An Approach to Predict Operational Performance of Airline Schedules Using Aircraft Assignment Key Performance Indicators

by

Robin Riedel

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of

Master of Science in Aeronautics and Astronautics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2006

Author .....
Department of Aeronautics and Astronautics
May 22nd, 2006

Certified by .....
John-Paul B. Clarke
Principal Research Scientist
Thesis Supervisor

Accepted by .....
Jaime Peraire
Professor of Aeronautics and Astronautics
Chair, Committee on Graduate Students

An Approach to Predict Operational Performance of Airline Schedules Using Aircraft Assignment Key Performance Indicators

by

Robin Riedel

Submitted to the Department of Aeronautics and Astronautics
on May 22nd, 2006, in partial fulfillment of the
requirements for the degree of
Master of Science in Aeronautics and Astronautics

Abstract

This thesis presents an approach for predicting operational performance of airlines on the basis of flight schedules and aircraft assignments. The methodology uses aggregate measures of properties of aircraft assignments, called Aircraft Assignment Key Performance Indicators (KPIs), and aims to find correlations between them and the operational performance of the airline. A simulation experiment is prepared to gather a large set of data points for analysis. A motivation is given for the use of control theoretic approaches in airline operations to utilize the KPIs as a basis for initial planning and corrective actions.

Thesis Supervisor: John-Paul B. Clarke
Title: Principal Research Scientist

Acknowledgments

There are a number of people who contributed to this work. I would like to express my gratitude to

- my parents, grandparents and siblings for their support over the years I spent at MIT.

- my thesis advisor, John-Paul Clarke, as well as all other participants in the Airline KPI Project, especially Michel Turcotte at Air Canada, Erik Andersson, Fredrik Johansson, Stefan Karisch, and Antonio Viegas Alves at Carmen Systems, Gerrit Klempert, Christoph Klingenberg, Michael Mederer and Brigitte Stolz at Lufthansa, and Terran Melconian, Praveen Pamidimukkala, and Michelangelo Raimondi at MIT, for the productive and pleasant collaboration.

- my friends and colleagues at MIT, especially Louis Breger, Nayden Kambouchev, Georg Theis, and Laura Waller for their continuous support in preparing this thesis.

- Cynthia Bernhart, David Darmofal, Amedeo Odoni, Jaime Peraire, Raul Radovitzky, and David Robertson at MIT for their exceptional teaching and their integrity, that made my time at MIT that much more enjoyable.

- Christian Käufer for his continuous support of my flying career which I could not have continued throughout graduate school otherwise.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

Delays are an expected part of air travel today. Over the last seven years, approximately one fifth (21.5%) of all domestic flights in the United States of America (USA) were either delayed[1] or canceled [22]. The situation in Europe is similar [11]. A study estimates that flight delays in scheduled European air traffic in 1999 cost airlines between EUR3.0 and EUR5.1 billion [9]. With the recent strengthening of passenger rights in the European Union [24, 20], the cost to airlines is likely to increase further.

One possible way to reduce delays during operation is to design robust airline schedules [4]. By rearranging existing and adding additional reserves or buffers into a schedule, delays can be reduced and their propagation can be limited. However, a tradeoff exists between cost efficiency and robustness. While the cost impact of adjustments to improve robustness is easily determined from the schedules and accounting data, the operational improvements are difficult to estimate. As a result, reactive approaches, i.e., airline recovery, are more common than proactive approaches, i.e., robust scheduling. It is relatively easy for the mid-level decision maker to justify costs created by bad weather. It is much harder to justify costs for proactive adjustments that will reduce costs in the case of bad weather in the future. This is especially true

[1] A delayed flight in this statistic is defined as a flight that arrives at the arrival gate 15 minutes or more after the scheduled time.

because it is usually uncertain how well the proactive measures will reduce costs on a bad weather day.

To reduce delays through robust scheduling, an understanding of the effects of possible airline schedule adjustments on the airline operation is required. An airline needs to be able to predict the performance of its schedule in order to make successful adjustments. Such predictions can be based on past performance, if the environment and the parameters of the operation remain similar. In fact, a structured analysis of the relationship between scheduling factors and operational performance should provide insight into the operation and assist with future predictions.

However, airline operations are complex. The airlines within the Lufthansa Corporation[2], for example, operate more than 1500 flights on a typical day transporting 128,000 passengers [3]. Not only does the airline schedule vary, but there are also varying uncontrollable external factors that influence the operation, such as weather conditions. Thus, variations exist and no day of operation is like another. Therefore, it is extremely unlikely that a large set of identical schedules in different weather conditions or different schedules in identical weather conditions would exist in historic data. Such sets are, however, required to conduct structured factor analysis to identify what properties of a schedule make it more robust to the uncertain conditions in which the schedule is operated in.

In this thesis, a simulation platform is developed for creating such a data set, an approach to analyze aircraft assignment measures and their correlation to operational performance is presented, and a set of simulation input data is described. In the proposed approach, the airline operation is simulated using an air network simulation and an airline recovery tool and multiple schedules are examined over a range of different weather scenarios to provide a level of granularity of data points that is unavailable in historic data. Multiple Leading Key Performance Indicators (LKPIs) aggregating information on the aircraft routing are presented and an analysis of their correlation to operational performance measured by output measures is suggested. Also, a view on airline operations, based on control theory, is introduced as a potential

[2]Deutsche Lufthansa, Lufthansa City Line, Lufthansa Cargo, Condor

framework for utilizing the LKPIs examined in the analysis.

1.2 Thesis Outline

This thesis is divided into the following chapters:

In **Chapter 2** an overview of airline processes is provided. This includes the airline scheduling, operation, and recovery processes.

In **Chapter 3** a general definition of the term Key Performance Indicator and definitions and examples of the Leading Key Performance Indicators and output measures that should be examined in the analysis are presented.

In **Chapter 4** an alternate view of airline processes is introduced. This view is based on control theory and is intended to provide a framework for the use of Key Performance Indicators in airline scheduling and recovery.

In **Chapter 5** the simulation platform and inputs are presented. This includes a description of the MIT Extensible Air Network Simulation (MEANS), the Integrated Operations Control System (IOCS), the simulation environment, and detailed information about the benchmark simulation inputs.

In **Chapter 6** the proposed analysis for determining the correlations between the Leading Key Performance Indicators and output measures based on the simulation results is presented.

In **Chapter 7** shortcomings of simulation environment and the proposed analysis are identified and suggestions are made for for further research leading to the application of the approach.

Chapter 2

Overview of Airline Scheduling, Operation, and Recovery

2.1 Airline Scheduling

An airline typically plans its operation ahead of time in an *airline schedule*. Generally, the airline schedule consists of three major parts: a) the flight schedule, b) the aircraft assignment and c) the crew schedule (Figure 2-1).

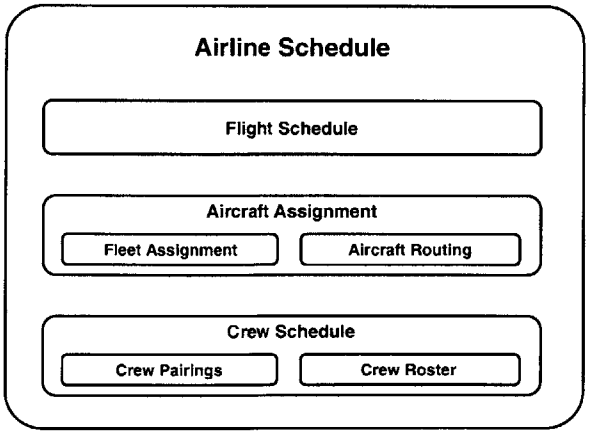

Figure 2-1: Parts of an Airline Schedule

2.1.1 Flight Scheduling

The basic "building block of a flight schedule is the flight leg. A *flight leg* is defined by a departure airport and time and an arrival airport and time. For the departure and arrival times there exist two different time references, shown in Figure 2-2: block time and flight time.

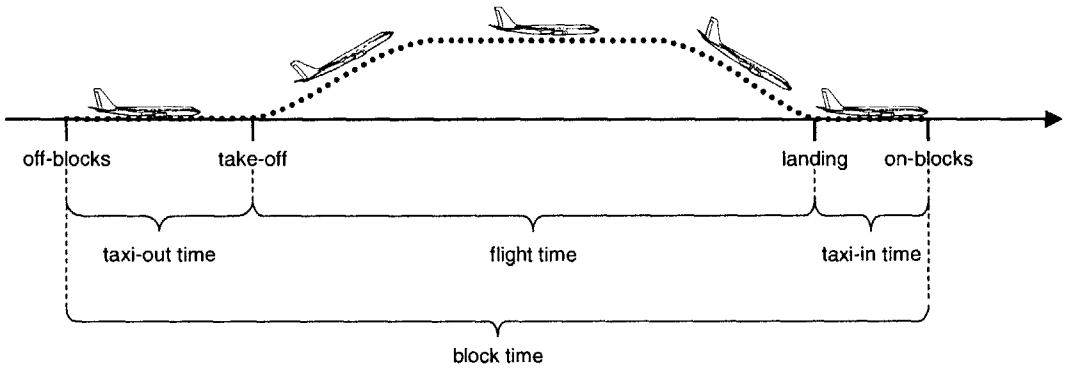

Figure 2-2: Time Terminology

Block time is the time between departure from the gate at the departure airport (*off-blocks*) and arrival at the gate at the arrival (*on-blocks*). The terms off- and on-blocks refer to blocks (or chocks) that are put in front of and behind the aircraft wheels on the ground to stop the aircraft from moving. The exact definition of the beginning and ending of block time varies. Often, off-blocks is defined as the time at which the aircraft begins to move on the ground and on-blocks as the time at which the aircraft comes to a complete stop at the final parking position on the ground. In some cases, off- and on-blocks times are defined in relation to the engine state, where off-blocks will be the time at which the first engine is started (which does not necessarily have to be prior to aircraft movement, e.g., in the case of a push back at the gate) and on-blocks will be the time at which the last engine is shut down. *Flight time* is the time between take-off (at which the aircraft leaves the ground and goes airborne) and landing.

A *flight schedule* is a list of flight legs. The design of a flight schedule is based on numerous factors, among which are market demands, expected revenues and costs, crew constraints, aircraft availability and performance, minimum ground times, legal

constraints, route restrictions, overall airline strategy, competitor behavior, and airline alliances. In the creation of flight schedules, the airline network strategy plays an important role.

