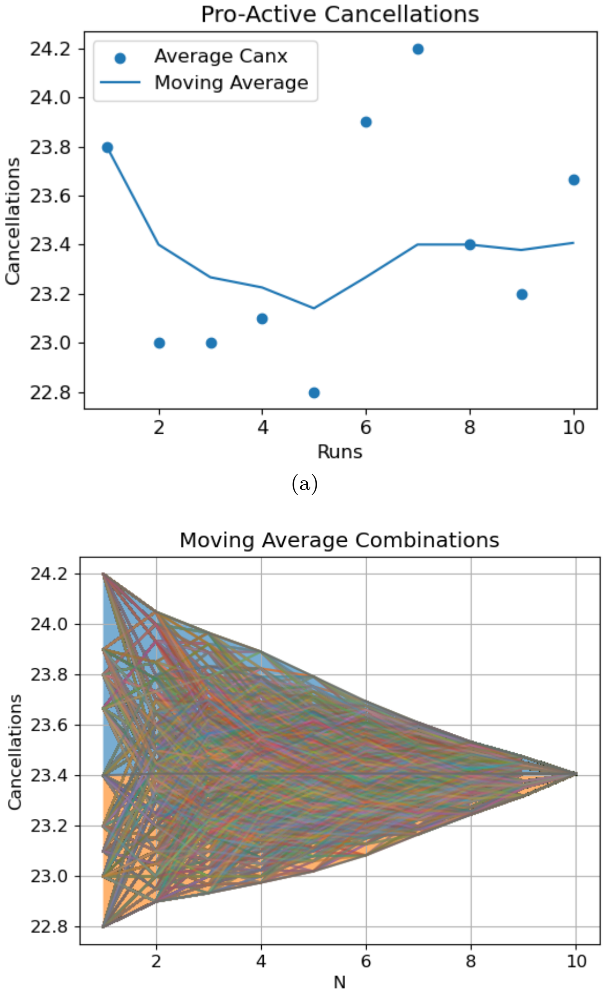

(a)

(b)

Figure 8: (a) Average amount of cancellations over scenario variations (Level 1) for 10 runs. (b) Moving average of the amount of cancellations with 1000 combinations of the set in (a) resulting in different moving average lines converging as N increases.

However, the moving average line will change if the order of points in Figure 8(a) is different. Since there are 10 runs, there are $10! = 3,628,800$ combinations in which these points can be ordered which has effect on the moving average line. From all these combinations, 1000 combinations have been randomly chosen. In Figure 8(b) it is clear how the upper and lower bounds of all the moving averages convergence into the average number of cancellations at $N = 10$. Now, of course it would be best to choose $N = 10$, or any higher number, as this is the convergence point. However as stated before, a trade-off will be made between a reliable result and computational time. Therefor, it was decided that N would be chosen at the point where the difference between the upper and lower bound is less than 1. From $N = 4$, the difference between the bounds is just below 1, therefor, $N = 4$ will be used in the framework. For this example, $N = 4$ means the minimum of the moving average is 22.97 and the maximum value 23.89 which in both cases will mean that 24 flights will be cancelled as only full rotations can be cancelled.

5.4 Case Study Results

In this section the framework will be used for an expected disruption and results will be reported and analyzed. The complete schedule at the airport is shown in Figure 10, where the dominance of the reference airline with their hub and spoke network can be clearly seen. There are several banks throughout the day with a peak of arrivals or departures, and quiter moments in between. For this case study a cold (freezing temperature) and misty (limited visibility) day is expected. This implies that there is a chance that both inbound and outbound capacities are reduced. Outbound capacity is reduced due to de-icing activities. At the airport, the de-icing capacity is limited which means that in peak hours, not all departing aircraft can be de-iced at their scheduled departure time which causing a limited outbound capacity. The inbound capacity may be reduced due to limited visibility. Lower visibility implies the need for increased spacing between aircraft which reduces the overall airport capacity. Translating the weather forecast for this day into a set of scenarios has been done by an expect in the industry keeping in mind the inbound capacity levels for different visibility levels and estimating the de-icing capacity. The set of scenarios is presented in Figure 9.

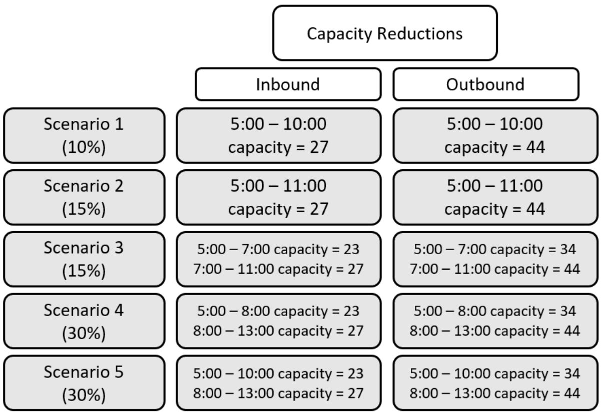

Capacity Reductions	Inbound	Outbound
Scenario 1 (10%)	5:00 – 10:00 capacity = 27	5:00 – 10:00 capacity = 44
Scenario 2 (15%)	5:00 – 11:00 capacity = 27	5:00 – 11:00 capacity = 44
Scenario 3 (15%)	5:00 – 7:00 capacity = 23 7:00 – 11:00 capacity = 27	5:00 – 7:00 capacity = 34 7:00 – 11:00 capacity = 44
Scenario 4 (30%)	5:00 – 8:00 capacity = 23 8:00 – 13:00 capacity = 27	5:00 – 8:00 capacity = 34 8:00 – 13:00 capacity = 44
Scenario 5 (30%)	5:00 – 10:00 capacity = 23 8:00 – 13:00 capacity = 27	5:00 – 10:00 capacity = 34 8:00 – 13:00 capacity = 44

Figure 9: Capacity reductions used for case study

Using these scenarios, and the same settings as listed in Table 5, with $N = 4$, the framework was run. This resulted in 38 proactive cancellations. For each scenario extra cancellations are required on the day itself. Despite some of these flights being cancelled in all scenarios, the model oughts these flights to be cheaper to cancel on the day of operations, hence no proactive cancellations for these few flights. Especially for scenario 4 and 5 a greater set of flights need to be cancelled on the day of operations to fit the amount of flights in the available capacity.

Out of the 19 rotations that should be cancelled a day before the disruption, 16 are identified as nightstops, indicating rotations with a night overlay at an outstation. The remaining three rotations involve flights that depart and arrive at the airport on the day of the disruption. This distribution can be explained by the disruption primarily occurring in the morning and causing a greater reduction in inbound capacity than outbound capacity. The cancellation of nightstop rotations can be clarified by their early arrival times, allowing flights to be cancelled in the early phase of the disruption. Similar findings are observed in the validation case, where the majority of the flights are cancelled in the beginning of the disruption.

To gain insight into the specific flights that are cancelled and the reasons behind their cancellation, an examination is conducted with respect to cancellation cost. It would make sense for the model to choose cheap flights however passenger and crew connections, amongst others, could play a role as well. Initially, the 25 cheapest flights on the day of the disruption are compared with the cancellation list, among these flights three rotations exist in the list of cheapest flights. Interestingly, these rotations do not have a night overlay at an outstation and thus are rotations on the day of disruption. Subsequently, the analysis is extended to the remaining 16 rotations, all of which are nightstops. Another comparison is made, this time considering the 25 cheapest flights arriving before 12:00. It was found that 14 out of the 16 nightstop rotations are present in this list. Thus, the majority of the cancelled rotations are associated with lower-cost flights. However, there are two cancellations involving more expensive rotations, each with distinct explanations.

One of the two rotations has numerous outbound connections that incurred higher costs when these connections would be missed, compared to the cancellation cost of the corresponding inbound flight. The cancellation of the other rotation is motivated by crew reasons, as there was no deadheading crew on the rotation resulting in mislocated crew members. Also, only the cockpit crew of the inbound flight has an onward connection, by cancelling this inbound flight, the crew is not able to perform the next flight, resulting in only using a reserve cockpit crew to perform the onward flight. From this analysis it can be concluded that the majority of the cancelled flight have been cancelled because of their cancellation cost, which was also expected. However, the results also show that not simply the set of cheapest flights have been chosen, but also considerations regarding passengers and crew are taken into account.

Previous results are obtained considering all nightstop rotations. However, the outbound leg of the nightstop rotations are already departing at the day before operations. The time at which the framework is used a day before operations is not fixed which means that if the framework is used later on the day, it could be possible that outbound flights have already departed or are departing soon. This will have effect on the results as these flight cannot be cancelled, or are not desirable to be cancelled. For this reason, the framework was used again, with the same input, but no nightstops could be cancelled representing a situ-

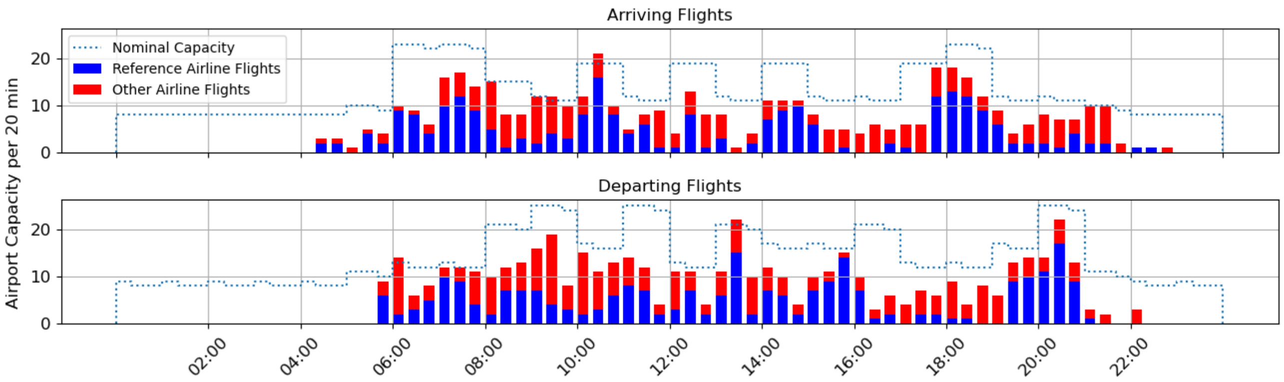

Figure 10: Overview of all scheduled flights at airport including nominal runway capacity at the day of the case study. Every bar represents a 20 minute window.

ation where the framework was used in the evening when all outbound flights of the nightstop rotations have already departed. Results show that instead of 36 proactive cancellations, the model only suggests 10 proactive cancellations however a lot more cancellations on the day of operations.

An interesting derived decision from the framework is the amount of available aircraft throughout the day. It was decided not to assign flights to aircraft tails however check only if there are enough aircraft of a certain type available to perform the flights that are scheduled at any moment in time. This is visualized in Figure 11, for aircraft type [redacted] For every scenario, as depicted in Figure 9 (5 in total) a line represents the available aircraft throughout the day. [redacted] Furthermore, a line is plotted for the available aircraft in the original schedule, the day without any disruption, to compare how this differentiates from the availability in the scenarios. When taking a closer look at the differences between the available aircraft in the original schedule and the scenarios, it is clear that there is a gap of 6 aircraft between the lines at 00:00. This originates from the fact that nightstop rotations have been cancelled yielding extra available aircraft. Then, for scenario 1,2,3 and the original schedule, there is a big drop of available aircraft at 9:00, which means that almost all aircraft are flying or at an outstation. Also at 9:00, the available aircraft of scenario 4 and 5 do not show this big drop. This follows from the fact that also in this period in the day, for these scenarios flights have been cancelled which leaves extra aircraft available.

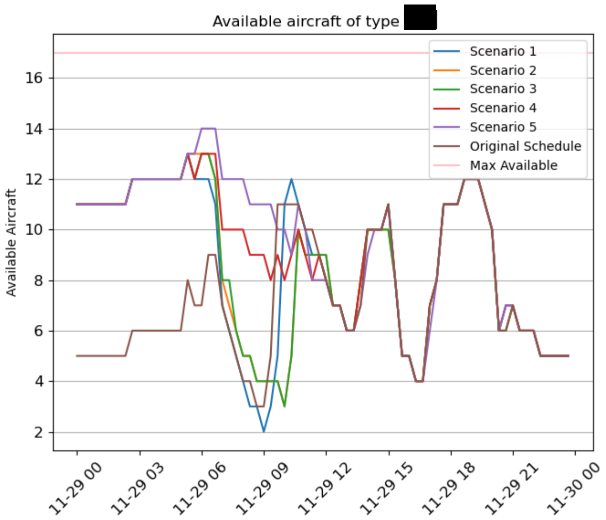

Figure 11: Available aircraft of type [redacted] throughout the day of disruption

5.5 Machine Learning Model Results

As outlined in Section 4, a machine learning model was used to estimate the cost of cancelling a flight on the day of the disruption when knowing the cost of cancelling the flight on the day before. A model was created for arriving and departing flights, with a different set of features. The complete dataset consists of almost 13,000 flights of which the cost of cancelling a day before operations and the cost of cancelling on the day of operations is known. From this, the cancellation factor can be calculated, which is defined as the cost of cancelling on the day of operations divided by the cost of cancelling on the day before operations. After removing outliers, based on the assumption that the data is normally distributed such that outliers are identified as values which fall outside the mean ± two standard deviations, a set of 12,000 flights remain.

[redacted]

Figure 12: Feature importance plot of both models

The features for each model have been listed in Section 4. This set of features have been chosen based on a feature importance plot. Variables with low feature importance have been removed from the model to simplify the model. The final feature importance plot is shown in Figure 12. It can be seen that similar features in the two models do not have the same feature importance. For example, the time of arrival seems to be much more important in the model for arriving flights than the time of departure in the model for departing flights. In both models, the percentage of connecting passengers on the flight is the most important feature. This was also expected as the cancellation cost is a cost related to passengers and depends on the rebooking possibilities for the passengers.

[redacted]

Figure 13: Accuracy results of machine learning model

Figure 13 shows the absolute error of the two models. Absolute error is a measure of how far off the predictions are from the actual values. By calculating the average error, the accuracy of the models can be determined. These results have been gathered using a 80/20 ratio for the train-test set. The results show that the arriving flights dataset model achieves an accuracy [redacted] These accuracies have been opted high enough to use the machine learning model in the framework. If more data is available, one potential approach for improving the accuracies of these models is to expand the dataset. By increasing the dataset size, the models can learn from a larger and more diverse set of examples. Another potential improvement could be a more extensive research on other relevant features that show a high importance.