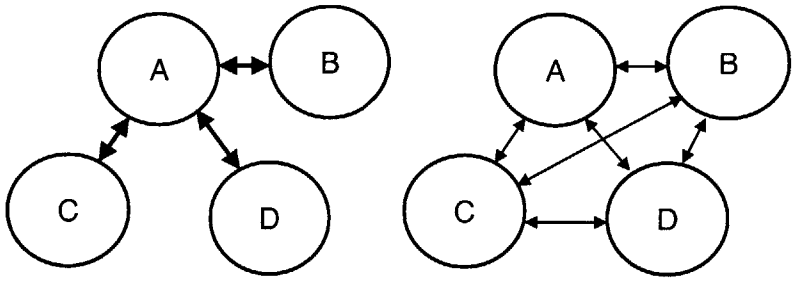

(a) Hub-and-Spoke (b) Point-to-Point

Figure 2-3: Network Strategies

For example, consider a typical hub-and-spoke network as shown in Figure 2-3(a). Airport A serves as a hub and passengers who want to travel from one spoke (airports B, C, or D) to another spoke have to connect through the hub. An alternative network strategy is that of the point-to-point network shown in Figure 2-3(b), where all airports are connected by direct flights. In the hub-and-spoke case the airline operates only three different routes compared to the point-to-point case where it operates six routes. Given the same availability of resources, the flight frequency in the hub-and-spoke network is higher than in the point-to-point network. The thickness of the flight connection arrows between the airports in Figure 2-3 indicates the number of flights per days. Large carriers usually operate a mixed network of both hub-and-spoke and point-to-point flights, where the hub-and-spoke portion is dominant. To provide service between different spokes, the flight schedule has to be created in such a way as to enable passengers to connect between flights. This results in *banks* at the hub in the schedule during which connections between important spokes are available.

For most major scheduled airlines, flight schedules are set and published multiple months before the day of operation. Although major scheduled airlines operate the majority of flight legs daily, some variation in flight schedules between different days exists. These variations could be, for example, driven by passenger demand or aircraft

availability. As a result, the flight schedules for any two days will almost never be identical.

2.1.2 Aircraft Assignment

During *aircraft assignment*, a specific aircraft is assigned to each of the flight legs in the flight schedule. Often, aircraft assignment is broken down into two sequential parts: fleet assignment and aircraft routing. The reason for this separation is that the required computing power is reduced significantly and that the fleet assignment can be created earlier because the required input data for the fleet assignment problem is available earlier than the required input data for the overall aircraft assignment problem. The fleet assignment process only requires aggregate aircraft availability and maintenance requirements which are available much earlier than detailed aircraft availability and maintenance requirements. As a result, fleet assignment can be carried out soon after flight schedule design.

The *fleet assignment* for a flight schedule assigns each flight leg in the schedule to a specific fleet. A *fleet* is a group of aircraft with a common property. Often, that common property is aircraft type (e.g., the Boing 747 fleet or the Airbus A300 fleet), stage length (e.g., long haul fleet and short haul fleet), or crew requirement (e.g., the Boeing 757 and 767 fleet can be operated by the same pilots). Fleets can be broken down into subfleets (e.g., the Airbus A320 series fleet could include the Airbus A320 and the Airbus A321 subfleets).

The modern fleet assignment process for large airlines utilizes mathematical programming methods to optimize an objective function (usually expected profits). This optimization is restricted by a number of constraints: *Cover constraints* require that each flight leg will be assigned to exactly one fleet. *Balance constraints* ensure that aircraft departing an airport have arrived at that airport before and have had at least the minimum ground time on the ground at the airport. *Fleet size constraints* ensure that the number of aircraft required in the assignment does not exceed the size of the available fleet. Additional operational constraints (e.g., minimum ground times for aircraft turnarounds or noise restrictions at airports prohibiting certain fleets to

operate there) may also need to be satisfied. The *minimum ground time* (MinGT) for an aircraft at an airport between two flights is the minimum time required between on-blocks of one flight and off-blocks of the next flights to prepare the aircraft for the next departure (a process also called *turnaround).* During the turnaround the aircraft is unloaded, refueled, serviced (e.g., cleaned and catered) and loaded again. MinGT depends on multiple factors, among which are airport, number of passengers and cargo, aircraft type (e.g., size, break cooling time) and services provided (e.g., catering, cleaning, fuel). The first published modern fleet assignment method was developed by Abara in 1989 [2] and many extensions and reformulations exist today.

During the *aircraft routing* (or aircraft maintenance routing) process, a specific aircraft (or tailnumber) of a given fleet type is assigned to each of the flight legs that was assigned to the fleet during the fleet assignment process. In addition to satisfying cover, balance and fleet size constraints, as in the fleet assignment, the maintenance routing also has to provide maintenance feasibility. Aircraft need to undergo routine maintenance checks prescribed by the airline's maintenance manual, which needs to be approved by the competent authority. These checks are usually time-based (e.g., daily or weekly checks), flight-time-based (e.g., 500 hour check) or cycle-based (5,000 cycle check). While time-based check requirements are known well ahead of time, the due dates for flight-time-based and cycle-based maintenance depend on the usage of the aircraft and, therefore, cannot be scheduled many months ahead. For example, the maintenance system of Lufthansa is shown in Table 2.1.As can be seen, some variability exists in both the intervals and ground times per event. More precise estimates of the values within these ranges become available as the time of the event draws near.

In addition to routine maintenance, aircraft will also need to undergo non-routine maintenance. Non-routine maintenance includes the repair of technical failures and the implementation of Airworthiness Directives issued by the authorities or service bulletins issued by the manufacturer. Non-routine maintenance requirements are usually not known well ahead of time. Due to this limited availability of maintenance information ahead of time, the aircraft maintenance routing is usually done a few

Table 2.1: Example Overview of an Airline's Maintenance System [3]

Event	Interval	Ground Time per Event	Work Hours per Event
Pre-Flight-Check	before every flight	30-60 minutes	1
Ramp-Check	daily	2 - 5 hours	6-35
Service-Check	weekly	2.5 - 5 hours	10-55
A-Check	ca. 350 - 650 flight hours	5 - 10 hours	45-260
B-Check	ca. 5 months	10 - 28 hours	150-700
C-Check	15 - 18 months	36 - 48 hours	650-1,800
IL-Check	5 - 6 years	ca. 2 weeks	up to 25,000
D-Check	5 - 10 years	ca. 4 weeks	up to 60,000

weeks before the day of operation.

2.1.3 Crew Scheduling

The crew schedule for a flight schedule is the assignment of crew members to flight legs. A typical crew for a passenger airline flight consists of a flight-deck (or cockpit) crew and a cabin crew. The flight-deck crew usually consists of a Captain (or Commander, Pilot in Command) and a First Officer (or Copilot, Second in Command). The cabin crew usually consists of a Purser (or Chief Flight Attendant, Chef de Cabin) and one or more Flight Attendants (or Cabin Attendants). Depending on aircraft type and purpose of flight, additional crew members such as a flight engineer, a load master, a translator or long range relief crew may be required.

The modern crew scheduling process consists of two phases: creation of a) crew pairings and b) crew rosters. A *crew pairing* is a list of flight legs that can be operated by a single crew member. A *crew roster* is the assignment of individuals to these pairings. In the creation of both crew pairings and rosters, multiple constraints need to be considered, among which are cover constraints, balance constraints, crew availability constraints, crew qualification and training requirements, legal constraints (e.g., maximum duty time rules), and labor agreements. Crew scheduling is usually carried out three to six weeks before the day of operation.

2.2 Airline Operation

Generally, the goal of airline operation is to execute the airline schedule with minimum costs and maximum passenger satisfaction. Airline operation is cost-driven; very few revenues can be raised within the operations domain of an airline since, by the time of operation, tickets have already been sold.

Execution of an airline schedule is complicated by the fact that assumptions made during the scheduling process are not always valid. There exists a wide range of influences that cannot be controlled, or even accurately predicted, by the airline. These influences include meteorological conditions, other air traffic (e.g., flights by other airlines, the military, and general aviation), Air Traffic Control (ATC) actions, equipment failures, and crew member absenteeism. They have the potential to cause significant differences between the environment for which the airline schedule was designed and the environment in which that schedule must be implemented. This change of environment can cause the airline schedule to become infeasible.

A *disruption* in an airline operation is an event that causes infeasibilities in the future of the airline schedule. For example, consider a snow storm at an airport which will require aircraft to de-ice before the flight. As a result, the time between off-blocks and take-off of the aircraft will be increased. Assuming that the off-blocks departure cannot be made earlier (since passengers will arrive for the scheduled off-blocks departure time) the take-off is also delayed and a disruption occurs. If the aircraft can compensate for the delay en route (by reducing the actual flight time) and land at the scheduled time, the disruption will have been recovered. However, if this is not possible, the resulting arrival delay may cause more infeasibilities in the airline's schedule (e.g., delay propagation to other flights) or in passenger travel plans (e.g., missed connection flights).

Arrival delays often cause schedule infeasibilities through delay propagation to other flights. *Delay propagation to other flights* occurs when an aircraft or crew arrives at a destination behind schedule and is scheduled for another flight which it can no longer serve at the scheduled time. As a result, the next flight of the aircraft

and/or crew will be delayed until the aircraft or crew are ready. This kind of delay is called *propagated delay* or *reactionary delay* [11]. For example, consider the timeline for a specific aircraft shown in Figure 2-4. The top row represents the schedule and the bottom row represents the actual events for the aircraft. The horizontal axis represents time. A disruption in time occurs during taxi out of flight 111 and the take-off of the flight is delayed. The delay propagates through to the next flight of the aircraft, flight 112, since the delay cannot be recovered during flight 111 (by flying or taxing faster) or on the ground between flight 111 and 112 (because the scheduled ground time is equal to MinGT).

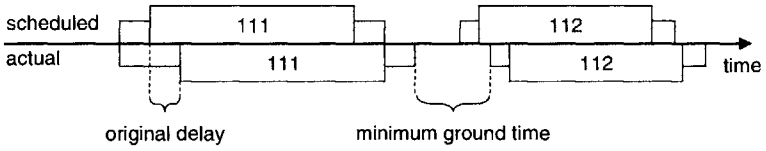

Figure 2-4: Example of Delay Propagation

Frequently, disruptions will span across resource boundaries such that a single disruption can have a negative impact on multiple resources. If, for example, a flight arrives delayed, both the aircraft and the crew are delayed. If the crew is scheduled to fly on other aircraft following the delayed flight, at least two additional flights could be delayed as a result: the next flight of the aircraft and the next flight of the crew (or even more flights if the crew does not stay together).

In the USA, approximately one quarter of all delayed flights (or about 4.5% of all flights operated) are delayed due to delay propagation to other flights (aircraft arriving late), where a delayed flight in this statistic is defined as a flight that arrives 15 or more minutes after its scheduled arrival time [22]. An overview of the causes of flight delays in the USA in October 2005 is given in Figure 2-5.

The distribution of causes for flight delays by delay minutes is shown in Figure 2-6. There are three major causes for flight delays: weather ($5.8\% + 33.9\% * 76.6\% = 31.8\%$), propagation delay (32.5%) and air carrier delay (27.6%).