5.6 Analysis of Sensible Input Parameters

This section presents a sensitivity analysis on two sensible input parameters. Actually, there is a third sensible input parameter, which is the start and end time of the disruption time. However, as explained in the set-up of the framework, this parameter has been included in the framework already and its sensibility is included in the results. The two other sensible parameters are the hourly division and the maximum delay for flights of the reference airline and delay of flight of other airlines which will both be discussed.

Hourly Division
The hourly division defines the length of a slot time in which flights are assigned to as well as the periods flights can be delayed. [redacted] In this sensitivity analysis it will be tested what the effect is of increasing the hourly division, ie. reducing the slot time. The slot times tested are 20 minutes, 15 minutes, 10 minutes and 5 minutes. The results are shown in Figure 14. Increasing the hourly division from 3 towards 12 yields some interesting results. First of all, when increasing the hourly division, the slot times decreases which means that the day is divided into more slots. This means that there are more delay options for a flight. As such, the amount of variables and constraints is increasing when increase hourly division. The increase is linear. The logical result of more variables and constraints is an increase of runtime. This is also visible in the plot, however this increase is not linear, as expected. Then looking at the amount of proactive cancellations. An increase of cancellations is visible with an increase of the hourly division. However, contrary to expectations, the amount of cancellations is not decreasing with increasing hourly divisions. Increasing hourly division results in shorter slot times so flights are provided with additional opportunities for rescheduling and accommodating delays. This flexibility allows flight schedules to be modified more precisely to minimize the impact of the disruption. However, with more options available for rescheduling, there is also an increased likelihood of finding scenarios where cancellation becomes the optimal choice for mitigating the effects of runway capacity

reductions.

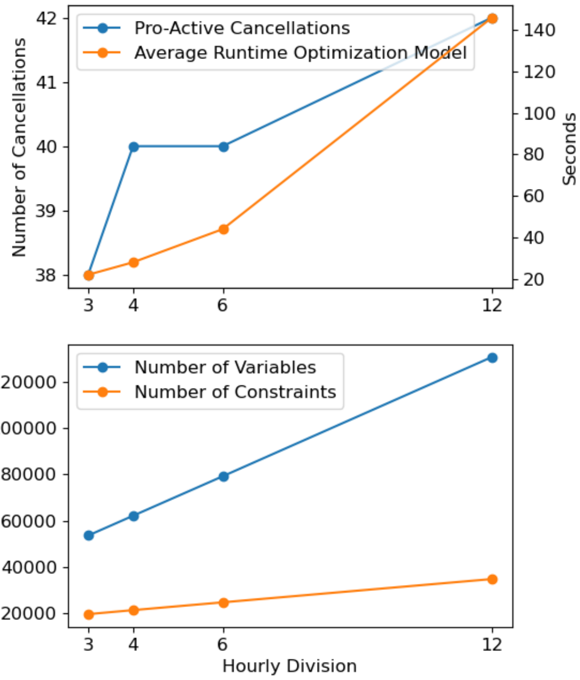

Figure 14: Sensitivity of hourly division on number of cancellations and runtime of optimization model

Maximum Delay
Another important parameter in the optimization model is the maximum delay. Flights cannot be delayed unlimited and the range is defined by the maximum delay. The maximum delay is defined for flights for the reference airline and other flights. [redacted] These numbers have been chosen by an expert in the field. During this sensitivity analysis, the effect of changing the maximum delay on the amount of proactive cancellations will be tested. Since the maximum delay is actually defined by two numbers, two tests will be performed. First the maximum delay of other flights than the reference airline will remain fixed, [redacted] The maximum delay of flights of the reference airline will range between 40 till 120 minutes with steps of 20 minutes. In the second test, the maximum delay of other flights will depend on the maximum delay of the reference airline. [redacted] Both results have been plotted in Figure 15. Results of this sensitivity analysis are as expected. In both cases, the amount of proactive cancellations decrease with increasing maximum delay. The slope of decreasing is different for both cases. Using a linear trend line, the slope of the lines can be estimated and compared. In case of the fixed maximum delay for other flights than the reference airline, increasing the maximum delay by 20 minutes results in an average reduction of 3.4 cancellations. In the case of the dependent maximum delay, this decrease is approximately 8 cancellations with an increase of 20 minutes of the maximum delay of the reference airline. The steeper slope is caused by the extended maximum delay of other flights than the reference airline. With an increase of 20 minutes of the maximum delay of the reference airline, the maximum delay of other flights increase with 40 minutes which gives these flights two extra available periods to be delayed to.

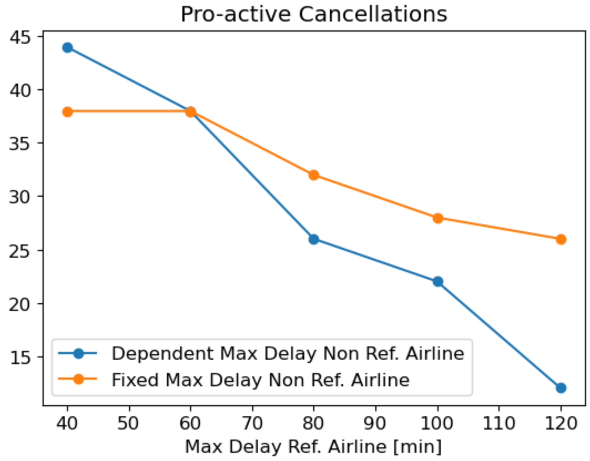

Figure 15: Sensitivity of maximum delay on number of proactive cancellations considering two options for maximum delay of flights other than the reference airline

6 Conclusions and Recommendations

This paper presented a decision support tool for day-before cancellations. In case of an expected disruption in the form of a runway capacity reduction at the hub airport of an airline, this framework provides the airline with proactive cancellations. One of the key contributions of this research is its novel approach to proactive disruption management. Surprisingly, the literature on this topic is lacking, with limited studies focusing specifically on evaluating expected disruptions one day prior to operations. The absence of existing models addressing this aspect emphasizes the importance and uniqueness of the framework presented in this paper.

The framework consists of three levels. Level 3 represents the optimization model to solve the disruption, fed by a set of scenarios. Level 2 extra random runs can be added to the model and level 1 add a variation of the disruption window to include the sensibility of the disruption window. The framework outputs a list of proactive cancellations with associated probability. This probability is determined by the existence of a flight in a list of cancellations of a solution. If the flight exists in all solutions, its probability will be 1 and the flight will be most likely to be cancelled in the final solution. This proactive strategy provides significant benefits, allowing airlines to optimize resources, reduce delays, and improve overall operational efficiency.

A case study was developed in collaboration with a major European airline and presented to validate the framework and show its capabilities. The goal of this case study was to recreate the preparation process for a disrupted day, allowing for a evaluation of the framework's performance. The study involved comparing the framework's results to two key sources: actual implemented airline cancellations and the outcomes generated by a deterministic model commonly used by the major European airline. The case study findings demonstrated the framework's capabilities in providing cost-effective cancellation options. In comparison to the airline's actual implemented cancellations and the cancellations of the deterministic model, the framework identified more cost-effective alternatives. Furthermore, the case study revealed that the framework outperformed the deterministic model in terms of overall cost savings when it came to dealing with disruptions. The framework's decision-making process proved to be more cost-effective than the deterministic model's outcomes by taking into account multiple factors and parameters specific to the airline's operations. This demonstrates the added value of using a proactive approach to disruption management, as made possible by the framework.

To further enhance the capabilities and applicability of the framework, several recommendations can be made. First, the model could be modified to include the uncertainty of arrival and departure times rather than using scheduled arrival and departure times. By considering factors such as delays, early arrivals, the framework can provide more accurate assessments of flight cancellations. Another recommendation is to allow for flight subtype swapping within the model. By introducing this flexibility, the model can identify more efficient options for flight scheduling which would mean a better utilization of available resourcing ultimately reducing the overall impact of disruptions. Incorporating maintenance flexibility is also advised. Currently, the model assumes fixed maintenance slots that cannot be delayed or canceled. However, by introducing the ability to adjust maintenance slots based on the severity of the disruption, the model can optimize flight schedules within maintenance schedules. Additionally, conducting a study on the relationship between cancellations of the reference airline and cancellations of other airlines during disruptions could add a lot of value to the model. Other airlines are reacting different to different types of disruptions. As for other airlines the airport is most likely not their hub airport, a disruption at this airport has a much smaller effect on the network of that airline. Therefor, an analysis on the cancellation strategy of other airlines would result in a more realistic implementations and therefor better decisions by the framework. In addition to the aforementioned recommendations, future studies could also consider the examination of consequences on the day after the disruption. Understanding how disruptions affect the following day(s) can provide useful insights for developing more robust and proactive mitigation strategies.

References

Abdelghany, A. and Abdelghany, K. (2018). *Airline Network Planning and Scheduling.* John Wiley Sons.

Aloulou, M. A., Haouari, M., and Mansour, F. Z. (2010). Robust aircraft routing and flight retiming. *Electronic Notes in Discrete Mathematics*, 36:367–374.

Ball, M., Barnhart, C., Nemhauser, G., and Odoni, A. R. (2007). Chapter 1 air transportation: Irregular operations and control.

Bouarfa, S., Blom, H., Curran, R., and Hendriks, K. V. (2014). A study into modeling coordination in disruption management by airline operations control. *AIAA AVIATION 2014 -14th AIAA Aviation Technology, Integration, and Operations Conference.*

Cook, A. (2015). European airline delay cost reference values - updated and extended values (version 4.1).

Coulston, J. W., Blinn, C. E., Thomas, V. A., and Wynne, R. H. (2016). Approximating prediction uncertainty for random forest regression models. *Photogrammetric Engineering & Remote Sensing*, 82(3):189–197.

Eurocontrol (2019). Annual network operations report 2018.

Eurocontrol (2020). All-causes delay and cancellations to air transport in europe 2019.

Ghaemi, N., Cats, O., and Goverde, R. M. P. (2017). Railway disruption management challenges and possible solution directions. *Public Transport*, 9:343–364.

Hassan, L., Santos, B., and Vink, J. (2021). Airline disruption management: A literature review and practical challenges.

Hoffman, R., Ball, M., Odoni, A., and Rifkin, R. (1999). The static stochastic ground holding problem with aggregate demands.

Lee, J., Marla, L., and Jacquillat, A. (2020). Dynamic disruption management in airline networks under airport operating uncertainty. *Transportation Science*, 54:973–997.

Mukherjee, A. and Hansen, M. (2007). A dynamic stochastic model for the single airport ground holding problem. *Transportation Science*, 41(4):444–456.

Odoni, A. R. (1987). *The Flow Management Problem in Air Traffic Control.* Springer Berlin Heidelberg, Berlin, Heidelberg.

Ogunsina, K. and Okolo, W. A. (2022). Artificial neural network modeling for airline disruption management. *Journal of Aerospace Information Systems*, 19(5):382–393.

Richetta, O. and Odoni, A. R. (1993). Solving optimally the static ground-holding policy problem in air traffic control. *Transportation science*, 27(3):228–238.

Richetta, O. and Odoni, A. R. (1994). Dynamic solution to the ground-holding problem in air traffic control.

SimpleFlying (2022). How airport slots are allocated and why they are worth so much. Last accessed 20 April 2023.

Szabo, S., Liptáková, D., and Vajdová, I. (2015). Robustness as a method of airline pro-active disruption management. *International Review of Aerospace Engineering (IREASE)*, 8(4):151–156.

Terrab, M. (1990). Ground holding strategies for air traffic control.

II

Literature Study

previously graded under AE4020

1

Introduction

In 2019, before the Corona crisis hit the aviation industry, passenger load factors were a record high of 82.6% world wide, caused by a capacity expansion of only 3.6% which is a slower pace than that of the demand (IATA, 2019). Demand growing faster than capacity can cause capacity-demand imbalance which can be the root cause of airport congestion. By carefully planning the scheduled flights, airports can fully utilize the maximum runway capacity and minimize delays. This careful planning is possible because the maximum runway capacities are known and the demand consists of predictable flight schedules. Greater capacity-demand imbalances will appear when the runway capacity experiences a sudden significant drop. The main cause for this is weather.

Big capacity-demand inequalities can cause arrival and departing aircraft to queue up and delays start to build up quickly. Minimizing disruptions is one of the main objectives of every airline because they are very expensive. In the latest estimations by Cook (2015), delays cost European airlines a total of 1.25 billion Euros in 2014. Disruptions occur daily in airline operations. Some disruptions are unexpected and unforeseen, such as technical problems, medical emergencies and delayed turn-around activities, to name a few. Greater disruptions occurs in cases of staff shortages, crew unavailability and airport capacity issues. Despite the fact that these disruptions are far more worse for the network, they are not unforeseen and allow the airline to prepare. Preparing an expected disruption properly can save a lot of costs and reduces workload on the day of the disruption. On the other hand, preparing for the disruption a day-before-operations introduces uncertainty which complicates the problem. In this study, it will be investigated how an airline could prepare for a specific disruption, runway capacity reduction, a day-before-operations.

The document is structured in the following way. chapter 2 will start with a short introduction of the problem followed by explaining disruptions in airlines. In chapter 3 the concept of runway capacity will be explained, furthermore models are presented to estimate runway capacity in nominal conditions as well as in bad weather conditions. chapter 4 presents models that solve the Ground Delay Program (GDP), which are initiated when runway capacity is significantly reduced. Deterministic and stochastic models are discussed. In addition, solution methods are discussed which are needed to solve stochastic problems. chapter 5 describes models that deal with slot allocation, another way to model the problem. Based on the previous chapters, chapter 6 points out some gaps found in the literature. Finally, the research question is presented in the last chapter as well.

2

Disruptions

This chapter discusses disruptions in airline operations. In section 2.1, the problem, as shortly outlined in the introduction, will be further explained and linked to a type of disruption. In section 2.2, the concept of irregular operations is outlined which involves disruption management techniques and planning periods.

2.1. Problem Description

In this section, the problem as shortly outlined in the introduction will be described in more detail to understand the type of disruption that is initialized by the capacity reduction and its impact on an airlines network. The problem has the focus on airlines which make use of the hub-spoke-network and have a hub airport where the airline is the main operator. As described in the introduction, airlines experience severe disruptions when the runway capacity of their hub airport is lowered significantly. This is due to way the airline build its network namely by the means of hub-and-spoke.