For every propagated delay there must be an original cause, i.e., the cause for the first delay in a propagation sequence. The portion of propagation delay in Figure

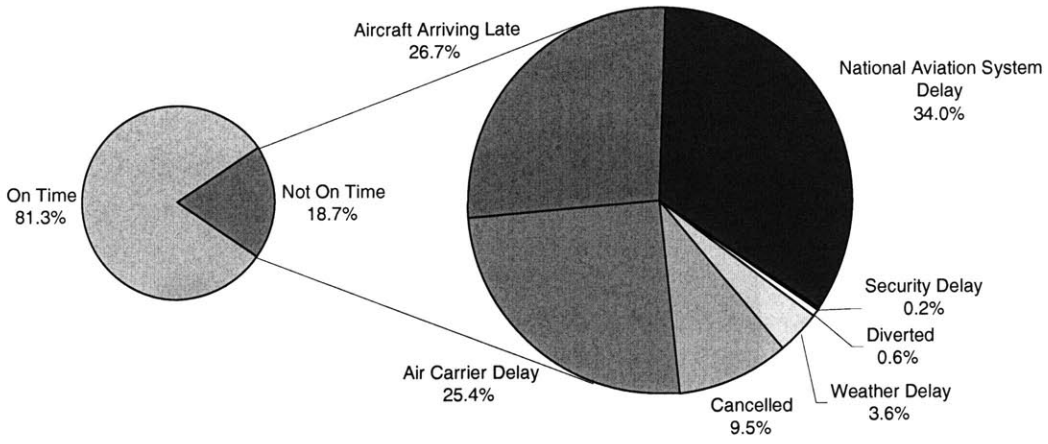

Figure 2-5: On-Time Performance and Delay Causes by Number of Operations in October 2005 (USA only)

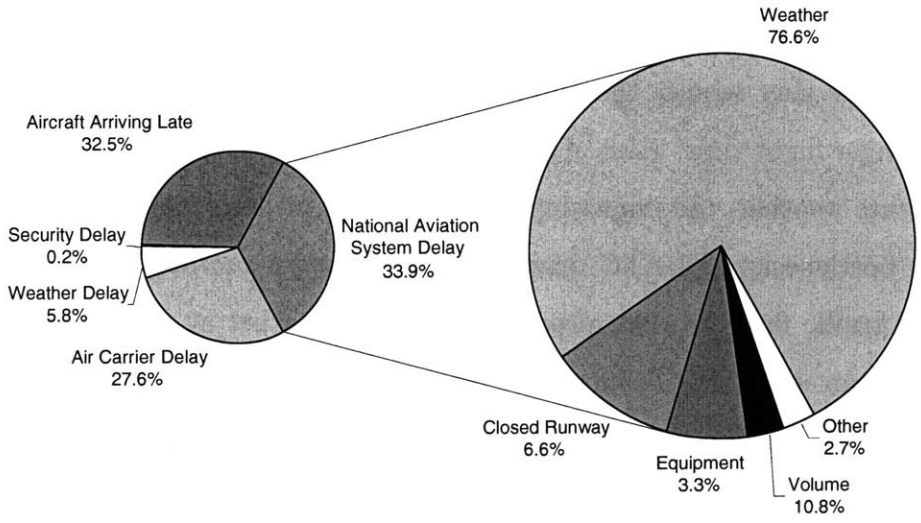

Figure 2-6: Delay Causes by Delay Minutes in October 2005 (USA only)

2-5 can be assigned to these original causes. If it is assumed that the distribution of original causes for the propagation delays is equal to the distribution of causes shown in Figure 2-6, excluding the portion due to aircraft arriving late, a distribution of original delay causes (shown in Figure 2-7) can be derived. Approximately half of

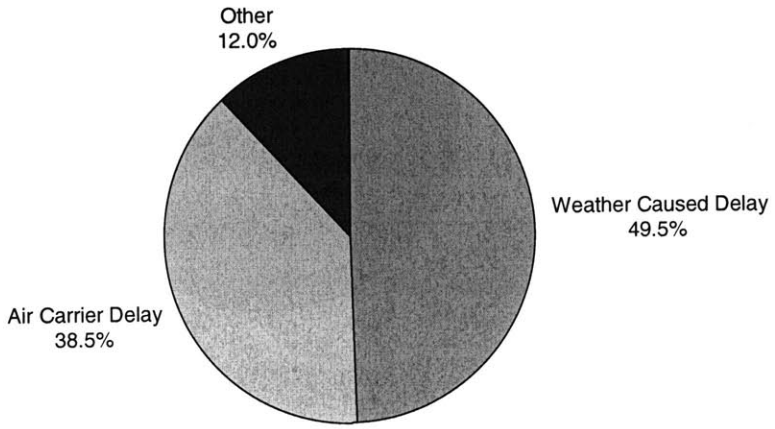

Figure 2-7: Original Delay Causes by Delay Minutes in October 2005 (USA only)

all delays in the USA are originally caused by weather (see Figure 2-7). Within this group of weather causes, there exist two subgroups depending on where the weather occurs. *En route weather* is weather that the aircraft encounters during the flight between airports. Examples of unfavorable en route weather are strong turbulence or hurricanes. They may require the pilot to take a less direct flight path, which may result in a longer flight time. Also, if certain sectors of the airspace become unusable due to en route weather, the remaining sectors will be utilized more which may lead to capacity bottlenecks and ATC driven delays. *Airport weather* affects departing and arriving traffic flows to that airport. When the weather at an airport becomes unfavorable it can either slow down certain segments of the flight process for individual aircraft (e.g., freezing rain that requires aircraft to de-ice) or reduce the capacity of the airport as a whole (e.g., heavy crosswinds that make runways unusable). In the latter case, the departing and arriving flows at the airport are affected as a whole. For example, Boston's Logan airport (BOS) operates on three runways in favorable weather conditions, however can be restricted to the usage of only one runway in the case of strong North-Westerly winds. The reduction of capacity in this case might be

as great as 50%, from approximately 120 operations per hour in favorable weather to as few as 60 operations per hour in poor weather [6].

In addition to the airline schedule, passenger travel plans may also become infeasible due to disruptions. If a passenger travels only on a single delayed flight leg, the passenger will also be delayed. if a passenger travels on an itinerary of multiple flights and needs to connect between flights, a delayed flight could cause the passenger to miss his or her connection and his or her itinerary will not be feasible any more. Missed connections are especially frequent in hub-and-spoke operations, because they rely on passengers to connect between flights.

2.3 Airline Recovery

Airline recovery is the management of disruptions. The goal here is to restore airline schedule feasibility while keeping costs and negative impact on customers low. To achieve this goal, airlines adjust their schedules. Common adjustments are intentional delaying of flight departures, swapping of aircraft, ferrying of aircraft (relocating an aircraft with a non-revenue flight), cancelation of flights, activation of reserve aircraft or crew, and rescheduling of aircraft or crew.

Due to complexity of the airline schedules and constraints on computing power, a general approach is to first attempt recovery within a single resource domain. For example, when a crew disruption occurs, the airline will try to recover this disruption within the crew domain without changing the aircraft routing or flight schedule. If a satisfactory solution cannot be found within a single resource domain, other resources will be included in the recovery. Today, this usually occurs sequentially: A change will be made in one resource domain and then the other resource domains are adjusted to the change sequentially. For example, a change could be made to the flight schedule and then the aircraft schedule and finally the crew schedules will be adjusted to this change.

If passenger travel plans become infeasible due to disruptions, airlines also need to recover their passengers. Common actions by the airlines are rebooking the passenger

onto another flight or itinerary (with the same or a different airline) and compensating the passenger for the disruption of their travel plans (e.g., through restaurant and hotel vouchers).

While recovery personnel and systems today understand a wide range of correlations between different resource domains, state-of-the-art recovery is rarely globally optimal due to the sequential approach. Multiple airlines and decision support software developers are developing a new generation of integrated recovery tools that will find a global optimum over multiple domains, including aircraft, crew and passengers.

Chapter 3

Key Performance Indicators and Output Measures

3.1 General Definitions

The term Key Performance Indicator (KPI) is used in a variety of contexts. It is commonly used within the business administration context [19], where it typically describes "a significant measure used on its own, or in combination with other key performance indicators, to monitor how well a business is achieving its quantifiable objective" [28]. Other areas in which the term KPI is used include product optimization [25], education [10, 21, 32] and public administration [12, 13]. KPIs are frequently used to benchmark a business against its competitors or its own past performance [19].

An indicator can be defined as "that which serves to indicate or give a suggestion of something; an indication of" [26]. An indication can be defined as "a hint, suggestion, or piece of information from which more may be inferred" [26]. Following this definition of an indicator, we define the term *Key Performance Indicator* as a metric used to indicate or estimate one or more critical success factors of an operation, where this operation is the airline operation, defined as the operation of aircraft to transport passengers between airports of their choice at the time of their choice.

Considering the temporal relation between KPIs and performance, there exist

three types of KPIs: lagging, coincident, and leading KPIs. The time difference between KPIs and the performance they indicate is called lead time or lag time. *Lagging* KPIs trail behind the measure. For example, consider the landing time of an aircraft as the performance (Figure 3-1). A lagging indicator for the landing time could be the on-blocks time. Knowledge of the on-blocks time provides information about the landing time, because both are related through the taxi time. However, the on-blocks time lags the landing time and, therefore, has no predictive value.

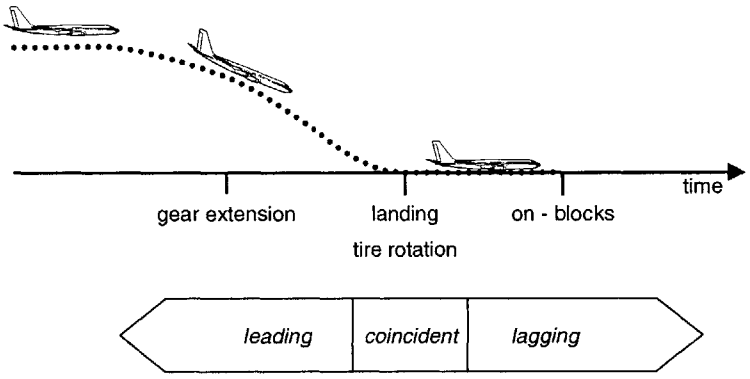

Figure 3-1: Temporal Relation of KPIs and the performance

Coincident KPIs occur concurrently with the performance. In the example above, the time at which the tires of the aircraft rotate for the first time since take-off is a coincident indicator for the landing time; they occur simultaneously, because the tires start rotating when they touch the ground. *Leading* KPIs change before the performance, giving them predictive value. In the example above, the extension of the landing gear is a leading indicator of the landing time, because the landing gear must be extended before landing. Knowing that the landing gear is usually extended a few minutes before the landing, the gear extension time can be used as a leading indicator (i.e., predictor) for the landing time.

3.2 Aircraft Assignment Related Key Performance Indicators

The flight schedules and aircraft assignments have a great influence on the operational performance of the airline (see Chapter 2). Therefore, the properties of the flight schedules and aircraft assignments are expected to have predictive value for the operational performance. The following KPIs, based on flight schedules and aircraft assignment, are suggested for further analysis of their correlation with output measures.

3.2.1 Fleet Assignment Buffer Statistics

The number of buffers within an airline schedule influences the airline's ability to recover from disruptions. A buffer, in the context of fleet assignment, is extra (or slack) time between the arrival and the departure of an aircraft during which the aircraft is idle. A timeline for an aircraft on the ground between two flights is shown in Figure 3-2. When an aircraft arrives, a post-flight service (e.g., disembarkment, unloading) is conducted at the airport. Before the aircraft departs again, a pre-flight service (e.g., embarkment, loading, refueling) is conducted. The sum of the time required for both post- and pre-flight service is the MinGT (see Section 2.2). While the post-flight service can be conducted without any knowledge of the next flight of the aircraft, the pre-flight service requires this information. It must be known where the aircraft will fly next to determine the required fuel and catering to be loaded.

In the actual operation, buffers exist in between post- and pre-flight service. However, an abstraction can be made by shifting the buffer, as shown in the lower part of Figure 3-2. For this abstraction to hold it must be assumed that the next flight of an aircraft is known early enough for the pre-flight service to be performed. Here we assume that this abstraction holds and we define a *buffer* as the time span between the arrival of an aircraft plus the minimum ground time of that aircraft at the airport and the next departure of the aircraft. A buffer can be expressed as $< a, f, t, d >$,

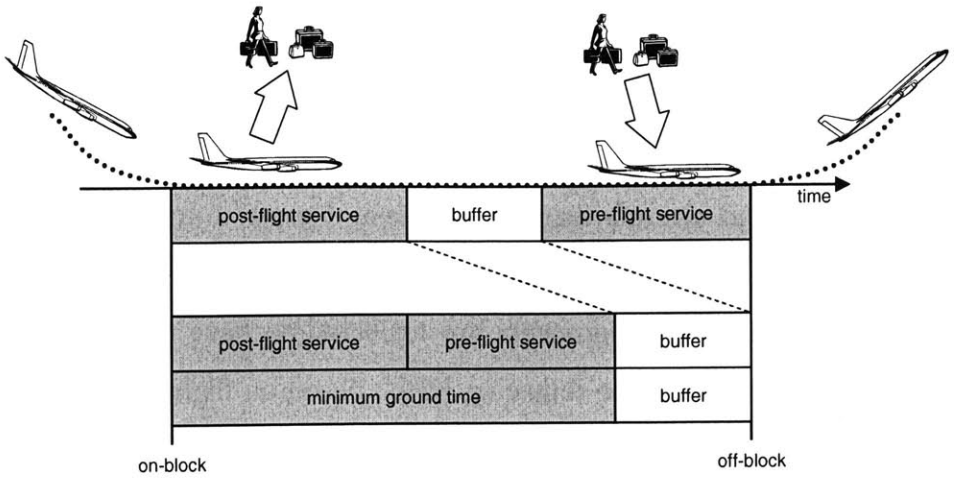

Figure 3-2: Timeline of the Ground Process Between Two Flights

where a is the airport at which the buffer occurs, t is the time at which it begins, d is its duration, and f is the fleet type of the aircraft.

To obtain the set of buffers from an airline schedule, the set of all aircraft ground arcs (i.e., the period when the aircraft is on the ground) is created from the arrival and departure times of the aircraft. In the set of ground arcs, the arrival times are shifted by the MinGT to find the set of buffers. If the aircraft routing for the schedule is unknown, a feasible aircraft routing is created using a feasible fleet assignment and a first-in first-out (FIFO) methodology. In this case, for each fleet, the aircraft arriving at an airport depart from the airport in the same order in which they arrived.

From the set of buffers the average weighted buffer per departure (WBD) can be computed using

$$\text{WBD} = \frac{\sum_{<a,f,t,d>\in B} w(a,f,t,d)}{|L|} \tag{3.1}$$

where B is the set of all buffers, L is the set of all flight legs, and $w(a,f,t,d)$ is a weighting function of a, f, t, and d. In a simple case, the weighting function is the

duration of the buffer, $w(a, f, t, d) = d$. In this case, Equation 3.1 is simplified to

$$\begin{aligned} \text{WBD} &= \frac{\text{total buffer minutes}}{\text{number of aircraft}} = \frac{\sum_{<a,f,t,d>\in B} d}{|L|} \\ &= \frac{\sum_{a\in A}\sum_{r\in R}\left(\sum_{i\in O_{a,r}}(dep_i) - \sum_{i\in I_{a,r}}(arr_i + MinGT_{a,r})\right)}{|L|} \end{aligned} \tag{3.2}$$

where $I_{a,r}$ and $O_{a,r}$ are the sets of all flights of aircraft r that are arriving at and departing from airport a respectively, dep_i and arr_i are the departure and arrival times of flight i respectively, R is the set of all aircraft, and $MinGT_{a,r}$ is the minimum ground time of aircraft r at airport a.

The weighting function may also be any function of any group of the parameters of a buffer, namely a, f, t, and d. Three examples of weighting functions are shown in Figure 3-3. The first example weighting function is a piecewise linear function of d (3-3(a)). It is based on the assumption that the value of an extra minute in the buffer decreases with the overall length of the buffer. This is because the probability that a delay is as long as the buffer decreases as the length of the buffer increases. Thus, an additional minute of buffer time will become less useful as the buffer becomes longer. The second example weighting function is a piecewise linear function of t (3-3(b)). Buffers in the early morning (e.g., before the first flights departs) are less useful than buffers later in the day because in the morning delays will not have occurred yet. Therefore, the weighting function is 0 up to 6am, at which point it rises to a value of 1 until 8am. During the day, buffers are useful to reduce delays and, therefore, the weighting function remains at a value of 1. The later a flight departs, the fewer flights will follow it and, therefore, the fewer flights can be affected by delay propagation. In the extreme case no other flights will follow a flight on that day. Buffers become less valuable as fewer flights can be affected by delay propagation. Therefore, the example weighting function value is again reduced from a value of 1 to a value of 0 between 3pm and 10pm. The third example weighting function is the product of the previous two example weighting functions (3-3(c)). Not shown are examples for

weighting functions that depend on airport or fleet.

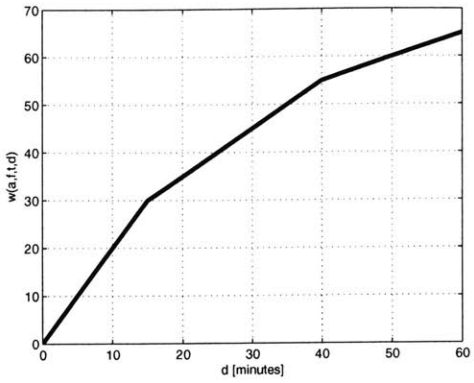

(a) function of duration d

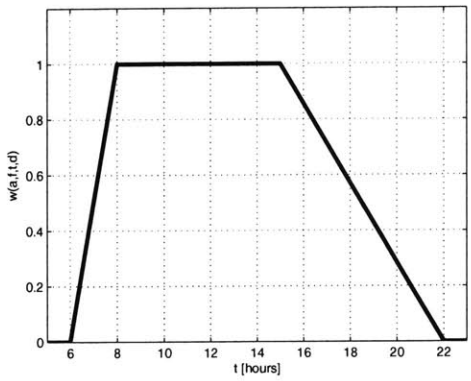

(b) function of time of occurrence t

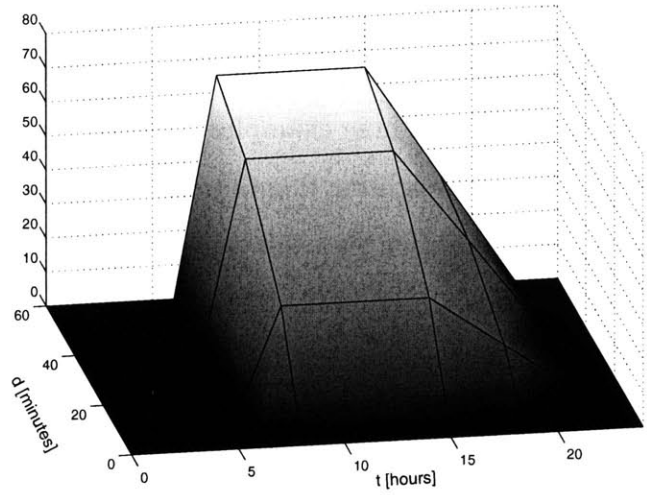

(c) function of both duration d and time of occurrence t

Figure 3-3: Example Buffer Weighting Functions

The following example illustrates the impact of different weight functions on the weighted buffer per departure. Consider the flight schedule for an airport served by two fleet types shown in Figure 3-4. Assuming the MinGTs are 0:40 and 0:50 for the Boeing 737-500 and the Airbus A300, respectively, and FIFO aircraft routing, the set of buffers is

$$B = [< \text{DUS, B, 7:30, 0:30} >, < \text{DUS, A, 8:30, 0:20} >, < \text{DUS, B, 8:50, 0:10} >]$$

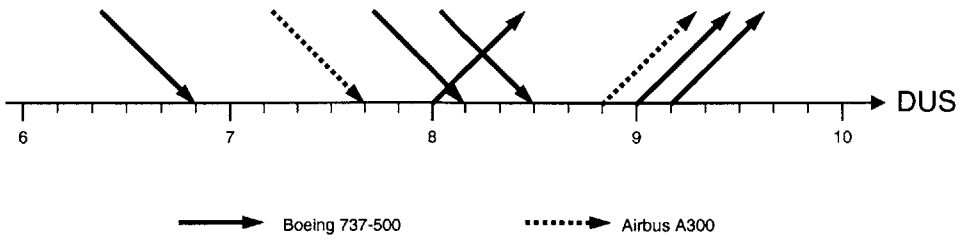

Figure 3-4: Timeline of Aircraft Arrivals and Departures

where A represents the Airbus A300 and B represents the Boeing 737-500. Note that the last arrival does not have any buffer and departs again after the MinGT.