A hub-and-spoke network allows the operation to be mainly focused on connecting passengers and by smartly assigning flights to so-called banks, the amount of connecting passengers with a minimum connection time can be maximized. This banking system is also known as the wave system which defines multiple peaks of arrivals and departures throughout the day. The idea behind this wave system is that a cluster of flights will arrive before a certain time and departs a little later allowing the passengers of the arriving flights to transfer to the departing flights. This banking system proves to be very efficient for airlines.

Of course there are also downsides to this type of network. During peaks, the ground handling experiences a high workload as well as passenger handling due to the peak of flights arriving and departing in a short time. This requires tight planning of resources. This type of network does also impose problems when flight schedules are disrupted. Having aircraft coming in late means passengers missing their connection and having to be rebooked to other flights. As will be discussed in chapter 3, weather issues might cause runway capacity reduction. Hub airports are most often operating with the banking system, which means that during the arrival and departure peaks the airport is operating at maximum capacities. If for any reason the runway capacity is reduced during a peak major delays will develop. These delays will propagate through the network and a disruption in the morning can still cause delays in the evening.

During these disruptions, the airline which has the airport as their hub airport will experience the most problems, as most of their flights will be delayed. These airlines can choose to do nothing and wait and see how the delays will start to develop and cancel flights on the day-of-operations when the delays are exceeding a certain time. Differently, airlines could prepare an expected disrupted day. This means, in case of a possible capacity reduction, measures will be taken to minimize the impact. Advantages of this way of operating is that flights can be cancelled already to reduce delays and passengers can be rebooked pro-actively. This reduces the workload on the day-of-operations.

The model to be created needs to determine which flights need to be cancelled the day before operations to get a feasible schedule on the day of the expected disruption. The conditions for a feasible schedule can be set by the airline itself where the maximum delay a flight experiences is one of the main factors as well as the average delay. If this maximum delay or average delay is too high, it is required to cancel flights to make space for other flights. In other words, it is required to find the most cost-efficient way to return to the original schedule in case of irregular operation.

2.2. Irregular Operations

Disruptions occur very regularly in airline operations resulting in not being able to follow the planned schedule. In literature, this is being referred to as irregular operations. That disruptions are very common is proved by Eurocontrol (2019), which showed that in 2019 only 59% of the flights arrived on scheduled arrival time or before. Within airlines it is common to say that an aircraft is still on-time if the delay is maximum 15 minutes. Taking this into account, a little over 77% of flights arrived on-time in 2019 (Eurocontrol, 2019), this still implicates that about a quarter of the flights is delayed each year.

Disruption management has been studied a lot in literature for the fact that disruptions occur regularly and cause serious problems in the network of airlines. Disruption management is not only a hot topic in airlines but also in other industries such as railway industry (Ghaemi et al., 2017). Actually all industries that are working with schedules have to deal with disruption management, some more complex than others. Carefully planned schedules can become infeasible due to disruptions. Factors causing disruptions in airlines are too many to list, but very generally they can be divided into *airline resource shortages* and *airport and airspace capacity shortages*. Solving these disruptions is known as disruption management (Ball et al., 2007). Lee et al. (2020a) classify disruptions into three categories which also captures the propagation of the disruption in the network (Lee et al., 2020a).

- **Contingent disruptions:** inefficiencies in airline and passenger operations such as late passengers, late crew, aircraft maintenance and late aircraft, to name a few. These disruptions are airline related.
- **Systematic disruptions:** imbalances in demand-capacity at airport or en-route.
- **Propagated disruptions:** disruptions spreading through the time-space network due to interrelations.

This paper emphasized the significance of propagated disruptions through the network of an airline. Again as described in the paper of Ball et al. (2007), hiding under the flight schedule, there are three other layers of schedules; aircraft schedules, crew schedules and passenger schedules. Given the fact that these are all related, suggests that a disruption on a single flight leg can have significant 'downstream' effects which can lead to several delays in other flight legs.

The urge to solve the disruptions as fast and efficiently as possible comes from the fact that disruptions are very expensive for airlines. The cost of disruptions can be divided into hard and soft costs (Hassan et al., 2021). Hard costs are costs that are made and should be directly paid to the customer, such as compensation, hotel costs, vouchers and rebooking costs. In addition to this, there is also the concept of soft costs. There are referred to as potential future revenue costs which are missed due to the disruption. These costs differ per passenger class and flown distance. Despite this, it is quite complex to estimate soft and hard costs and they are most of the time estimated in literature by using models of Cook (2015).

2.2.1. Disruption Management

As described above, hidden under the flight schedule are the schedule of aircraft, crew and passengers. These three schedules need to be recovered in case of disruptions. Teodorovi and Guberini (1984) were the first to address the aircraft recovery problem. In the model the authors considered cases with aircraft unavailability with the goal of minimizing passenger delays. Swapping aircraft and delaying flights were recovery options. This study initialized aircraft recovery models which have been expanded over time. Several disruption types and other recovery options have been researched in the past years. Other recovery options are cancelling flights and using the concept of aircraft ferrying. Including several disruption types and more recovery options significantly complicates the problem. Hassan et al. (2021) created an overview of the literature on aircraft recovery and they concluded that for over 80% of the publications heuristics were needed to solve the problem in a reasonable time. Clausen et al. (2010) states that solving the ground delay problem (GDP), as discussed in detail in chapter 4, is a special type of aircraft recovery where the problem is modelled as an assignment problem (assigning aircraft to available slots).

Crew recovery problem is considered to be much more complicated than the aircraft recovery problem. In crew recovery problems, crew is re-assigned to a set of flights given disruptions and the cost of re-assigning them is minimized in the models. Complexity in this problem originates from the fact that there are a lot of regulations and restrictions regarding crew (Hassan et al., 2021). Interesting to note, in the period 2009-2018 there have only been six publications on the crew recovery problem. There are almost no models that solve the problem exact and other types of solution methods are used to solve the problem, such as heuristics and multi-agent systems due to the complexity of the problem.

Lasty, passenger recovery should be addressed. For airlines this is a very relevant problem as there are high costs related to this, as described in previous section. A disrupted passenger is defined as a passenger which is not able to complete its booked itinerary or a passenger who arrives delayed on the final destination (Bratu and Barnhart, 2006).

Studying these recovery problems in isolation leads to solving the problems in a sequential way. First recover the aircraft, based on that solution re-assign crew and based on that solve the passenger recovery problem. Sequential optimization techniques frequently produce sub-optimal recovery solutions because they do not properly account for the interactions between the aircraft, crew, and passengers (Hassan et al., 2021). Despite this, computational times to find the optimal value are reasonable as it is the sum of the computational time of the three problems combined. To overcome the downside of finding a sub-optimal value, the problem can also be solved by solving the problems simultaneously, in an integrated way. This model will find the optimal value however computation times are significantly higher than solving the problem in a sequential way. For this reason, it is uncommon in literature to integrate aircraft, crew and passenger recovery. Studies that were able to solve the problem are not useful for operational decisions since the computation times are reaching several minutes, even for smaller problems (Hassan et al., 2021).

2.2.2. Planning Periods

Within airline planning there are several stages in which different processes are performed. A general airline planning process can be seen in Figure 2.1. The time horizon runs from long term to short term. In this context long term could be from ten years to one year before operation of the flight. Short term is from one year until one day before operations. In this planning period, planning starts strategic and is tactical closer to the day of operations. During this planning periods, the main objective of the airline also changes. Within strategic planning, the airline defines its strategy and vision and tries to match the forecasted demand with their resources, with the main objective to match the demand with supply. Within the tactical phase, airline's main focus is to maximize the revenue by readjusting to the market. Close to the day of operations, the airline mainly wants to minimize the costs, this period starts about four weeks before operations. On the day of operations, obviously additional costs are minimized (Bouarfa et al., 2014). As visualized in Figure 2.1, generally, disruption management is done on the day of operations where decisions are operational (minimize costs) instead of tactical (maximize revenue).

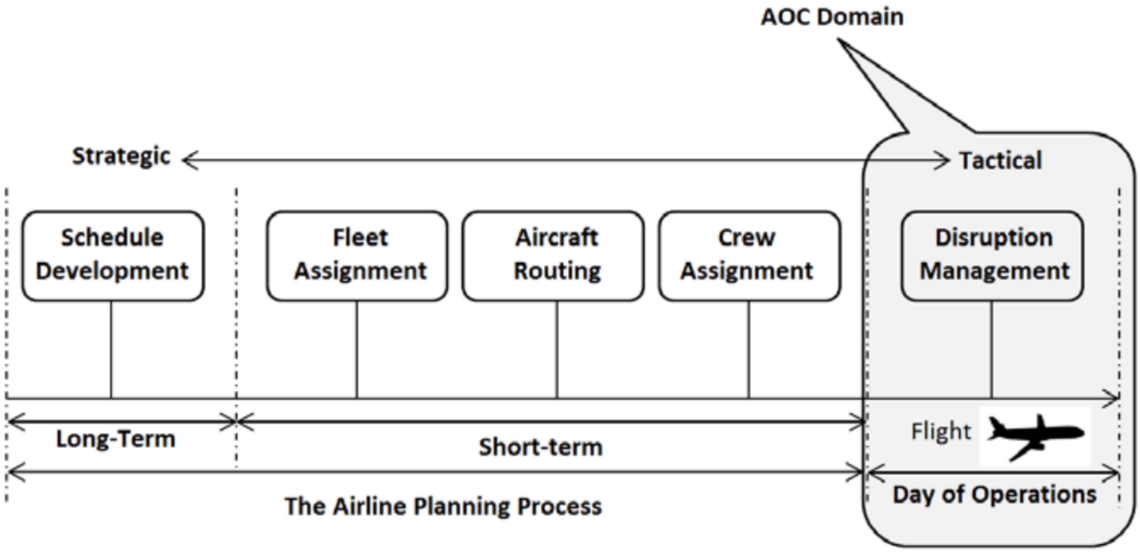

Figure 2.1: The airline planning process highlighting the strategic and tactical phase and corresponding tasks performed by an airline. Bouarfa et al. (2014)

As became clear in the previous section, disruption management is about mainly recovery techniques and recover aircraft, passengers or crew in case a disruption has happened. As the name 'recovery' suggests, this is also known as a reactive approach. Hassan et al. (2021) found that apart from a single article, all disruption management techniques follow a reactive approach. This leaves a clear gap in literature of proactive disruption management.

Proactive disruption management is still a wide topic in which several studies can be done. The single article found in literature about proactive disruption management is an article which combines solving current disruption by anticipating on future disruptions as well (Lee et al., 2020b). By taking the flight schedule and airport capacities as an input, a stochastic dynamic queuing model is created which can forecast any congestion disruptions. Based on this observations, disruption management techniques are implemented for aircraft recovery. According to the author, the results of the complete model motivate for future research on optimizing airline recovery under uncertainty. The current model could be extended by also recovering passengers and crew and could be made dynamic by including updated information.

In literature, often the proactive approach to disruption management is referred to as techniques that avoid disruption in the first place, and limit the consequences. Very often, the term robustness is used where robust planning is recognized as a proactive disruption management technique (Szabo et al., 2015). Robust planning aims to produce feasible plans that can withstand uncertainties while also keeping extra costs as low as possible. Especially for hub-and-spoke airlines, a robust schedule is very important to limit the impact of disruptions. Cadarsoa and Marín (2011) proposes a model to include robustness in the schedule of hub-and-spoke airline with the focus on connecting passengers. By allowing some extra connection time between high demand itineraries, passenger are more probable to get their connection in case of a disruption. Of course, this robustness comes at a price which is also captured by the model. As becomes clear from this model, it is focused on schedule development, which is, when referring to Figure 2.1, in the strategic planning phase. Disruption management approaches are seen as proactive as long as they are implemented before the actual disruption occurred. This also means that uncertainty is a term which is closely related to proactive disruption management and robustness (Szabo et al., 2015). As the exact nature and impact of the disruption is unknown, uncertain parameters can be modelled into scenarios which can be used as an input for a stochastic model. Stochastic modelling will be discussed later in the report.

Conclusion

In recent years, the literature trends to solve disruption related problems in an integrated way, where passenger/crew recovery are modelled as a part of the aircraft recovery. The second trend to be observed is the use of more functionalities which are thereby allowing more detail in the model which leads to more realistic outputs. Both of these trends also relate to the fact that airline disruption management models are more often used in airline operation control centers where integrated models and more detailed formulations are required to have a model that is better able to represent the real-world. On the other hand, these two trends come at the expense of computational time. Currently this is the biggest challenge for researchers. The first papers have been published which use machine learning algorithms to solve the problem. Vink et al. (2020) is a great example of the use of machine learning for a selection algorithm. Furthermore, Hassan et al. (2021) suggest that there is another aspect in disruption management which requires more detail, which relates to proactive disruption management. Currently most papers present a reactive approach which solves the disruption most efficiently after occurrence. Proactive disruption management can be used to avoid disruption or limit the consequences. Currently, proactive disruption management techniques are mainly implemented in the strategic planning phase by means of robust scheduling. This can be extended towards the tactical planning phase, however there is very limited literature on this topics. This is a gap in the current literature related to disruption management.

3

Estimating Runway Capacity

An estimate of the airport runway capacity is one of the most important elements in traffic management initiatives (TMIs) used today. Poor predictions or estimations, both under- and over-predictions will result in unnecessary delays. These delays might occur in the air or on-ground. However, it is challenging to predict airport capacity accurately since it depends on a variety of interconnected elements, including operating standards and procedures, runway configuration, meteorological conditions, and estimated air traffic demand mix (Kicinger et al., 2016). Estimating nominal runway capacity will be discussed first in section 3.1 where air traffic demand mix and runway configuration will be included. The effect of meteorological conditions on airport runway capacity will be discussed in section 3.2.

3.1. Nominal Airport Capacity

The runway capacity determines how many aircraft can arrive or depart from a single runway on an airport. This runway capacity is mainly expressed in amount of movements per hour, however this can also be made more precise by looking at the amount of movements per 15 minutes. Mainly for single runway airports, arrivals and departures have to take place on the same runway which means that there should be a trade-off between the amount of arrivals and departures. This trade-off can be visualized in Figure 3.1. On airports with multiple runways, a single runway can be used for arrivals and another for departures which inherently increases the airport capacity.