The buffer per departure using $w(a, f, t, d) = d$ according to Equation 3.2 is 0:15 $(0.25(((8:00+9:00+9:10) - (6:50+8:10+8:30+3*0:40))+(8:50 - (7:40+0:50))))$. Using the weighting function shown in Figure 3-3(a) the weighted buffer per departure is 25 $(0.25(w(0:30) + w(0:20) + w(0:10)))$. Using the weighting function shown in Figure 3-3(b) the weighted buffer per departure is 0.6875 $(0.25(w(7:30) + w(8:30) + w(8:50)))$. Using the weighting function shown in Figure 3-3(c) the weighted buffer per departure is 22.1875 $(0.25(w(7:30, 0:30) + w(8:30, 0:20) + w(8:50, 0:10)))$

The choice of the weighting function depends on the airline and its operation. Potential weighting functions can be evaluated with historic and/or simulation data for the correlation of the WBD with the operational performance. Also, parameters of those weighting functions that provide the best correlation between the WBD and the operational performance can be determined by a best fit optimization of the historic and/or simulation data.. Only values of the WBD computed using identical weighting function can be used to compare and benchmark different airline schedules.

3.2.2 Global Fleet Composition Indicators

When aircraft must be swapped between flights as part of a recovery measure, it is usually less complicated to switch two aircraft of the same fleet type than of different fleet types. In this case the capacity, the performance, and the crew requirement of the aircraft servicing each of the flights remains unchanged. If two aircraft of different fleet types are swapped, passengers may be disrupted if, after the switch, the capacity of

the aircraft operating a flight is less than the number of booked passengers. Additional fuel stops may be necessary or other performance based restrictions may apply. Also, wider crew recovery may be required due to the initial crew assignment no longer being valid since the fleet types on some flight legs have changed. Therefore, in a homogenous fleet aircraft swaps are usually less complicated than in a heterogenous fleet.

Global fleet composition indicators (GFCIs) give an overview of the fleet homogeneity of an airline as a whole. A general formulation for the global fleet composition indicator is

$$GFCI = \frac{\sum_{f \in F} (w_f)^2}{\left(\sum_{f \in F} w_f \right)^2} \tag{3.3}$$

where F is the set of fleets and w_f is the weight that fleet f carries. The definition of different fleets within F may be varied (see 2.1.2). The weight w_f for each fleet type can be defined in multiple ways: by number of aircraft, number of total seats, or number of scheduled departures of fleet f.

In Figure 3-5(a) the GFCI with w_f = number of aircraft for an airline with two fleet types is shown. The GFCI value is plotted in the vertical axis and the percentage of aircraft that are of one of the two fleet types is plotted on the horizontal axis (the percentage of the other fleet type is 100% minus the value shown). In Figure 3-5(b) the GFCI with w_f = number of aircraft for an airline with three fleet types is shown. The percentages that two of the three fleets take are shown on the two horizontal axes. From the plots it can be seen that the maximum value of the GFCI is 1. A maximum occurs whenever one fleet holds 100% of the weight. The minimum value of the GFCI in the two and three fleet scenario is 1/2 and 1/3, respectively. Generally, the minimum GFCI value is $1/|F|$ which occurs when all fleets share the weight equally.

The definition of the weight w_f a fleet type carries has a significant effect on the GFCI as the following example shows. Consider an airline that operates Boeing 737-500 and Airbus A321 aircraft. These aircraft types differ in size and range. The

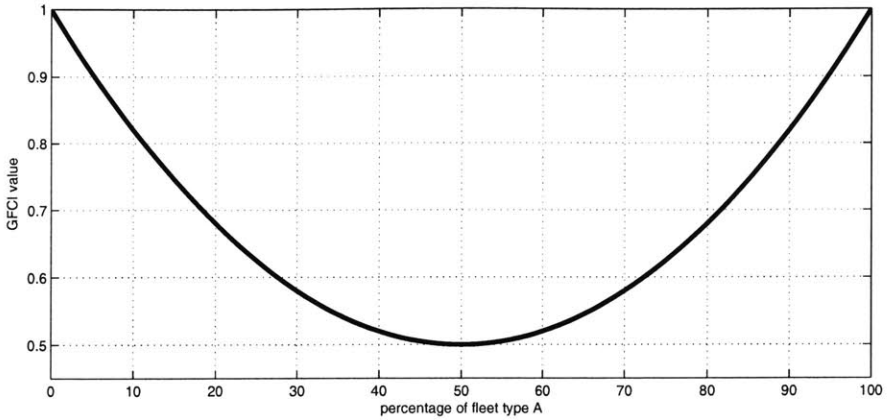

(a) Fleet with two Aircraft Types

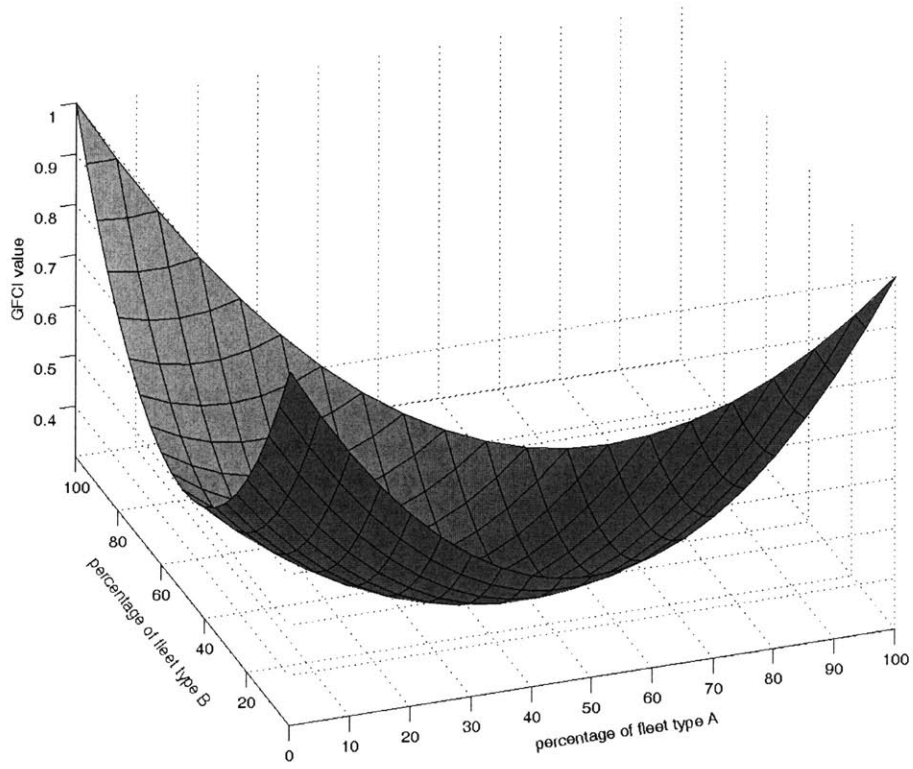

(b) Fleet with three Aircraft Types

Figure 3-5: GFCI examples

Airbus A321 has a 77% greater seat capacity than the Boeing 737-500 and due to its almost 70% greater range, the Airbus A321 operates on longer flights and thus has fewer departures per day than the Boeing 737-500 [3]. The properties of these two aircraft in this example are shown in Table 3.1.

Table 3.1: Properties of Aircraft for Example Airline

fleet	number of a/c	seats per a/c	departures per day per a/c
Boeing 737-500	10	103	6.2
Airbus A321	variable	182	4.7

In Figure 3-6 the GFCI values as a function of the number of Airbus A321 aircraft are shown for three different definitions of the weight w_f. The three curves have the same basic shape and minimum GFCI value (0.5), however, are spread out along the horizontal axis differently due to the different weights, which in this case are linearly related.

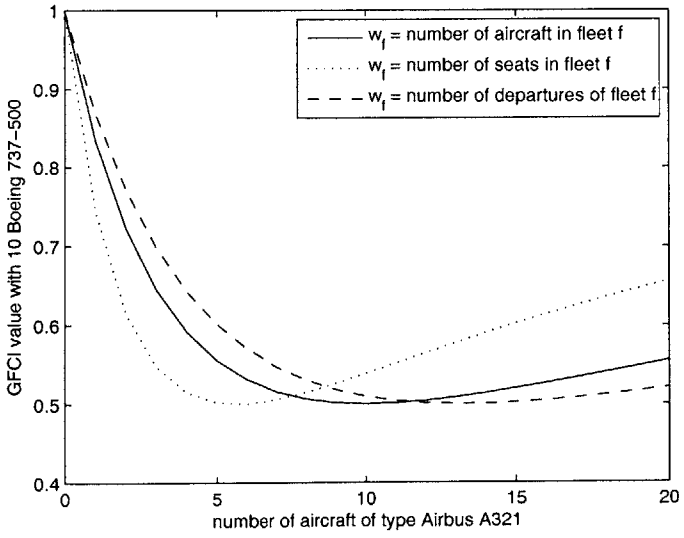

Figure 3-6: Influence of Choice of w_f on GFCI

3.2.3 Swap Option Indicators

If multiple aircraft are at the ground at the same airport at the same time, these aircraft could potentially be swapped between flights[5, 29]. The global single swap options per departure (GSD) are calculated using Equation 3.4. The term *global* indicates that it is assumed that swapping can occur globally within or between fleets. The term *single* indicates that single swap options are considered, i.e., swaps between only two aircraft at a time. The GSD is given by

$$GSD = \frac{\sum_{a \in A} \sum_{g \in G_a} \sum_{h \in G_a / \{g\}} (\delta_{g,h})}{2 \times |D|} \tag{3.4}$$

where A is the set of all airports serviced, D is the set of all departures, G_a is the set of all ground arcs at airport a, and $\delta_{g,h}$ is equal to 1 if ground arcs g and h overlap in time and 0 otherwise.

Swapping two aircraft of the same type will likely be less complicated than swapping two aircraft of different type (see 3.2.2). To account for this difference in ease of swapping aircraft, each swap option is multiplied by a weighting coefficient c_{f_1,f_2}, which is representative of the ease of swapping aircraft of fleet type f_1 and f_2. Incorporating these weighting coefficients into Equation 3.4 yields the weighted single swap options per departure (WSD)

$$WSD = \frac{\sum_{a \in A} \sum_{g \in G_a} \sum_{h \in G_a / \{g\}} \left(c_{f(g),f(h)} \times \delta_{g,h}\right)}{2 \times |D|} \tag{3.5}$$

where c_{f_1,f_2} a measure of the ease of swapping two aircraft of fleet type f_1 and f_2, and $f(g)$ is the fleet type of ground arc g. Larger and lower values of c_{f_1,f_2} indicate that swapping of aircraft of fleet types f_1 and f_2 is less and more complicated, respectively. Note that $c_{f_1,f_2} = c_{f_2,f_1}$, because a switch from fleet type f_1 to fleet type f_2 on one flight also means a switch from fleet type f_2 to fleet type f_1 on another flight.