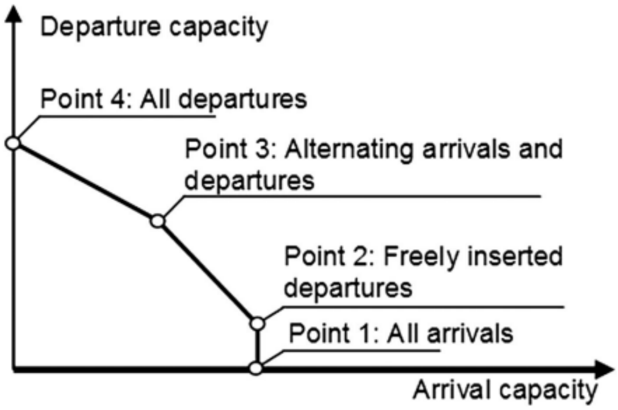

Figure 3.1: Schematic overview of the trade-off between arrivals and departures for a single runway (Kicinger et al., 2016)

The exact number of arrivals and departures that are allowed on a runway per hour depend on a few factors. The arrival capacity is expressed by the airport arrival rate (AAR), and the departure rate is expressed by the airport departure rate (ADR). The effect of weather related factors will be explained in the following section. Delaura et al. (2014) states that elements that are used to calculate the capacity are fleet mix, runway occupancy time and wake vortex separation distance. In nominal conditions the maximum amount of movements on a single runway can be calculated relatively easy using these elements. In case an airport uses multiple runways, the capacity of certain runways are dependent on each other. This holds when there are parallel runways and runways which cross each other (Janic, 2008). The reduction in capacity depends on the distance between the runways in case of parallel runways.

Next to this calculation of the nominal runway capacity, there are also capacity limits to adhere to for an airport. These rules limit runway maximum capacity throughout the day. In The Netherlands, this capacity declaration is set up by ACN (2022). The capacity declaration, as shared by the local airport slot coordinator determines the nominal capacity of an airport on which airline schedules will be based. In Europe, members of the EU are required to appoint an independent entity who is in charge of slot allocation (Santos and Robin, 2010). In Europe, airports can be classified in three categories based on slot coordination.

- Level 1: non-coordination airports
- Level 2: schedule facilitated airports
- Level 3: fully coordinated airports

In non-coordinated airports, supply never exceeds demand, in that way slot coordination is never needed. Schedule facilitated airports are airports where there is a chance that demand will exceed supply in some periods. This congestion is solved by voluntary schedule changes. In schedule facilitated airports, the slots are not actually allocated. In fully coordinated airports, demand exceeds supply more often and the congestion problems cannot be solved by voluntary schedule changes. All arrivals and departures are therefor allocated to a slot. Other papers describe the three levels of coordination as 'uncongested', 'mildly congested' and 'congested'. In 2017, 103 airports in Europe are classified as fully coordinated airports where mostly hub-airports are amongst them (Ribeiro et al., 2019). By slot allocation for those airports, demand is matched to supply, so in normal operating conditions congestion will not occur. This changes when runway capacity is reduced and this has the greatest impacted on busy and congested airports where demand matches supply throughout the day.

3.2. Meteorological Effect on Runway Capacity

As explained above, in normal conditions, arrival and departure capacities of an airport are as designed, and airlines are able to operate according to their schedule. In case of changing weather conditions, the capacity of the runway may decrease due to heavy winds or reduced visibility.

Traffic Management Initiatives (TMIs) can be initiated to control the flow of air traffic. If a TMI is used effectively, it will minimize delay caused by external conditions like weather. The purpose of a TMI is to balance the traffic demand with the capacity of the system, in this case the airport (Chuang, 2017). Chuang (2017) describes four different types of TMIs, of which for all of them the key input is the capacity of the airport. This proves that having an accurate estimation of the airport capacity is extremely important to get the most out of the TMI and thus reduce the delays as much as possible.

However, it is challenging to predict airport runway capacity accurately since it depends on a variety of interconnected elements, including runway configuration, meteorological conditions, expected air traffic demand mix and operational procedures (Neufville and Odoni, 2003). Despite this, there have been many studies on estimating runway capacity under uncertainty of other elements, mostly weather. Several techniques have been used to make an estimation, these will be discussed below.

William (1981) presents one of the first models to calculate the airport capacity. Nominal airport capacity is calculated considering aircraft mix and the required distance between aircraft as well as the runway occupancy time. The model can be amended for 15 different types of runway configurations as a runway system with multiple runways may introduce separation issues. Weather conditions are included in a deterministic matter. Three different conditions can be distinguished by the model, these different conditions affect the separation requirements and Air Traffic Management (ATC) procedures and thus have an effect on the capacity (William, 1981). These weather conditions are Visual Meteorological Conditions (VMC), Marginal Meteorological Conditions (MMC) and Instrument Meteorological Conditions (IMC).

Stamatopoulos et al. (2004) also proposed a model that used deterministic weather conditions. The model proposed is a bit more complex however it is build on the same theory as the model from William (1981). A model is provided which can calculate the nominal capacity of several points on the runway capacity envelope as shown in Figure 3.1. Complexity is added in the integration of more elements of the airfield. In this model, the capacity of the runway, taxiways and apron are integrated. According to the author, until the moment of writing these elements have only been addressed individually and it will be attempted to integrate these elements in a single tool. In order to keep the tool fast, the different elements are integrated in a macroscopic way. Incorporating weather into the model is done in the same way as done by William (1981), however instead of using meteorological conditions, the weather categories are defined as flight rules conditions. The different conditions used are; Visual Flight Rules (VFR) Conditions, Instrument Flight Rules (IFR) Conditions and Low Instrument Flight Rules (LIFR) Conditions.

Amongst William (1981) and Stamatopoulos et al. (2004), there are more papers published that use deterministic weather information. Using deterministic weather data results in predictions that do not take into account the uncertainty of the weather forecast. According to Kicinger et al. (2014), weather uncertainty is definitely something to take into account as weather uncertainty can be very large especially with greater look-ahead intervals, already for time horizons of four to eight hours.

As deterministic models started to show predictions that were far off, the Integrated Airport Capacity Model (IACM) was developed to address the gap of using probabilistic weather information by Kicinger et al. (2011). It integrates weather forecasts to produce probabilistic capacity predictions. In general, the components of the IACM can be summarized as shown in Figure 3.2. In most models, the Terminal Capacity Model is excluded and assumed that there is enough capacity at the terminal. Later, the model was extended by Kicinger et al. (2012) by replacing the weather forecast with ensemble forecasts. This is a weather prediction type that generates a sample of future states of the atmosphere (Kicinger et al., 2012).

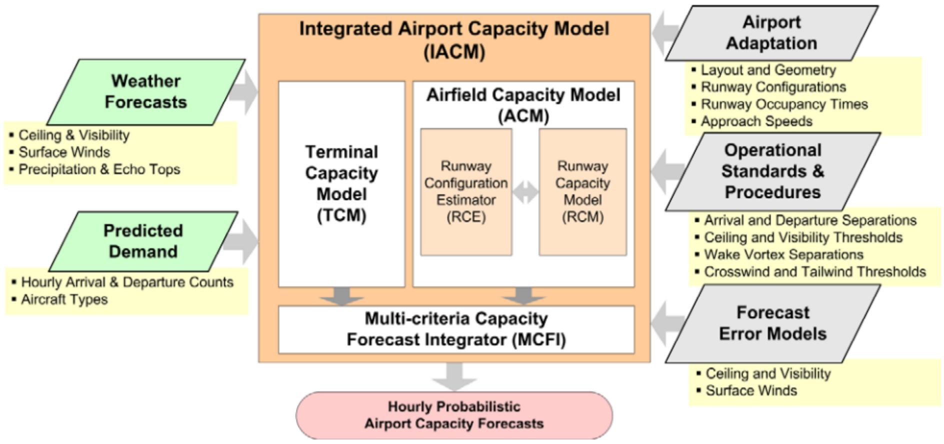

Figure 3.2: Components of the Integrated Airport Capacity Model (IACM) (Kicinger et al., 2014)

Although the process of estimating the runway capacity has been generalized by the IACM, there are still many different methods for estimating the different components of the model such as the Runway Configuration Estimator (RCE) or the Runway Capacity Model (RCM). Below a few models will be discussed which use different methods to, in the end, end up with a runway capacity estimation. They all make use of the principles of the IACM.

Hesselink et al. (2014) shows that it is possible to create a high quality runway capacity forecast based on probabilistic inputs with a lead time of up to two days. As explained before, runway configuration is also an important factor for airport capacity. In case of a single runway use, runway capacity will not be affected by the configuration. However, with the use of two or more runways capacities may be reduced as enough separation is required. This mostly holds for parallel runways. For this reason it is also important to estimate which runway configuration will be used. The model by Hesselink et al. (2014) first forecasts the runway configuration and uses the predicted configuration to estimate the runway capacity.

All available runways are determined based on the wind direction, wind speed, and open runways. It will be checked to make sure that the cross- and tailwinds do not exceed the limit. The optimal runway layout is then determined using an overlay with the preferential runway system and data on runway availability. This overlay has been developed by Hesselink and Nibourg (2011). A probabilistic forecast of the runway configuration that will be used is the outcome of their model. Turning the runway configuration into a runway capacity is done by using the nominal capacities of runway configurations.
Kicinger et al. (2016) presents an analytical stochastic model for estimating airport capacity with a look-ahead horizon suited for tactical traffic flow control. The impact of terminal weather and its uncertainty are explicitly integrated into the model, which expands on earlier studies on estimating airport capacity. For the purpose of creating distributions of expected arrival and departure capacity for each runway configuration at an airport, various types of weather forecast inputs, such as deterministic forecasts, deterministic forecasts with forecast error models, and ensemble forecasts, are investigated, in this way the impact of uncertainty is tested. Interestingly, the author does not find any statistical differences in the results when using different forecast uncertainty representations. In other words, the accuracy of the predictions were comparable. Despite these interesting findings, the author did not do other tests to confirm these findings however advised this as a recommendation. Even with several contributions to the literature about runway capacity predictions ((Kicinger et al., 2011), (Kicinger et al., 2012) and (Kicinger et al., 2016)) the author did not confirm his findings.

Leege and Janssen (2016) created a model which makes use of a machine learning technique to provide a 30-hour probabilistic forecast. Using Meteorological Aerodrome Reports (METAR) and data on runway use, the model was trained for predicting runway configurations. The machine learning technique that is used is multinominal logistic regression. The output is the probability that a certain runway combination will be used. The probabilistic runway use was combined with the meteorological probabilistic forecast and with the use of Monte Carlo simulation a prediction for the capacity could be made that accounts for the uncertainties in the runway use and weather forecast.

The model as developed by Kicinger et al. (2016), is referred to as the ’business rule tool’ by Raju et al. (2021). Raju et al. (2021) aims to find out if machine learning could improve the runway capacity predictions comparing to the ’business rule tool’. Multiple machine learning models were developed and compared to the ’business rule tool’ with the objective of finding out if machine learning is more accurate. Secondly, the author want to find which type of machine learning model would yield the best results. Four machine learning models were developed; random forest, logistic regression, catboost and neural network. It was found that only logistic regression yields worse results than the ’business rule tool’ and the models yield significantly better results than the tool. This shows the relevance of using machine learning for runway capacity estimations.

Also in literature, recently airport capacity estimations trend to use machine learning. This might be due to the fact that several researches wanted to fill the gap of using machine learning for airport capacity estimations as the use of machine learning became more popular recently. Examples are papers of Leege and Janssen (2016) and Dhanasekaran (2014). Leege and Janssen (2016) proposed a model that could be used in the OCC of KLM. The machine learning technique used is logistic regression, it is not stated why this type of model is used. Accuracy ranges from 76% when predicting for one hour ahead to 69% when predicting 27 hours in advance. According to the author this accuracy is high enough to support the decision makers in the OCC. The authors lack in testing other machine learning models as learned from Raju et al. (2021) that logistic regression is the least accurate type. Dhanasekaran (2014) was able to test two machine learning models (nearest neighbors and neural networks) to test the accuracy of the predictions. It was found that both models have a satisfactory accuracy however neural networks showed the best accuracy. This is also in line with what was found by Raju et al. (2021).

Conclusion
Concluding, models predicting runway capacity exist already for a long time. (William, 1981) started with a model for a single runway. Later different types of runway configurations were considered as well as the aircraft mix. Weather conditions were included in a deterministic matter, mostly in three scenarios. The Integrated Airport Capacity Model was developed to fill the gap of moving from deterministic weather conditions to probabilistic weather information leading to much better estimations of runway capacity. In recent years

the first models using machine learning to estimate the runway capacity were introduced. Accuracy of these models are reasonable but according to authors high enough to be used in operational environments. As said, this area has been very popular in research, therefor the gaps are very limited. Many different types of models have been tested, with the latest addition of machine learning. Improvement could only be made to the accuracy of the models by coming up with a new solution methods.

4

Ground Delay Programs

The first efforts to describe the ground delay problem dates back to 1987 when Odoni (1987) released a paper on the flow management problem in Air Traffic Control. During that period, the aviation sector experienced a significant growth and the amount of air traffic increase rapidly while the airport capacity remained mainly the same. As expected this resulted in congestion issues at airport and it was estimated by Odoni (1987) that this congestion costs several billions of US dollars each year. This congestion already occurred with maximum airport capacity. However, during days with bad weather, it is not unusual that the capacity drops to significant lower numbers which results in severe congestion, which could also propagate to other airports.

Odoni (1987) proposed solutions which vary based on the planning horizon. On the long term, new airports and runways could be build to increase capacity. On the mid-term, he proposed to introduce regulations that would spread traffic over less congested airport. His short-term solution focuses on minimizing the delay costs due to the congestion by controlling the aircraft flow. These delay costs are seen as unavoidable. The short-term horizon can vary from minutes to a whole day where the flow of aircraft will be matched to the available capacity.

Within the short term approach, actions to be taken are *tactical* or *strategic*, relating to the planning periods as explained in subsection 2.2.2. Tactical actions can be done when the aircraft is already airborne. These actions include for example, high altitude holding, modifying en-route flights plans, control en-route speed or sequencing aircraft to maximize aircraft arrival rate (Terrab, 1990). Although these actions can help to control the flow of traffic, their effect is limited since the aircrafts are already airborne and scheduled to land at the congested airport. Strategic actions have a much greater potential for regulating traffic (Odoni, 1987) as it includes the modification of flight plans before take-off or delaying the departure of flights. Delaying flights on their departure airport has the biggest impact on controlling the flow of traffic. Furthermore, delaying aircraft on the ground is proven to be less expensive than delaying flights that are already airborne, as this will impose extra fuel costs as well as maintenance costs (Richetta and Odoni, 1994).

Odoni (1987) created the so-called generic Flow Management Problem which focuses on strategic flow actions and it is assumed that only the airports have capacity restrictions so no en-route capacity issues. The solution deals with a trade-off between ground-holds and airborne delays. This is also where the name of the problem comes from, Ground Delay Program (GDP). In papers Ground Delay Program and Ground Holding Problem (GHP) are used interchangeably. The main goal of these problems is to solve the trade-off between ground-holds and airborne delays in case of runway capacity reduction.

According to Terrab (1990) there are three different ways to solve the ground holding problem with increasing complexity. First, the problem can be solved deterministic. Deterministic models do not take into account randomness and uncertainty. Specifically for the GHP, this means that the capacities of an airport are known and fixed. The problem can also become static-stochastic. In stochastic models, randomness and uncertainty are taken into account. In static-stochastic models, the capacities are random variables with a probability distribution that represents a forecast for these capacities. Since the model is static, probability distribution is known before the departure of any of the flights. The problem becomes even more complex when a dynamic-stochastic model is used. Also in this model the capacities are random variables, however the probability distribution of those capacities change over time so this probability distribution can change after the first departure of any set of flights. This makes the model dynamic.

In this chapter, the three different models will be discussed. First the formulation of deterministic models will be discussed in section 4.1. After that, uncertainty will be introduced and stochastic programming is needed to solve the problem. These models will be described in section 4.2.

4.1. Deterministic Models

One of the first deterministic models were introduced by Odoni (1987) and Terrab (1990). The model is a first approximation to the single airport flow management problem. As this is the first approximation, the model is very basic, however it is able to solve the problem. The objective is to minimize the total ground delay cost. The objective function is given by Equation 4.1.

$$\min \sum_{i=1}^{N} \sum_{j=P_i}^{P+1} C_{ij} x_{ij} \tag{4.1}$$

According to Terrab (1990), the constraint matrix of this integer problem is uni-modular, the integrality constraints can be relaxed and the well-known simplex method can be used to solve this model.

There exist several additions to this model to make it a bit more realistic and thereby complex. One of them is presented by Hoffman and Ball (2000). Hoffman and Ball (2000) added a parameter in the objective function to allow super-linear growth. In this way, the model favors delaying two flights by a moderate amount over delaying one flight a large amount and another by a small amount. According to the paper this makes the assignment of delays more realistic. The objective function is described by Equation 4.2.

$$\min \sum_{f \in F} \sum_{t=1}^{T} C_f (t - a_f)^{\sigma} X_{ft} \tag{4.2}$$

where F is a set of flights, T a set of time periods, C_f a constant of flight f, a_f the scheduled arrival time of flight f and X_{ft} is the decision variable, which is 1 if flight f is assigned to time interval t.

Both of these models are still very basic. They do not account for departing aircraft, and the link between arriving and departing flights with the same aircraft. For example, aircraft 1 comes in with a delay of one hour, but in its original schedule it had to depart 1 hour after its original arrival time for its second flight. This means that there is not enough time for the turnaround of the aircraft, and the second flight needs to be delayed even more. Manyem (2018) presents a deterministic model that is more complex than the previous discussed models. It accounts for the predecessors and successors of each examined flight and makes sure that there is enough time between two flights if they are connected. Exactly what is lacking the previous discussed models. The model presented by Manyem (2018) is has originally been developed by Navazio and Romanin-Jacur (1998), however it has been slightly altered.

Also in deterministic models, the impact of missed connections of passengers was introduced. Using the simple model from Terrab (1990), Soldner and Barnhart (2009) expanded the model to account for passenger delays. Normally passenger delays are quantified by 'passenger delay-minutes' which is the multiplication of the amount of passengers by the delay in minutes. According to Soldner and Barnhart (2009) this metric makes sense for flights were passengers are not connecting. This metric does not account for the extra delay that is caused because a connecting passenger misses their connection. Therefor they introduce a more accurate metric which takes into account how late a passenger arrives at his or her destination, which might not be the destination of the flight. The delay cost of a passenger is determined by the departure time of the next flight. It might be possible that delaying the flight by two periods is free of cost since the passengers can still make the connection. After this, the price of delaying linearly increases. The actual price for a passengers missing its connection is unique for each passenger as this depends on the re-booking possibilities. An example can be seen in Figure 4.1.

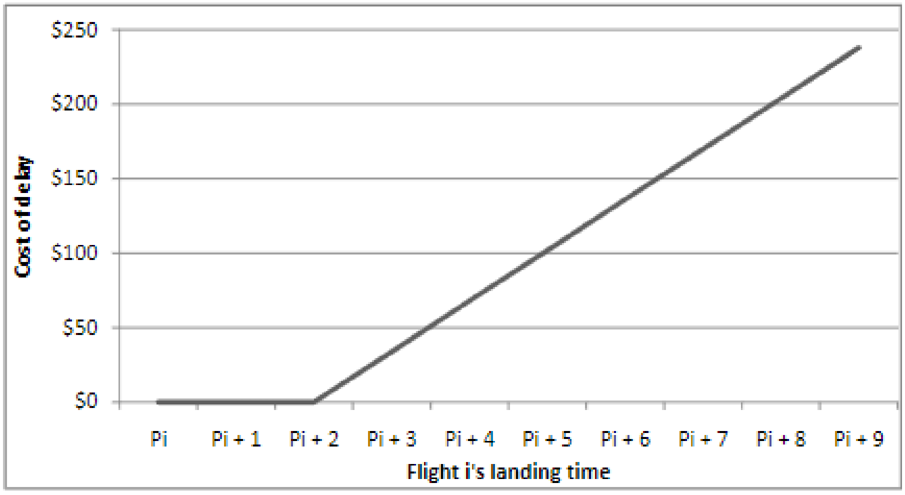

Figure 4.1: Delay cost for passengers who will miss their next flight if the first flight is delayed by more than two periods (Soldner and Barnhart, 2009)

The following model was proposed by Soldner and Barnhart (2009). Since this model builds on the model of Terrab (1990), it uses the same notation as the first model in this section. The objective function is formulated as:

$$\min \sum_{i=1}^{N} \sum_{j=1}^{P+1} C_{ij} x_{ij} + \sum_{i=1}^{N} \sum_{j=1}^{P+1} D_{ij} x_{ij} \tag{4.3}$$

The second term in this objective function is the addition of Soldner and Barnhart (2009). D_{ij} is the passenger delay cost of assigning flight i to time period j. As explained above, it might be possible that this delay cost is zero if a delay does not exceed the minimum connection time of that passenger. This model can also be extended by for example adding a constraint on the maximum delay that is allowed.

These examples of deterministic models can be used to solve the problem and can be expanded to account for crew as well. However, as explained, these models do not take into account randomness and uncertainty. As forecasts always have an uncertainty factor, this also has to be taken into account in the model. This can be done by creating stochastic models. Stochastic modelling has also been applied to solving GDP problems.

4.2. Stochastic Programming

Stochastic programming, commonly referred to as stochastic optimization, is a mathematical framework that is used to model decision-making under uncertainty (Li and Grossmann, 2021). The first publication of stochastic programming, "Linear Programming under Uncertainty" was published by Dantzig (1955), who is recognized as the father of the simplex algorithm for linear programming. One of the reasons for developing the stochastic programming modeling framework, according to Dantzig (1955), was "to include the case of uncertain demands for the problem of optimal allocation of a carrier fleet to airline routes to meet an anticipated demand distribution". Beale (1955) is another early study on stochastic programming.

4.2.1. Two-stage Stochastic Programming

After the introduction of stochastic models in literature, several papers have been written solving the ground holding problem (or ground delay problem) with two-stage stochastic modelling. A two-stage optimization problem with uncertainty has two sets of decision variables that are divided according to the conventional two-stage stochastic programming paradigm. Before the uncertain parameters are actually realized, the first-stage variables are chosen. Once the random events have formed, it is then possible to choose, at a cost, the values of the second-stage, or recourse, variables. The second-stage variables are typically seen as corrective measures, or recourse, to solve infeasibilities that might arise due to the uncertainty. The second-stage, however, may also be an operational-level decision problem that arises from a first-stage design and the realization of uncertainty. The second stage cost is a random variable as a result of the uncertainty. The goal is to select the first-stage variables in a way that the total of first-stage costs and the estimated value of randomly generated second-stage costs is as low as possible. The concept of recourse can be applied to linear, integer, and non-linear programming (Sahinidis, 2004).

One of the simplest two-stage models to solve GDP has been created by Hoffman et al. (1999). It is an integer programming model for a single-airport ground holding problem. The model will be shortly discussed below.

The model assumes that the arrival capacity of the airport is the only element of uncertainty. The demand, and also arrival times of those aircraft are exactly known and thereby deterministic. In the model it is assumed that the uncertain capacity is available in a set of scenarios, just as explained before. The arrival capacities for each scenario are known, as well as the probability of the scenario occurring. In order to reduce uncertain airborne delays, the aim of the model is to assign ground delays. As previously explained, ground delays are cheaper than airborne delays for airlines. The objective of the model is to minimize the expected value of the sum of the costs of the ground and airborne delays. The objective function and constraints are:

$$\min \sum_{t=1}^{T} c_g G_t + \sum_{q=1}^{Q} \sum_{t=1}^{T} c_a p_q W_{q,t} \tag{4.4}$$

The objective function (Equation 4.4) consists of two parts. The first part is the first stage (without uncertainty) and the second part is the part where the scenarios are introduced (Q). In the first stage, the goal is to minimize the cost of ground holding aircraft, c_g is the cost of ground holding an aircraft and G_t is the number of flights which are delayed by one or more time steps. Where in the second part, the cost of airborne delays is minimized under each scenario, c_a is the cost of one period of airborne delay. It should be noted that $c_a > c_g$, otherwise there is no need for ground holding as aircraft could just wait in the air as that would be cheaper. p_q is the chance of scenario q to occur and $W_{q,t}$ are the number of flights delayed in the air under each scenario.

Richetta and Odoni (1993) also described the problem as a two-stage stochastic problem in which the notation of the problem is very comparable to the model of Hoffman et al. (1999). The difference is that Richetta and Odoni (1993) divides the aircraft is cost classes according to their maximum take-off weight. In the way a heavier aircraft will be more expensive to delay than a lighter aircraft. Despite these papers being the first publications on stochastic models for ground delay programs, both papers recommend to turn the problem in a dynamic problem to include the dynamic nature of the problem.
Glover and Ball (2013) also describes a two-stage model for the ground delay program. Despite Richetta and Odoni (1993) advising to turn the problem into a dynamic (multi-stage problem), the author decided to keep this a two-stage model. The authors model the weather in a way that there are only two possible states. This collapses the scenario tree significantly and allows the problem to be solved wit a two-stage model (Glover and Ball, 2013). It is clear that the authors choose computational time over modelling the problem perfectly. Regardless of the choice of the two-stage model over a dynamic multi-stage model, computational times are not discussed in the paper.

4.2.2. Dynamic Stochastic Programming

Until now, a few models have been discussed which take this uncertainty into account and minimize the total cost of delay however these models do not use updated information on the development of the airport capacity. For that reason, those models are known as static stochastic models. If updated information is taken into account, such models are known as dynamic stochastic models. The first (partly) dynamic model was formulated by Richetta and Odoni (1994). In this model, delays are assigned to flights when their departure time is approaching, rather than assigning the delays to all the flights at once. Richetta and Odoni (1994) choose for this approach for the reason that the decisions would be made with the most up-to-date data. A probabilistic binary decision tree with a finite number of scenarios is used to depict the uncertainty in airport arrival capacity. Branches of the tree are realized as the day progresses, which yields more accurate information about future capacities. A weakness of the model is the fact that once a delay is assigned to a flight, it is not possible to revise this, even though this would be possible in real life. Richetta and Odoni (1994) tried to minimize the effect of this by assigning the delay to the flight short before the departure time.

A model that overcomes the limitations of the discussed model is the model proposed by Mukherjee and Hansen (2007). They propose the so-called Dynamic Revisable Ground Holding (DRGH) model. As the name says, the model allows for the revision of ground delays for aircraft that have not departed yet with respect to updated information. The availability of information throughout the day is modelled using a scenario tree as depicted in Figure 4.2. In this example, there are four different scenarios depicted and the time of the day is represented on the horizontal axis. At the start of the day, all scenarios are still possible. Then, at time τ_1, the tree branches and either scenario 1 and 2 are realized or scenario 3 and 4. New information becomes available at τ_2 and τ_3 as well.