The following illustrates the computation of the GSD and WSD for the example

introduced in Figure 3-4. In this example, the GSD and WSD are given by

$$GSD = \frac{1+3+2+2}{2 \times 4} = \frac{8}{8} = 1$$

and

$$WSD = \frac{c_{B,A} + 3c_{A,B} + c_{B,A} + c_{B,B} + c_{B,A} + c_{B,B}}{2 \times 4} = \frac{6c_{A,B} + 2c_{B,B}}{8}$$

For $c_{A,B} = c_{B,B} = 1$, the values for the GSD and the WSD are identical.

The choice of the weighting coefficients depends on the airline and its operation. Sets of weighting coefficients can be evaluated with historic and/or simulation data for the correlation of the WSD with the operational performance. Also, the set of weighting coefficients that provides the best correlation between the WSD and the operational performance can be determined by a best fit optimization of the historic and/or simulation data. Only values of the WSD computed using identical sets of weighting coefficients can be used to compare and benchmark different airline schedules.

In the above indicator, the possible swaps are assessed once for the duration of each ground arc. A different way to assess the possible swaps is to assess all possible swaps at each time step. Here we use consecutive minutes as time steps, because airline time data is typically available in a resolution of minutes. The number of possible swaps within fleet f at airport a at time step t, $N_{a,t}^{f,f}$, is given by

$$N_{a,t}^{f,f} = \begin{cases} \frac{n_{f,a,t}!}{2(n_{f,a,t}-2)!} & \text{if } n_{f,a,t} \geq 2; \\ 0 & \text{if } n_{f,a,t} < 2. \end{cases} \tag{3.6}$$

where $n_{f,a,t}$ is the number of aircraft of fleet type f that are on the ground at airport a at time step t.

The number of possible swaps between two different fleets f and g at airport a at

time step t, $N_{a,t}^{f,g}$, is given by

$$N_{a,t}^{f,g} = \frac{n_{f,a,t}!}{(n_{f,a,t}-1)!} \times \frac{n_{g,a,t}!}{(n_{g,a,t}-1)!} = n_{f,a,t} \times n_{g,a,t} \qquad \text{if} f \neq g. \tag{3.7}$$

The total number of possible swaps at airport a at time step t, $N_{a,t}$, is given by

$$N_{a,t} = \sum_{f \in F} \left(N_{a,t}^{f,f} + \frac{1}{2} \sum_{g \in F/\{f\}} N_{a,t}^{f,g} \right) \tag{3.8}$$

The factor of $\frac{1}{2}$ is introduced in Equation 3.8, because every possible swap between different fleets is counted twice, once each way.

Evaluating Equation 3.8 for all airport and all time steps gives an overall measure of possible swaps. This number divided by the number of time steps and the number of departures gives the global time averaged swap options per departure (GTSD)

$$GTSD = \frac{\sum_{a \in A} \sum_{t \in T} N_{a,t}}{|T| \times |D|} \tag{3.9}$$

$$= \frac{\sum_{a \in A} \sum_{t \in T} \sum_{f \in F} \left(N_{a,t}^{f,f} + \frac{1}{2} \sum_{g \in F/\{f\}} N_{a,t}^{f,g} \right)}{|T| \times |D|} \tag{3.10}$$

where A is the set of all airports, T is the set of all time steps and F is the set of all fleets.

Equation 3.10 is in a form in which weighting coefficients can be applied which yields to weighted time averaged swap options per departure (WTSD)

$$WTSD = \frac{\sum_{a \in A} \sum_{t \in T} \sum_{f \in F} \left(c_{f,f} \times N_{a,t}^{f,f} + \frac{1}{2} \sum_{g \in F/\{f\}} c_{f,g} \times N_{a,t}^{f,g} \right)}{|T| \times |D|} \tag{3.11}$$

The following example illustrates the computation of the WTSD. The shaded scale in Figure 3-7 represents the number of aircraft on the ground in the example introduced in Figure 3-4 as a function of time. The number of Boeing 737-500 and Airbus A300

aircraft on the ground are shown above and below the time axis, respectively.

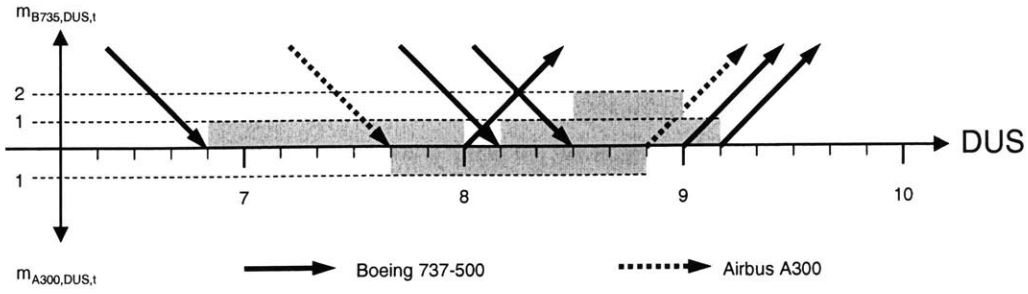

Figure 3-7: Timeline of Ground Arcs of Aircraft of Different Fleet Types

The WTSD evaluated for DUS from 6:30 to 9:30 in one minute time steps is

$$\begin{aligned} WTSD &= \frac{1}{2 \times 180 \times 4}(80c_{A,B} + 30c_{B,B} + 80c_{B,A}) \\ &= \frac{1}{48}c_{B,B} + \frac{1}{9}c_{A,B} \end{aligned}$$

As before, For $c_{A,B} = c_{B,B} = 1$, the values for the GSD and the WSD are identical. The choice of the weighting coefficients depends on the airline and its operation. Sets of weighting coefficients can be evaluated with historic and/or simulation data for the correlation of the WTSD with the operational performance. Also, the set of weighting coefficients that provides the best correlation between the WTSD and the operational performance can be determined by a best fit optimization of the historic and/or simulation data. Only values of the WTSD computed using identical sets of weighting coefficients can be used to compare and benchmark different airline schedules.

3.3 Output Measures

The operational performance of an airline can be assessed in different ways. Here we will assess it by measures of on-time performance and delays. On-time performance and delays can generally be assessed by four different references: off-blocks, take-off, landing and on-blocks. Here we will use on-blocks on-time performance and delays

because it is the most meaningful for aircraft and passenger connections. The on-blocks time determines the next possible earliest departure time for aircraft (after MinGT) and for passengers (after minimum connection time). Thus, it determines whether a passenger or an aircraft can make a connection. Other time references have additional non-deterministic time spans in between them and the next possible departure (e.g., the taxi-in time in case of the landing time). Therefore, they are less predictive of the next possible departure time of a aircraft or passenger than the on-blocks time.

3.3.1 Delay Minutes

The number of delay minutes per flight is defined as the difference between the actual on-blocks time and the scheduled on-blocks time of the flight. If a flight arrives prior to the scheduled arrival time, we will consider the number of delay minutes for this flight to be zero. For some airlines, early arrivals may result in operational difficulties (e.g., if no gates are available at the earlier time) and in these cases early arriving aircraft need to be considered in the output measure. Here we assume that no penalty for early arrivals exists. The number of total delay minutes (TDM) is defined as the sum of the number of delay minutes of all flights.

Two basic statistics can be computed for the set of the delay minutes for all flights: the mean and the standard deviation. The mean of the set provides the average number of delay minutes per flight arrival and is defined as the ratio of TDM to the total number of flight arrivals. The variance of the set provides a measure for how wide the distribution of the delay minutes for all flights is.

The distribution of on-blocks delays for an example day at a hypothetical airline is shown in Figure 3-8(a). The large spike at zero occurs since all flights that arrived at or before the scheduled on-blocks time are included in this bin and not only the flights that arrived at the scheduled on-blocks time. The mean of the distribution, and thus the average flight delay per aircraft, is 8.1 minutes.

3.3.2 On-Time Performance

On-time Performance (OTP) is typically measured as the ratio of delayed flights to overall flights. A flight is on time if it arrives at the gate (on-blocks) before, at or within a certain cutoff time, t, after the scheduled on-blocks time. Correspondingly, a flight is delayed if it arrives more than t minutes after its scheduled on-blocks time. Canceled flights are not included in the denominator or the numerator. Often, airline OTP is evaluated using a cutoff of 15 minutes because the Federal Aviation Authority (FAA) in the USA measures OTP in this way [22].

The OTP as a function of cutoff time t for the delay distribution presented in Figure 3-8(a) is shown in Figure 3-8(b). The OTP will always be monotonically increasing with increasing cutoff time, because an larger cutoff time will never result in fewer flights to be on time than a smaller cutoff time. The OTP approaches 100% as the cutoff time approaches the largest flight delay in the set of flight delays considered (here 51 minutes).