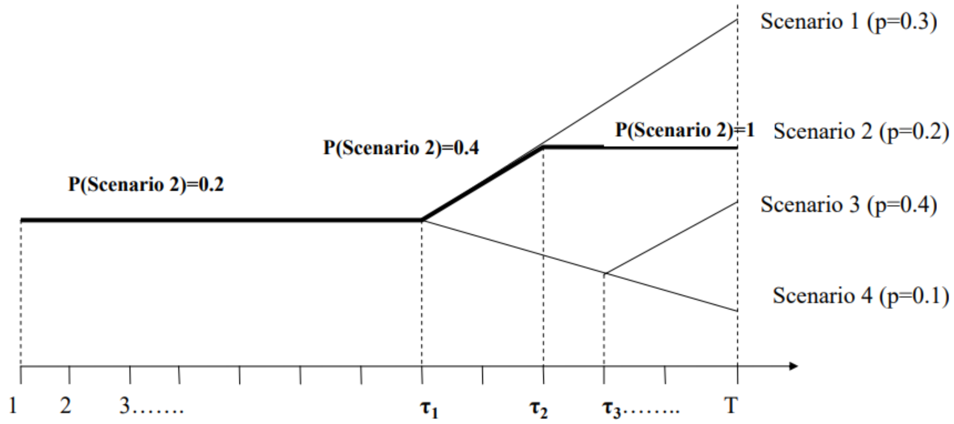

Figure 4.2: A scenario tree showing different scenarios with probabilities over a rolling horizon (Mukherjee and Hansen, 2007)

Also for this model, the objective function minimizes the expected sum of delays, ground and air. The objective function is shown in

$$\text{Min} \sum_{q\in\Theta} P(q)\left[\sum_{f\in\Phi}\sum_{t=Arr_f}^{T+1}(t-Arr_f)X^q_{f,t}+\lambda\sum_{t=1}^{T}W^q_t\right] \tag{4.5}$$

Comparing this objective function to the one of the two-stage model (Equation 4.4) as discussed in subsection 4.2.1, it can be seen that this objective function is a bit more complex. This is mainly due to the fact that this model is dynamics and thus allows for the revision of delays which is not possible in the model of Hoffman et al. (1999). $X^q_{f,t}$ is a decision variable which is equal to 1 if flight f is planned to arrive during time period t under scenario q and zero otherwise. W^q_t is a variable which saves the amount of aircraft which are in the arrival queue at the end of time period t under scenario q. So the first part is related to ground delay and the second part to airborne delay. In the article, Mukherjee and Hansen (2007) prove that their model produces lower expected delay cost compared to the model of Richetta and Odoni (1994). With realistic GDP settings, the expected delay cost of this model are 11% lower than the model of Richetta and Odoni (1994).

4.2.3. Chance Constrained Modelling

An assumption of the previously discussed model is that the exact scenarios are known of which the probabilistic scenario trees can be build. According to Chen and Sun (2018), exactly knowing all scenarios is practically impossible and to add, when the number of scenarios increases significantly, computational challenges may arise. For that reason, an alternate method to include probability in an optimization model is proposed. The idea of this new method is to constrain the chance of a constraint violation. The service level presents the chance that a constraint is violated. One of the major advantages of this method is the ability of a robust solution as this service level is user-defined (Chen and Sun, 2018).

Chen and Sun (2018) chose to modify an existing model and change it into a chance constrained model. The model before modifications is very comparable to the model of Hoffman et al. (1999) as outlined in subsection 4.2.1. The objective function is:

$$\text{Min} \sum_{t=1}^{T} \sum_{i=1}^{M} \left(c_g + Y_t^i \sum_{q=1}^{Q} p_q c_a Z_{t,q}^i \right) \tag{4.6}$$

Comparing Equation 4.6 to Equation 4.4, it can be seen that actually the only difference in the fact that this model considers multiple airports whereas the other model just considers one airport. Chen and Sun (2018) wants to transform this model from a Single Airport Ground Holding Problem (SAGHP) into a Metroplex Ground Delay Problem (MGDP) by considering multi-airports in a metropolitan area.

To get rid of the scenarios, a probabilistic constraint is introduced on the landing capacities, which is one of the constraints in the original model. The constraint is formulated as follows:

$$\mathbb{P}\left(Z_{t-1}^i + X_t^i - Z_t^i \leq \xi_t^i\right) \geq \alpha,\ \forall\, i \in M \,\text{and}\, \forall\, t \in T \tag{4.7}$$

where Z_t^i is the amount of air-delayed aircraft for airport i in time period t and X_t^i the amount of arriving aircraft for airport i in time period t. ξ_t^i are random parameters that represent the stochastic landing capacities at all airports. Lastly, α is the service level, as discussed before. This constraint requires the inner constraint to be satisfied with high probability which is expressed as α. This constraint is modelled as a joint chance constraint as it has a set of inner constraints. Alternatively, the chance constraint could also be expressed as a single inner constraint however joint chance constraints are preferred since they have more modeling power. On the other hand, single chance constraints are easier to work with (Geng and Xie, 2019). In chance constraint optimization, the feasible region of the solution is restricted, such that in the solution, the confidence level is high.

Generally, the chance constrained models are not easy to solve (Chen and Sun, 2018), traditionally the Sample Average Approximation (SAA) approach is used to obtain a solution. SAA will be later discussed in Figure 4.3. Downsides of using SAA is the fact that problems might become intractable with a high number of scenarios, and the fact that the solution found by SAA is not always optimal but sometimes only feasible. For this reason, Chen and Sun (2018) proposes to use a convex approximation method to solve the chance constraint model with the advantage of being able to solve large problems efficiently as well.

Solutions to the chance constraint optimization problem can be found by relaxing the constraints in deterministic constraints and decouple the decision variables and random variables, however, this only works in simple cases. If it is not possible to decouple the decision variables and random variables, the problem is currently impossible to solve (Li et al., 2008).

4.2.4. Robust Optimization

Next to using stochastic programming to deal with uncertainty in optimization, robust optimization is a possibility as well. An important assumption for stochastic programming is the fact that the probability distribution of the uncertain data has to be known. On the other hand, robust optimization does not need the exact probability distributions, but it assumes that the data is in a so-called uncertainty set.

The first publications of robust optimization models date back to the 1970s, when Soyster (1973) published an article on a different way of interpreting the feasible region. Despite this first publication quite some time ago, robust optimization is still a relatively new method as the main developments have been done in the last 20 years where robust optimization has been applied in many fields such as scheduling, engineering, healthcare, supply chain and finance (Gorissen et al., 2015).

With the robust optimization method the optimal solution has two main characteristics. The optimal solution is feasible for any of the uncertain parameters. This is because the optimization problem is solved for the worst-case condition of the uncertain parameters. As explained above, only limited data about the uncertain parameters needs to be known which includes the mean, minimum and maximum. The second characteristic is the fact that the optimal solution for the worst-case is optimal for all other realizations of the uncertain parameters. Again because the worst-case is considered in the decision making (Nazari-Heris and Mohammadi-Ivatloo, 2018).

In literature, there are no papers which solve the ground delay problem with the robust optimization technique. Always solving for the worst-case always causes the cancellation of too much flights which is very expensive. The idea of the model is to find a solution that fits the different scenarios the best.

4.3. Scenario Generation

When a stochastic problem is scenario-based, or when the random vector for the problem has a finite discrete distribution, the problem can usually be solved. For instance, when the underlying random vector is discrete, stochastic linear programming are transformed into large-scale linear problems. The process of creating scenarios (scenario generation) can involve either explicitly modeling the unknown quantities as discrete random variables or discretizing a continuous probability distribution. The number of scenarios in a set determines how much computing power is needed to solve the problem. The main challenge of scenario generation is how to express the uncertainty while limiting the number of scenarios to make the task computationally feasible (Fairbrother et al., 2022).

Before starting the scenario generation, it should be known what information becomes when available. Scenario trees can have multiple stages and the length of time periods can vary. In case of a two-stage scenario tree, the tree starts at a single node and branches from that node into different scenarios with different probabilities, an example of this can be seen in Figure 4.3(a). In case of dynamic models, there are multiple time periods at which uncertain data becomes available and the nodes branch multiple times.

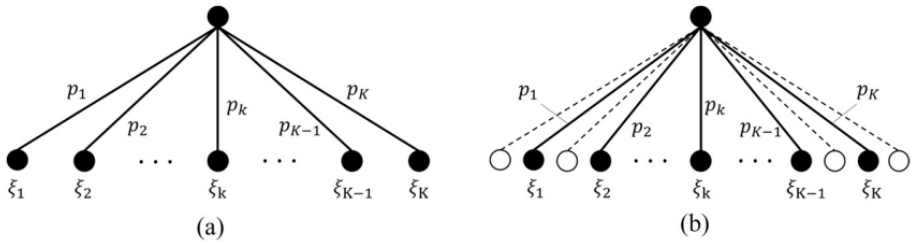

Figure 4.3: Scenario tree for discrete and continuous distribution of uncertain data (Bounitsis et al., 2022)

Solving the Ground Delay Program scenario-based is common in literature. On the other hand, there is not a single way how scenarios are generated. Chen et al. (2008) present an approach using historical data on airport capacity and applying non-parametric method. They use statistical clustering to retrieve patterns of arrival capacity profiles classified by arrival capacity data (Chen et al., 2008). This approach is can only be used if enough historical data is available.

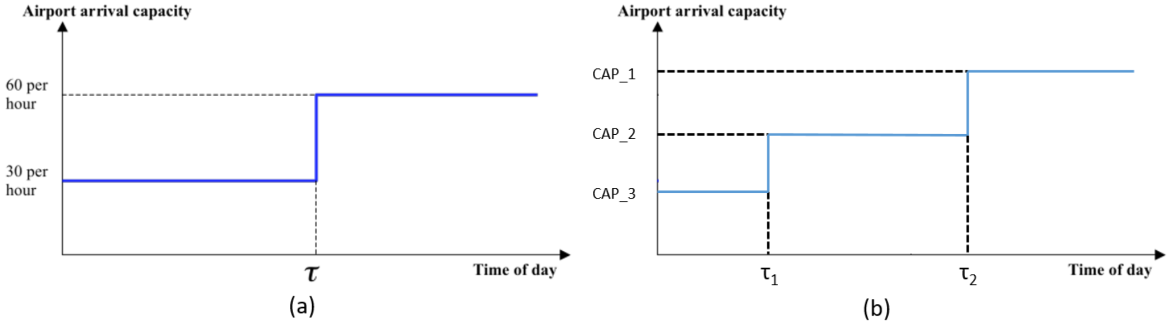

Figure 4.4: (a) Capacity profile under uncertain clearance time τ (Mukherjee et al., 2009), (b) Extended capacity profile under uncertain clearance times τ_1 and τ_2 and uncertain capacities (CAP_1, CAP_2 and CAP_3)

Mukherjee et al. (2009) present a method to generate scenarios based on forecasted data. Uncertain fog clearance times are the reason why different scenarios are examined. The uncertain fog clearance time τ is visualized in Figure 4.4(a). This is a very simple way of scenario generation with only one uncertain parameter. This can be extended into a capacity profile with multiple different runway capacities and multiple times at which the runway capacity increases stepwise. This is a more realistic way at which capacity returns back to normal however due to the extra uncertain parameters, the amount of scenarios might also increase significantly. This stepwise increase of capacity reduction is shown in Figure 4.4(b).

If the random data vector ξ has a lot of components or a lot of possible value, the amount of scenarios will grow to a very large number. Assessing all these scenarios in the optimization problem might cause computation issues (Linderoth et al., 2006). There exist multiple scenario reduction methods, visualized in Figure 4.3(b) where the empty nodes denote scenarios that are not examined due to the scenario reduction approach. A way to reduce the number of scenarios is Sample Average Approximation which will be explained in the next subsection.

Sample Average Approximation

Sample Average Approximation (SAA) is a widely used stochastic optimization technique using Monte Carlo. SAA is frequently used to solve large scale problems. In this method, a sample average estimate produced from a random sample is used to approximate the expected objective function of the stochastic problem. The SAA problem that results from this estimate is solved using deterministic optimization techniques. This process of selecting an estimate from a random sample and solving this problem using deterministic optimization techniques is repeated to get candidate solutions. Alongside the candidate solutions, statistical estimates of the optimality gaps are obtained too (Verweij et al., 2003). Optimality gaps is a way to measure the quality of the solution. In general, the Sample Average Approximation Problem is formulated as follows:

$$z_N = \min_{x \in X} c^T x + \frac{1}{\Psi} \sum_{\psi=1}^{\Psi} Q(x, \xi(\omega^{\psi})] \tag{4.8}$$

The process of obtaining the optimality gap has been described in detail by Kenan et al. (2018). To start off, a sample of size Ψ has to be generated, and the SAA model as described in Equation 4.8 has to be solved to obtain the optimal first stage solution $\hat{z}_{\Psi}^{m}$ and corresponding optimal value $\hat{\omega}_{\Psi}^{m}$. As described before, this process is repeated for M times. The 'true' objective value for the optimal first stage ($\hat{\gamma}_{\Psi'}(\hat{z}_{\Psi}^{m})$) can be estimated by generating a sample of size Ψ' where $\Psi' >> \Psi$. From this, the variance can be calculated. The optimality gap can be computed by $\hat{\omega}_{\Psi}^{m} - \hat{\gamma}_{\Psi'}(\hat{z}_{\Psi}^{m})$. The variance of the optimality gap is estimated by $\sigma^2_{\hat{\omega}_{\Psi}^{m}} + \sigma^2_{\hat{\gamma}_{\Psi'}(\hat{z}_{\Psi}^{m})}$. Formulas for parameters that have not been explained can be found in the report by Kenan et al. (2018). The computed optimality gap can be used as a stopping criteria for the model.

4.4. Solution Methods

This section will give an overview of the solution methods that can be used to solve the model to be created. Generally, the models discussed in this report are solved either using commercial solvers or by using a heuristic method. These methods will be explored in subsection 4.4.1 and subsection 4.4.2 respectively.

4.4.1. Exact Methods

When a model is formulated as a linear programming model, an easy way to solve it is to model it into a commercial solver and let this solver find the optimal solution. Meindl (2012) compares different commercial solvers that can be used to solve (mixed integer) linear optimization problems. The most popular and well-known commercial solvers are; CPLEX, Xpress and Gurobi. CPLEX and Xpress can both be used to tackle mixed integer linear optimization problems whereas Gurobi can also be used to solve other types of problems, for example non-linear problems (Meindl, 2012).

The chance constraint model used to solve the GDP as described in chapter 4 by Chen and Sun (2018) is solved a commercial solver (Gurobi). As described before, the authors modified a formulation of the problem by Hoffman et al. (1999) and implemented chance constraint modelling. Solving the original problem with 500 scenarios only takes 5 seconds. After implementing the chance constraints, the model can be solved in 45 seconds by the commercial solver. According to the author, the computation time is longer because the model needs to find a polynomial approximation on the distribution which can be used for the chance constraints. Increasing the polynomial degree from 5 to 15, the computation time significantly increases to over 200 seconds. This shows that making the problem more complex will significantly increase the computation times when using commercial solvers.

Chang et al. (2016) used CPLEX to solve their two-stage stochastic model for determination of how many aircraft to send to an airspace given a weather forecast and possible airspace capacity reductions. By using a set of scenarios the optimal value for the objective function is found. When considering a set of 1024 scenarios, the problem is solved in 4 seconds. However, when considering a significant amount of extra scenarios (32768) the computation time increases to 3880 seconds. This also shows that when a lot of scenarios have to be evaluated, computation times might increase significantly.

4.4.2. Heuristic Methods

One way to reduce computation time for solving an optimization problem is to use heuristics. Heuristic methods provide sub-optimal solutions against much lower computational times. Next to solving the problem exact, Chang et al. (2016) also implemented a heuristic method and showed that for the case of solving for 32768 scenarios, the solution was found in only 8 seconds with a gap of only 0.02% from the optimal value found in the exact method. Next to this example, others have used heuristic methods to solve problems related to ground delay programs as well to reduce computational times.

Examples are heuristic methods described by Richetta (1995), Luo and Yu (1997) and Brunetta et al. (1998). They all show that solving the problem exactly takes multiple hours and by implementing different heuristics, the problem can be solved in the range of seconds to minutes. Taking a closer look at the heuristic method as described by Brunetta et al. (1998). The heuristic algorithm used is the ABG algorithm and is based on a priority rule. Flights are dynamically grouped together in a number of classes. The priority rules specify which flight has priority over the other for two flights in two different classes. In one of the largest problem solved, the exact method takes over 5000 seconds to find the optimal value. With the heuristic method, an optimal value is found under 4 seconds with only being 2.5% off the actual optimal value. For other tests computation times are even lower and errors smaller. These examples clearly show that if the problem becomes computationally hard to solve, implementing heuristics significantly decreases the computation time by only while still being very close to the actual optimal value.

Conclusion

In this chapter several types of models have been discussed which solve the ground delay problem. It is clear that the GDP is a topic which is already widely addressed in literature and most solution methods have been used to solve the problem, one more successful than the other. Clearly, deterministic models are currently not the way to model problems related to the airline industry as they are lacking the ability to add uncertainty to the problem. Stochastic models are more realistic than deterministic models however they also impose problems, such as the need for scenario generation and high computational time. It is clear that authors do not agree on the fact if models should be extended to multi stage dynamic problem or a two-stage optimization model is good enough. A trade-off shall be made taking into account computational times and accuracy of the problem. The main take-away from this chapter are the different types of models that can be used to solve a similar problem. The GDP problem needs to be extended to be able to fill a gap in literature as the complete problem has been studied extensively.

5

Slot Allocation

Solving the problem like is done in GDP problems is not the only way. It is also possible to solve the problem as a slot allocation problem. Here, in reduced capacity scenarios, a reduced amount of slots is available and flights are assigned to these slots in such a way to minimize costs.

The single-objective formulation of Zografos et al. (2012) was the first publication to address Airport Slot Allocation (ASA) from a mathematical modeling standpoint. The model, was the first attempt to address the modeling complexity of the IATA WSG slot scheduling context. IATA WSG are administrative rules outlined by IATA. Furthermore it minimized total/schedule displacement by taking rolling runway capacity and aircraft turnaround time constraints into account when allocating series-of-slots for an entire scheduling season. By providing two bi-objective formulations that minimize total and maximum displacement and total displacement and the number of violated slot assignments, respectively, Konstantinos et al. (2018) expanded the formulation of Zografos et al. (2012). Utilizing a displacement criterion known as the "maximum acceptable displacement", the number of violated slot assignments was implemented.

Significant recent research has been done on single-airport slot allocation models. First, some research has concentrated on creating optimization models to figure out the right level of the declared capacity to reduce delays, increase airline profitability, and increase passenger welfare. Research in this area have been done by for example Swaroop et al. (2012) and Barnhart et al. (2012). Barnhart et al. (2012) discusses the challenges of managing air transportation demand and capacity due to capacity constraints, resulting in congestion and low schedule reliability, which incur significant costs for airlines and passengers. It suggests various approaches to improve the situation through marginal capacity increases, better management of demand and available capacity, and strategic and tactical measures. The paper also identifies research challenges in specifying, allocating, and utilizing capacity, faced by infrastructure providers and airlines. Both these papers only describe the problems and possible solutions. To quantify their effects on airline schedules and airport on-time performance, optimization model were created for US airports.

Jacquillat and Odoni (2015) presents a schedule optimization approach that selects a set of flights to reschedule in order to reduce demand-capacity mismatches and minimize interference with competitive airline scheduling. The approach uses an iterative solution algorithm that integrates queuing, capacity utilization, and scheduling intervention models. The algorithm converges quickly and has been tested with computational results at JFK Airport, where it achieved substantial reductions in delays through limited changes to airline schedules. The authors conclude that the model is able to significantly reduce arrival and departure delays however steps can be taken to improve the model. Mainly implementing mechanisms that enable airline collaboration to get close to on optimal solution for each individual airline instead of an optimal solution for the airport. This relates to Collaborative Decision Making (CDM).

Another paper, publised by Pyrgiotis and Odoni (2016) around the same time also presents an optimization model to reschedule flights. The paper presents a demand smoothing (DS) optimization model that generates a modified flight schedule at busy US airports in response to scheduling limits. The DS is an integer program. The model produces a feasible modified schedule that obeys the limits, respects aircraft itineraries, and preserves all passenger connections. The paper demonstrates, using an example based on Newark Liberty International Airport, how the DS model can be combined with a network queuing model to evaluate the impact of different scheduling policies on congestion, delays, and passenger convenience. The DS model can be a valuable tool for managing demand and capacity at busy airports. The research's US-centric focus, however, did not encourage consideration of slot series or some IATA regulations.

The paper of Ribeiro et al. (2018) fills the gaps which remain from the previous researches. It extends the literature in two major ways. First the model complies with all criteria in the IATA rule set, which is an extend to the model of Zografos et al. (2012). Second, the model is a major improvement in computational time. According to the authors, the volume of traffic considered in the case study of this research is twice as big as previously covered in literature. In conclusion, this work presents a model-based tool that may assist and optimize schedule coordination decisions based on quantitative objectives, leading to considerable improvements in slot allocation processes in congested airports.

Papers discussed until now had a great focus on the actual slot allocation process and the rules related to this. There are also models which are considered slot allocation models but do not have the rules of IATA implemented. The goal of these models is to assign flights to (time) slots in order to minimize delays. The papers of Jacquillat and Odoni (2015) and Pyrgiotis and Odoni (2016) are examples of these kind of models.

Brunner (2014) published a research in this area. He proposes a model that has to do with the rescheduling of flights during a GDP where the connections of passengers and crew are considered. This model is already quite sophisticated but it does not fully solve the problem needed. As the focus of this model is on the GDP, only arriving flights are considered, therefor the itineraries flown by aircraft are not considered. Despite this, this model can serve as a good foundation for the actual model to be built. According to Brunner (2014) there is currently no published work on the flight rescheduling problem which takes into account crew and passenger connections with the use of a linear programming formulation.

Important to note, this model only includes arriving flights and thereby only a set of arrival slots is incorporated into the model. The arrival delay is determined by the difference between the scheduled arrival time and the allocated slot. The departure delay is then calculated based on the time of the allocated slot, taking into account the minimum turnaround time of the aircraft. The model does not take into account the departure capacity of the airport.

Cheung et al. (2021) proposes a model which does take departing flights into account as well. By slightly adjusting the demand in relation to a dynamic capacity generated from an analytical capacity model, the proposed Mixed Integer Programming (MIP) model levels off peak airport traffic. Although this is not exactly the problem that should be solved, the formulation can still be useful since the model takes into account the departure aircraft and capacity as well. The authors believe that the proposed schedule optimization model is a valuable decision support tool for planners and can complement existing IATA administrative scheduling procedures. This result is significant as the smaller number of flights displaced and smaller displacement from the original schedule can substantially reduce the administrative burden of negotiating schedules at congested airports. This paper, instead of implementing the existing rules of IATA, chose to develop a model that can complement the existing IATA scheduling rules. In literature, this is the only paper found that proposes to complement the IATA rules set.

Conclusion

Research on slot allocation has a much shorter history than GDP problems. There are papers that implement the rule set of IATA for slot allocation as well as papers that describe the need for optimization models that could reschedule flights to reduce delays and fit in the declared capacity. Models that reschedule flights are available in literature, they range from problems that only allow small deviations from the schedule to smoothen the current schedule to problems that can reschedule flights in case of significant reduced runway capacities. Certainly, there is still a gap to fill here to create a model for rescheduling flights that integrate passengers, crew and aircraft recovery.

6

Research Gap and Opportunity

This chapter describes the research gap (section 6.1) that has been identified from this literature study. Following from the research gap, the research question will be presented as a result from the research opportunity in section 6.2

6.1. Research Gap

As became clear from chapter 4, the Ground Delay Program concept has been widely addressed in literature. In case of a capacity reduction at an airport, there exist many models, deterministic, static stochastic and dynamic stochastic, to minimize the impact of the disruption and minimize the associated delay costs. Launching a GDP is done by the Federal Aviation Administration (FAA), and thereby something which is typically done in the US. In Europe adjusted arrival and departure times are provided by Eurocontrol. In GDP models, only adjusted arrival times of aircraft are taken into account, whereas departure capacity can be reduced as well due to weather conditions.

As described by Bouarfa et al. (2014), the current practise in Operations Control Centers (OCC) is to wait for information from ATC about the affected flights in the GDP. Then, the OCC evaluates the impact of new arrival or departure times and reactive actions will be taken. Reactive actions can be swapping an aircraft, use a reserve aircraft, swapping crew and using reserve crew, to name a few. At last, the option of flight cancellation is considered as this is a very expensive option.

In literature, the models designed to recover from disruptions are mainly focused on tactical planning. This means that, once a disruption occurred, solutions will be found and decisions will be made. This is a reactive approach to disruption management. Tactical planning is done on the day-of-operations. Lacking in literature is research on proactive approaches for specific expected disruptions. One of the proactive approaches which is implemented widely in the airline industry already is robustness of schedules, this is an approach for disruption management, however on the long-term when specific details of the disruption are not yet known. This is known as strategic planning. Specifically lacking in literature are proactive approaches to disruptions in-between long term strategic planning and tactical planning, emphasized by the red area in Figure 6.1. The focus will be to create a model which can be used a day before operations but after the fleet and crew assignment. As such, it shares similarities with the robust airline planning literature, but it deals with a tactical disruption recovery problemas opposed to a strategic planning problem.

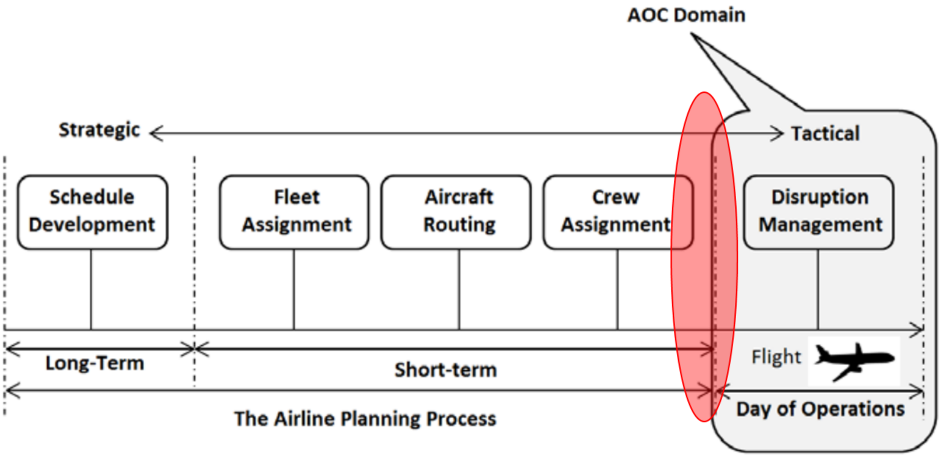

Figure 6.1: Strategic and tactical planning phases in the airline industry (Bouarfa et al., 2014)

Due to the reason that the model will be created for a day before operations, limits the actionable decisions that can be taken at that moment. Simulated delays cannot be assigned yet as the exact outcome of the scenarios considered is not known yet. Only cancellations are considered as actionable decisions, this is also not yet found in literature.

To add, rescheduling flights due to a capacity reduction is either done from an airport perspective or from an airlines perspective however only considering flights of that airline. In this model, rescheduling will be done from an airline's perspective but also considering other flights at the hub-airport of the airline. This introduces uncertainty about the decisions of other airlines as well as gap in the available data.

6.2. Research Opportunity

Based on the research gap presented in the previous section, the goal of research can be formulated in the following research question.

Research Question

How can an uncertain airport runway capacity reduction, at the hub-airport, be translated into actionable decisions for an airline, a day before operations?

As this research question covers the whole problem in once, sub-questions can be formulated to break the problem apart. Sub-questions are:

- How can the uncertain runway capacity be modelled into scenarios a day-before-operations?
- How will the actions of other airlines be included in the decision-making?
- What disruption recovery strategies can be used to model the daily operations realistically?
- In what way can heuristics be implemented in the model to limit computational time?

Bibliography

ACN. Capacity declaration amsterdam airport schiphol: Summer 2022, 2022.

M. Ball, C. Barnhart, G. Nemhauser, and A. R. Odoni. Chapter 1 air transportation: Irregular operations and control, 2007. ISSN 09270507.

C. Barnhart, D. Fearing, A. Odoni, and V. Vaze. Demand and capacity management in air transportation. *EURO Journal on Transportation and Logistics*, 1(1-2):135–155, 2012.

E. M. L. Beale. On minimizing a convex function subject to linear inequalities. *Journal of the Royal Statistical Society: Series B (Methodological)*, 17:173–184, 1955. doi: 10.1111/j.2517-6161.1955.tb00191.x.

S. Bouarfa, H. Blom, R. Curran, and K. V. Hendriks. A study into modeling coordination in disruption management by airline operations control. *AIAA AVIATION 2014 -14th AIAA Aviation Technology, Integration, and Operations Conference*, 2014. doi: 10.2514/6.2014-3146.

G. L. Bounitsis, L. G. Papageorgiou, and V. M. Charitopoulos. Data-driven scenario generation for two-stage stochastic programming. *Chemical Engineering Research and Design*, 187:206–224, 11 2022. ISSN 02638762. doi: 10.1016/j.cherd.2022.08.014.

S. Bratu and C. Barnhart. Flight operations recovery: New approaches considering passenger recovery. *Journal of Scheduling*, 9:279–298, 6 2006. ISSN 10946136. doi: 10.1007/s10951-006-6781-0.

L. Brunetta, G. Guastalla, and L. Navazio. Solving the multi-airport ground holding problem. *Annals of operations research*, 81:271–288, 1998.