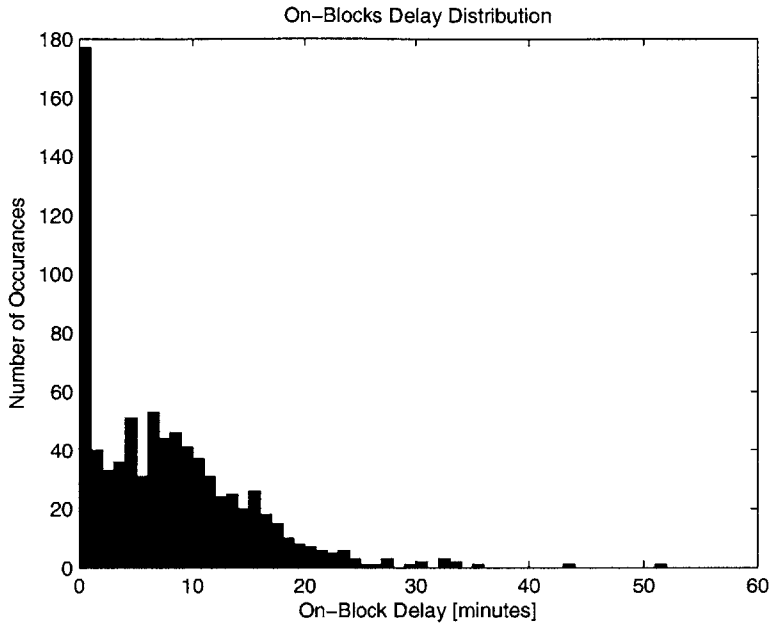

(a) Example On-Blocks Delay Distribution

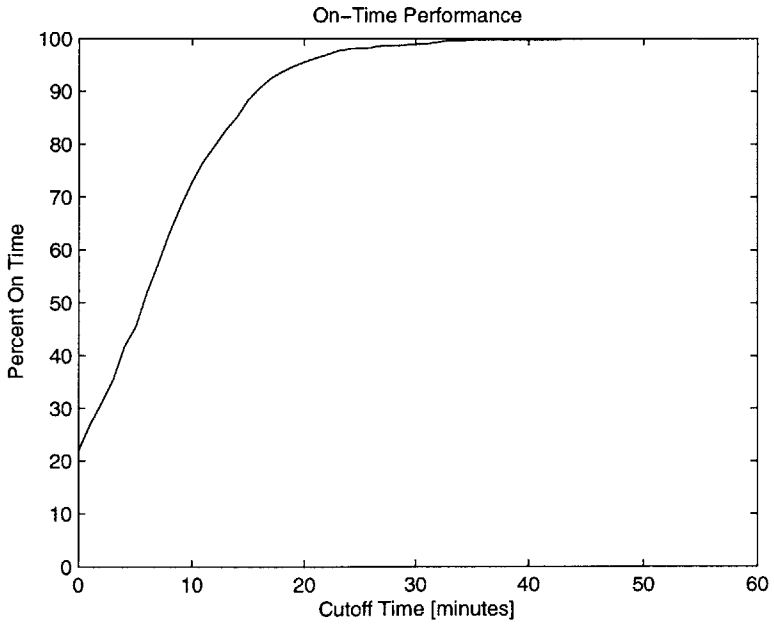

(b) Example On-Time Performance vs. Cutoff Time

Figure 3-8: Example On-Blocks Delay Distribution and On-Time Performance

Chapter 4

Application of Control Theory in Airline Operations

Control theory is a field of study that examines the control of the behavior of dynamic systems [30, 23, 8]. Control engineering is the application of control theory to design systems that change real world dynamic plants so that their behavior becomes more desirable. To achieve this, the plants are modeled as mathematical abstractions and controllers for these models are designed within the framework of control theory.

4.1 Background

In control theory two major groups of controlled systems exist: Open loop and closed loop controlled systems. The general block diagram for an open-loop controlled system (enclosed by the dashed line) is shown in Figure 4-1(a). The input to the system is a desired output and the output of the system is the actual output. Within the system a controller creates a control signal based on a model of the plant to achieve a desired actual output from the plant. Disturbances may affect the plant in its behavior.

A simple example of an open loop system is a heater that does not consider the current temperature in a room. The desired temperature (output) is entered through the thermostat. The controller opens or closes the hot water valve to the heater

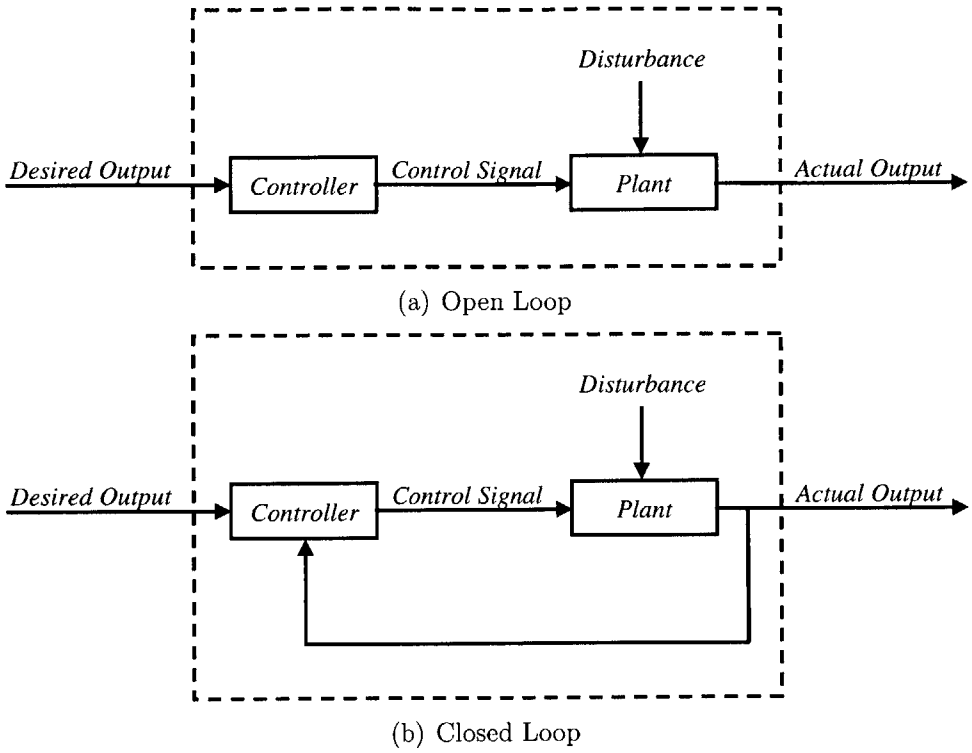

Figure 4-1: Control System Block Diagrams

according to the thermostat setting and based on a model for how heat is lost to the outside world. The room is the plant and its output is the room temperature. A disturbance could be an open window that leaks cold air into the room at an unknown rate and temperature. If the thermostat is set to a certain setting with the window closed, the room will reach and maintain a certain temperature due to the heater. If now the window leaks cold air into the room, the heater will still output the same amount of heat as before, however now the cold air from the outside will cool the room and the new room temperature will be colder than the room temperature in the case of the closed window. In this open loop case the hot water value depends only on the thermostat setting and not on the actual temperature in the room.

If, as in the example above, unknown disturbances exist or if the model of the plant which the controller is based on is not accurate, the actual output of the closed loop system will deviate from the desired output. Since the controller does not know about the current state of the system, it is unable to adjust the control signal accordingly.

This problem is overcome by closing the loop, i.e. feeding the actual output back to the controller. The resulting system is called a closed-loop (or feedback) control system (see Figure 4-1(b)). In the example above, this feedback may be completed with a thermometer that measures the current room temperature. A closed loop controller would take into account not only the thermostat setting but also the current temperature of the room and open or close the hot water valve considering information from both sources.

In classical control engineering, a three step approach is taken to implement a controlled system. The first step is the system identification in which the input-to-output characteristics of the plant are determined. The plant is probed in a structured way and its response is analyzed. For this step, the plant needs to be built or simulated. In the case of an airline the likely choice is simulation, since it would be very expensive to probe an actual airline for its response to extreme inputs. The second step is the system modeling in which a mathematical model of the plant is developed based the knowledge of the system gained during system identification. The third and final step is the controller design. A controller is developed to control the plant such that the overall system has a desirable behavior.

4.2 Example of an Application of Control Theory in Airline Operations

The following simplified example demonstrates a possible application of control theory to an airline operation[1]. Suppose that the management of an airline sets an acceptable level of delay minutes per flight arrival (DMPA) (see 3.3). The management might base the decision for this level on its understanding that a lower DMPA will be expensive to reach and a higher DMPA will make customers unhappy. Initially, the acceptable level DMPA is 14 minutes. The airline's scheduling department (SD) therefore adjusts delay buffers to achieve the desired DMPA as well as possible over time.

[1]Note that all variables used in this example are scalar and that the system is linear.

In this example, the following abstraction of the airline operation (plant) is used: The number of DMPA at a given discrete time t is DMPA_t, where we will use the abbreviated notation $t+n$ to indicate the time at n time steps after t. DMPA is measured in units of minutes. Due to delay propagation it is assumed that DMPA_t depends on DMPA_{t-1}. Here we assume that delays cannot be reduced during flight so that $\mathrm{DMPA}_t \propto \mathrm{DMPA}_{t-1}$. For this example we also assume that delay can be reduced by buffers such that $\mathrm{DMPA}_{t+1} \propto -\mathrm{BPD}_t$ where BPD_t is the buffer per departure at time t in minutes. A third component that influences DMPA_t is the weather. Here, the weather impact, $\mathrm{W_t}$, expresses the delay equivalent effect of the weather at time t. Using the approximation that $\mathrm{DMPA}_t \propto \mathrm{W_t}$, a high value of $\mathrm{W_t}$ represents bad weather that will cause more delays. $\mathrm{W_t}$ is also measured in units of minutes. Combining the components affecting DMPA yields the following plant model:

$$\mathrm{DMPA}_{t+1} = \mathrm{DMPA}_t - \mathrm{BPD}_t + \mathrm{W_t} \tag{4.1}$$

The airline has no influence over the weather $\mathrm{W_t}$. It also cannot change the past and therefore, has no influence on DMPA_t. It can, however, adjust the buffers in the schedule by rerouting aircraft and/or making use of reserve aircraft. Thus, BPD_t can be considered the control signal in this example. The block diagram shown in Figure 4-2(a) illustrates this airline as an open loop system. The airline scheduling department is not aware about the current state of the operation, i.e., DMPA_t. Flight scheduling is conducted ahead of time and can therefore only be based on estimates of the operational environment. The scheduling department can therefore only provide open loop control. In this example this means that the airline cannot adjust to weather changes or recover from weather driven delays; the scheduling department will not have knowledge about the weather changes or the delays.

In addition to the Scheduling Department, airlines have an Operations Control Center (OCC), which monitors the operation and takes recovery actions if necessary. The OCC is aware of the current state of the operation and thus a feedback loop is

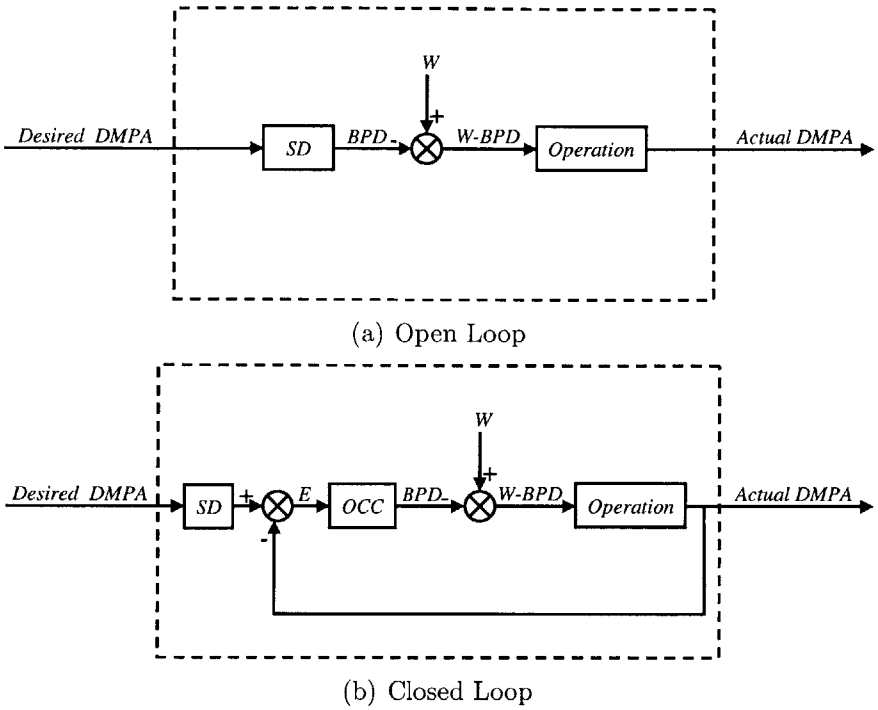

Figure 4-2: Airline Delay Control System

added to the system (see Figure 4-2(b)). If the desired DMPA and the actual DMPA differ, the error $E_t = \text{DMPA}_{desired} - \text{DMPA}_t$ is nonzero and the OCC should choose a BPD value to reduce the error. In this example, a possible implementation of the OCC is

$$\text{BPD}_t = k \times E_t = k \times \text{DMPA}_{desired} - k \times \text{DMPA}_t$$

where k is a constant gain. The controller in this case is called a proportional controller, because it creates a control signal proportional to the error. Using this implementation of the OCC, Equation 4.1 becomes

$$\begin{aligned} \text{DMPA}_{t+1} &= \text{DMPA}_t - (k \times \text{DMPA}_{desired} - k \times \text{DMPA}_t) + \text{W}_\text{t} \\ &= (1+k) \times \text{DMPA}_t - k \times \text{DMPA}_{desired} + \text{W}_\text{t} \end{aligned} \tag{4.2}$$