J. O. Brunner. Rescheduling of flights during ground delay programs with consideration of passenger and crew connections. *Transportation Research Part E: Logistics and Transportation Review*, 72:236–252, 12 2014. ISSN 13665545. doi: 10.1016/j.tre.2014.10.004.

L. Cadarsoa and A. Marín. Integrated robust airline schedule development. *Procedia - Social and Behavioral Sciences*, 20:1041–1050, 2011. ISSN 1877-0428. doi: https://doi.org/10.1016/j.sbspro.2011.08.113. URL `https://www.sciencedirect.com/science/article/pii/S1877042811014935`.

Y. Heng Chang, S. Solak, J. P. B. Clarke, and E. L. Johnson. Models for single-sector stochastic air traffic flow management under reduced airspace capacity. *Journal of the Operational Research Society*, 67:54–67, 1 2016. ISSN 14769360. doi: 10.1057/jors.2015.53.

J. Chen and D. Sun. Stochastic ground-delay-program planning in a metroplex. *Journal of Guidance, Control, and Dynamics*, 41:231–239, 2018. ISSN 15333884. doi: 10.2514/1.G002964.

P Chen, B. Liu, M. Hansen, and A. Mukherjee. Scenario-based air traffic flow management: From theory to practice. *Transportation Research Part B: Methodological*, 42:685–702, 2008. ISSN 01912615. doi: 10.1016/j.trb.2008.01.002.

W. L. Cheung, R. Piplani, S. Alam, and L. Bernard-Peyre. Dynamic capacity and variable runway configurations in airport slot allocation. *Computers and Industrial Engineering*, 159, 9 2021. ISSN 03608352. doi: 10.1016/j.cie.2021.107480.

C. Y. Chuang. An analysis of the relationship between traffic management initiatives, 2017. URL `https://drum.lib.umd.edu/bitstream/handle/1903/19514/Chuang_umd_0117N_18028.pdf?sequence=1&isAllowed=y`.

J. Clausen, A. Larsen, J. Larsen, and N. J. Rezanova. Disruption management in the airline industryconcepts, models and methods. *Computers Operations Research*, 37(5):809–821, 2010. ISSN 0305-0548. doi: https://doi.org/10.1016/j.cor.2009.03.027. URL `https://www.sciencedirect.com/science/article/pii/S0305054809000914`. Disruption Management.

A. Cook. European airline delay cost reference values - updated and extended values (version 4.1), 2015.

G. B. Dantzig. Linear programming under uncertainty. *Management Science*, pages 197–206, 1955. doi: 10.1287/mnsc.1.3-4.197.

R. A. Delaura, R. F. Ferris, F. M. Robasky, S. W. Troxel, and N. K. Underhill. Initial assessment of wind forecasts for airport acceptance rate (aar) and ground delay program (gdp) planning, 2014.

D. Dhanasekaran. Improved prediction of runway usage for noise forecast development and comparison of runway usage models to improve accuracy of runway usage prediction for noise forecast, 2014. URL `http://repository.tudelft.nl/.`

Eurocontrol. Coda digest all-causes delay and cancellations to air transport in europe, 2019.

J. Fairbrother, A. Turner, and S. W. Wallace. Problem-driven scenario generation: an analytical approach for stochastic programs with tail risk measure. *Mathematical Programming*, 191:141–182, 1 2022. ISSN 14364646. doi: 10.1007/s10107-019-01451-7.

X. Geng and L. Xie. Data-driven decision making in power systems with probabilistic guarantees: Theory and applications of chance-constrained optimization. *Annual Reviews in Control*, 47:341–363, 1 2019. ISSN 13675788. doi: 10.1016/j.arcontrol.2019.05.005.

N. Ghaemi, O. Cats, and R. M. P. Goverde. Railway disruption management challenges and possible solution directions. *Public Transport*, 9:343–364, 7 2017. ISSN 16137159. doi: 10.1007/s12469-017-0157-z.

C. N. Glover and M. O. Ball. Stochastic optimization models for ground delay program planning with equity-efficiency tradeoffs. *Transportation Research Part C: Emerging Technologies*, 33:196–202, 2013. ISSN 0968-090X. doi: https://doi.org/10.1016/j.trc.2011.11.013. URL `https://www.sciencedirect.com/science/article/pii/S0968090X11001653`.

B. L. Gorissen, I. Yanikolu, and D. den Hertog. A practical guide to robust optimization. *Omega (United Kingdom)*, 53:124–137, 6 2015. ISSN 03050483. doi: 10.1016/j.omega.2014.12.006.

L.K. Hassan, B.F. Santos, and J. Vink. Airline disruption management: A literature review and practical challenges, 3 2021. ISSN 03050548.

H. Hesselink, J. Nibourg, and P. Lezaud. Airport capacity forecast : Short-term forecasting of runway capacity, 2014. URL `https://hal.archives-ouvertes.fr/hal-01330101`.

H. H. Hesselink and J. M. Nibourg. Probabilistic 2-day forecast of runway use, 2011. URL `www.nlr.nl`.

R. Hoffman and M. O. Ball. Comparison of formulations for the single-airport ground-holding problem with banking constraints. *Operations Research*, 48:578–590, 2000. ISSN 0030364X. doi: 10.1287/opre.48.4.578.12417.

R. Hoffman, M. Ball, A. Odoni, and R. Rifkin. The static stochastic ground holding problem with aggregate demands, 1999. URL `https://www.researchgate.net/publication/2619039`.

IATA. Air passenger market analysis, 2019.

A. Jacquillat and A. R. Odoni. An integrated scheduling and operations approach to airport congestion mitigation. *Operations Research*, 63(6):1390–1410, 2015.

M. Janic. Modelling the capacity of closely-spaced parallel runways using innovative approach procedures. *Transportation Research Part C: Emerging Technologies*, 16:704–730, 2008. ISSN 0968090X. doi: 10.1016/j.trc.2008.01.003.

N. Kenan, A. Jebali, and A. Diabat. An integrated flight scheduling and fleet assignment problem under uncertainty. *Computers and Operations Research*, 100:333–342, 12 2018. ISSN 03050548. doi: 10.1016/j.cor.2017.08.014.

R. Kicinger, C. Cross, T. Myers, J. Krozel, C. Mauro, and D. Kierstead. *Probabilistic airport capacity prediction incorporating the impact of terminal weather.* 2011.

R. Kicinger, J. Krozel, M. Steiner, and J. Pinto. *Airport capacity prediction integrating ensemble weather forecasts*. 2012.

R. Kicinger, J. T. Chen, M. Steiner, and J. Pinto. Probabilistic airport capacity prediction incorporating weather forecast uncertainty, 2014.

R. Kicinger, J. T. Chen, M. Steiner, and J. Pinto. Airport capacity prediction with explicit consideration of weather forecast uncertainty. *Journal of Air Transportation*, 24:18–28, 2016. ISSN 23809450. doi: 10.2514/1.D0017.

Z. Konstantinos, A. Konstantinos, and M. Madas. Minding the gap: Optimizing airport schedule displacement and acceptability. *Transportation Research Part A: Policy and Practice*, 114:203–221, 2018.

J. Lee, L. Marla, and A. Jacquillat. Dynamic disruption management in airline networks under airport operating uncertainty. *Transportation Science*, 54:973–997, 2020a. ISSN 15265447. doi: 10.1287/trsc.2020.0983.

J. Lee, L. Marla, and A. Jacquillat. Dynamic disruption management in airline networks under airport operating uncertainty. *Transportation Science*, 54(4):973–997, 2020b.

A. De Leege and C. Janssen. Probabilistic runway and capacity forecasting using machine learning to support decision making, 2016. URL `https://www.researchgate.net/publication/310639400`.

C. Li and I. E. Grossmann. A review of stochastic programming methods for optimization of process systems under uncertainty. *Frontiers in Chemical Engineering*, 2, 1 2021. doi: 10.3389/fceng.2020.622241.

P. Li, H. Arellano-Garcia, and G. Wozny. Chance constrained programming approach to process optimization under uncertainty. *Computers and Chemical Engineering*, 32:25–45, 1 2008. ISSN 00981354. doi: 10.1016/j.compchemeng.2007.05.009.

J. Linderoth, A. Shapiro, and S. Wright. The empirical behavior of sampling methods for stochastic programming. volume 142, pages 215–241, 2 2006. doi: 10.1007/s10479-006-6169-8.

S. Luo and G. Yu. On the airline schedule perturbation problem caused by the ground delay program. *Transportation Science*, 31(4):298–311, 1997.

P. Manyem. Disruption recovery at airports: Integer programming formulations and polynomial time algorithms. *Discrete Applied Mathematics*, 242:102–117, 6 2018. ISSN 0166218X. doi: 10.1016/j.dam.2017.11.010.

B. Meindl. Essnet project on common tools and harmonized methodologies for sdc in the ess, 2012. URL `http://neon.vb.cbs.nl/casc/ESSNet2index.htm`.

A. Mukherjee and M. Hansen. A dynamic stochastic model for the single airport ground holding problem. *Transportation Science*, 41(4):444–456, 2007.

A. Mukherjee, M. Hansen, and S. Grabbe. Ground delay program planning under uncertainty in airport capacity, 2009.

L. Navazio and G. Romanin-Jacur. The multiple connections multi airport ground holding problem: models and algorithms. *Transp. Sci. 32*, pages 268–276, 1998.

M. Nazari-Heris and B. Mohammadi-Ivatloo. *Application of robust optimization method to power system problems*. Elsevier Inc., 6 2018. ISBN 9780128124420. doi: 10.1016/B978-0-12-812441-3.00002-1.

R. De Neufville and A. R. Odoni. *Airport systems : planning design, and management*. McGraw-Hill, 2003. ISBN 9780071770583.

A. R. Odoni. The flow management problem in air traffic control. In Amedeo R. Odoni, Lucio Bianco, and Giorgio Szegö, editors, *Flow Control of Congested Networks*, pages 269–288, Berlin, Heidelberg, 1987. Springer Berlin Heidelberg. ISBN 978-3-642-86726-2.

N. Pyrgiotis and A. R. Odoni. On the impact of scheduling limits: A case study at newark liberty international airport. *Transportation Science*, 50:150–165, 2 2016. ISSN 15265447. doi: 10.1287/trsc.2014.0564.

R. Raju, R. Mital, B. Wilson, K. Shetty, and M. Albert. *Predicting Runway Configurations and Arrival and Departure Rates at Airports: Comparing the Accuracy of Multiple Machine Learning Models*. 2021.

N. A. Ribeiro, A. Jacquillat, A.P. Antunes, A.R. Odoni, and J. Pita. An optimization approach for airport slot allocation under iata guidelines. *Transportation Research Part B: Methodological*, 112:132–156, 2018.

N. Antunes Ribeiro, A. Jacquillat, A. Pais Antunes, and A. R. Odoni. Improving slot allocation at level 3 airports. *Transportation Research Part A: Policy and Practice*, 127:32–54, 2019. ISSN 0965-8564. doi: https://doi.org/10.1016/j.tra.2019.06.014. URL https://www.sciencedirect.com/science/article/pii/S0965856418304804.

O. Richetta. Optimal algorithms and a remarkably efficient heuristic for the ground-holding problem in air traffic control. *Operations Research*, 43(5):758–770, 1995.

O. Richetta and A. R. Odoni. Solving optimally the static ground-holding policy problem in air traffic control. *Transportation science*, 27(3):228–238, 1993.

O. Richetta and A. R. Odoni. Dynamic solution to the ground-holding problem in air traffic control, 1994.

N. V. Sahinidis. Optimization under uncertainty: State-of-the-art and opportunities. volume 28, pages 971–983, 6 2004. doi: 10.1016/j.compchemeng.2003.09.017.

G. Santos and M. Robin. Determinants of delays at european airports. *Transportation Research Part B: Methodological*, 44(3):392–403, 2010. ISSN 0191-2615. doi: https://doi.org/10.1016/j.trb.2009.10.007. URL https://www.sciencedirect.com/science/article/pii/S0191261509001313. Economic Analysis of Airport Congestion.

M. J. Soldner and C. Barnhart. Passenger-centric ground holding: Including connections in ground delay program decisions, 2009.

A. L. Soyster. Technical noteconvex programming with set-inclusive constraints and applications to inexact linear programming. *Operations Research*, 21:1154–1157, 10 1973. ISSN 0030-364X. doi: 10.1287/opre.21.5.1154.

A. Stamatopoulos, G. Zografos, and A. R. Odoni. A decision support system for airport strategic planning. *Transportation Research Part C: Emerging Technologies*, 12:91–117, 2004. ISSN 0968090X. doi: 10.1016/j.trc.2002.10.001.

P. Swaroop, B. Zou, M. O. Ball, and M. Hansen. Do more us airports need slot controls? a welfare based approach to determine slot levels. *Transportation Research Part B: Methodological*, 46(9):1239–1259, 2012.

S. Szabo, D. Liptáková, and I. Vajdová. Robustness as a method of airline pro-active disruption management. *International Review of Aerospace Engineering (IREASE)*, 8(4):151–156, 2015.

D. Teodorovi and S. Guberini. Optimal dispatching strategy on an airline network after a schedule perturbation. *European Journal of Operational Research*, 15(2):178–182, 1984. ISSN 0377-2217. doi: https://doi.org/10.1016/0377-2217(84)90207-8. URL https://www.sciencedirect.com/science/article/pii/0377221784902078.

M. Terrab. Ground holding strategies for air traffic control, 1990.

B. Verweij, S. Ahmed, A. J. Kleywegt, G. Nemhauser, and A. Shapiro. The sample average approximation method applied to stochastic routing problems: A computational study, 2003.

J. Vink, B.F. Santos, W.J.C. Verhagen, I. Medeiros, and R. Filho. Dynamic aircraft recovery problem - an operational decision support framework. *Computers Operations Research*, 117:104892, 2020. ISSN 0305-0548. doi: https://doi.org/10.1016/j.cor.2020.104892. URL https://www.sciencedirect.com/science/article/pii/S0305054820300095.

J. William. Upgraded faa airfield capacity model, 1981. URL https://apps.dtic.mil/sti/pdfs/ADA104154.pdf.

K. Zografos, Y. Salouras, and M. Madas. Dealing with the efficient allocation of scarce resources at congested airports. *Transportation Research Part C: Emerging Technologies*, 21(1):244–256, 2012.