The response of the proportionally controlled system to a change of the desired DMPA in good weather with $k = -0.5$ is shown in Figure 4-3. At time $t = 15$, man-

agement decides to fight the airline's reputation of being often delayed and changes the desired DMPA from 14 minutes to 10 minutes. In this example, it is assumed that the weather is constantly good and thus $W_t = 0$ for all t. When the change in

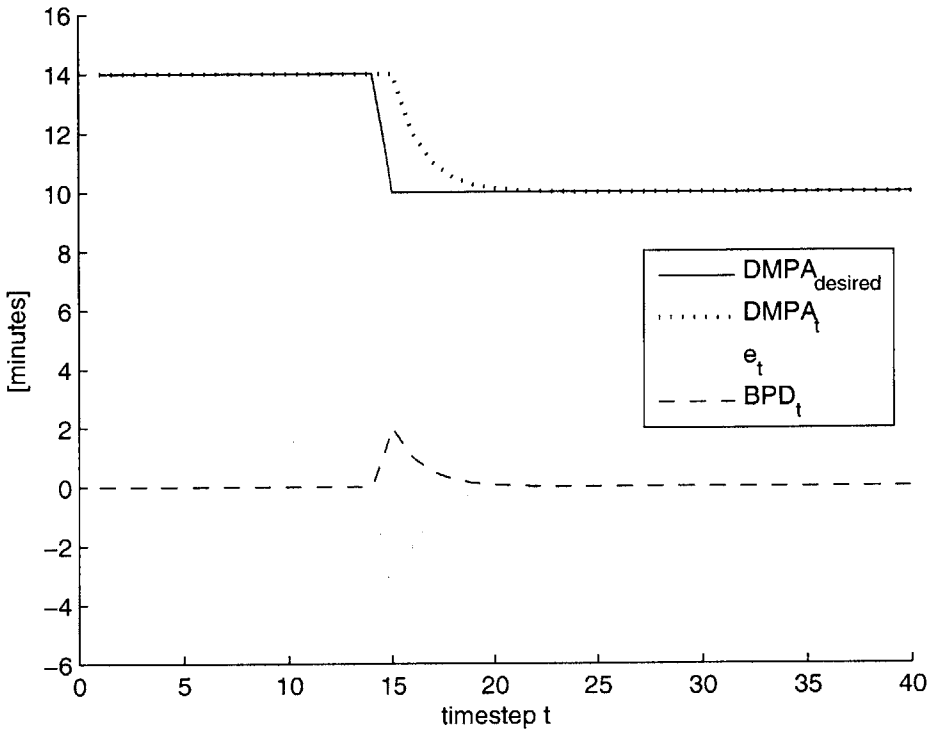

Figure 4-3: Response of the P-controlled System to a Change in $\text{DMPA}_{desired}$

$\text{DMPA}_{desired}$ occurs, the error becomes negative ($E = -4$). As a result, the P controller issues a positive control signal (BPD = 2). Conceptually speaking, the OCC implements buffers to reduce delays. Over the next few time steps the actual DMPA approaches the new desired DMPA and the error becomes zero, as does the control signal.

The choice of the controller gain k determines the performance of the system. There exist a range of values for k at which the system will reach the desired DMPA over time. In this example, the range is $-2 < k < 0$. For any value of k within this range, the system will be stable and the actual DMPA will approach the desired DMPA (see Figure 4-4(a)). For $k = -1$ Equation 4.2 becomes $\text{DMPA}_{t+1} = \text{DMPA}_{desired} + \text{W}_\text{t}$ and in the absence of weather effects this system will have the

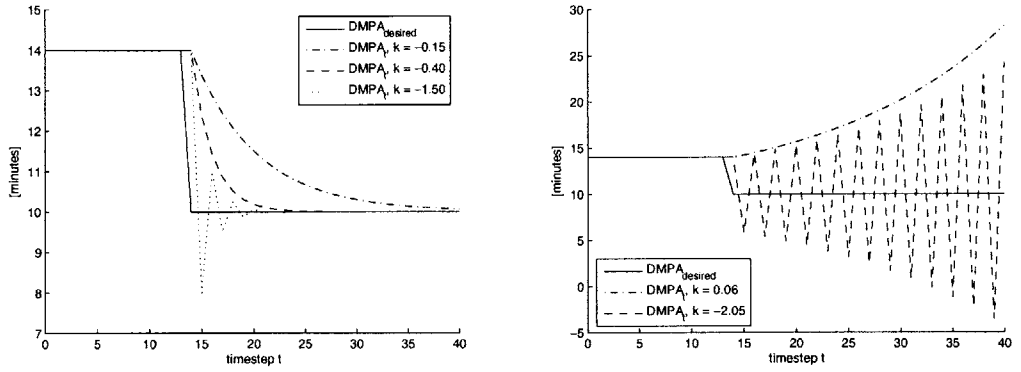

(a) stable (b) unstable

Figure 4-4: Response of the System for Different Gain Values

most desirable performance in terms of fast, accurate tracking of $\text{DMPA}_{desired}$.

This controller is called an inversion of the plant. Unfortunately, for complex systems it is often costly or impossible to construct a controller capable of inverting the plant. Outside the range $-2 < k < 0$ the system becomes unstable. For a positive value of k (e.g., $k = 0.03$), the error increases exponentially, and the desired DMPA is never achieved (see Figure 4-4(b)). For values of k smaller than -2 (e.g., $k = -2.1$), the system oscillates with increasing magnitude.

In the above example it was assumed that $W_t = 0$ for all t. Now we introduce negative weather effects. In a simple weather scenario, initially constant good weather ($W_t = 0$) exists for $t < 15$ and constant bad weather with $W_t = 2$ exists for $t \geq 15$. The response of the system with $k = -0.5$ is shown in Figure 4-5(a). The proportional controller is unable to compensate for the weather disturbances and the system settles with a steady error $E = -2$. A way to improve the steady state performance of the system is to introduce an additional term to the controller. A possible addition to the proportional term that will improve the system's performance is an integral term. The resulting controller is called a proportional-integral (PI) controller. In the integral term, all errors in the past and the error at the current time are summed up (or in a continuous system the integral of the error is taken). I_t is the sum of all errors from the beginning of the controlling of the system, i.e., $I_t = \sum_{u \leq t} E_u$. Multiplied with a

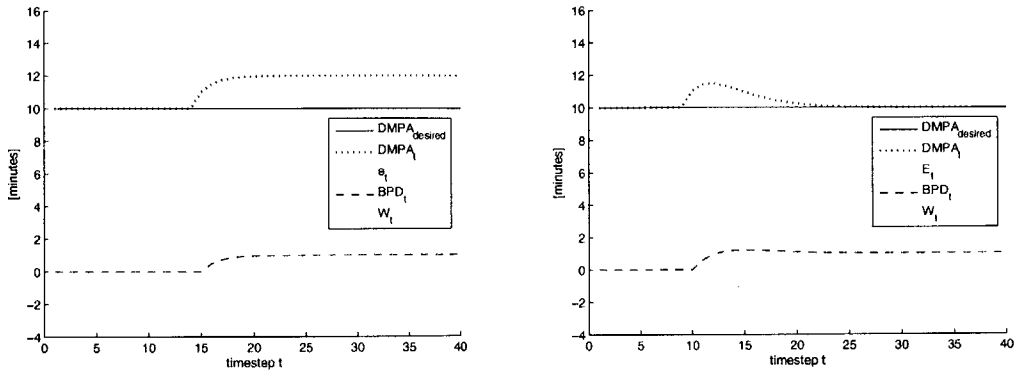

(a) P (b) PI

Figure 4-5: Response of the System to Constant Weather

gain m this sum is used as part of the control signal to eliminate steady state offsets in the system. The PI controller is described by

$$\text{BPD}_t = k \times E_t + m \times I_t$$

The response of the system to the change in weather described above with the PI controller with parameters $k = -0.5$ and $m = 0.1$ is shown in Figure 4-5(b). The response of the proportional and the proportional-integral controlled systems to a random weather disturbance that is uniformly distributed between $W_t = 1$ and $W_t = 3$ for $t \geq 10$ is shown in Figure 4-6. While the P controlled system's output DMPA

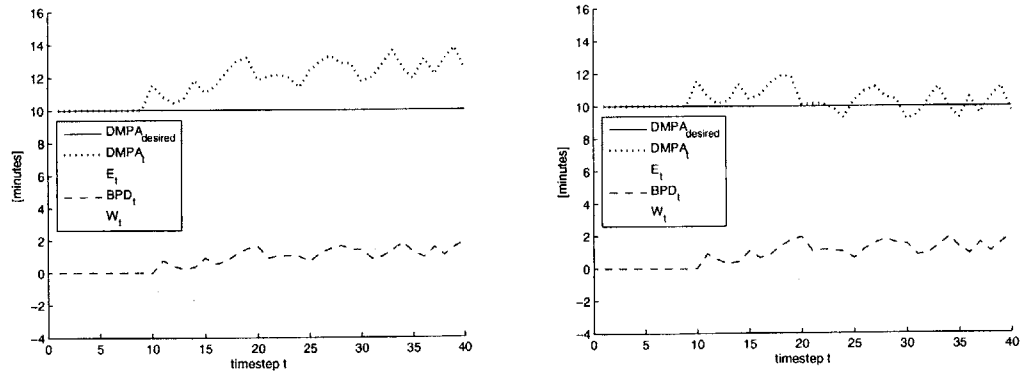

(a) P (b) PI

Figure 4-6: Response of the System to Random Weather

stays constantly above the desired DMPA, the PI controlled system operates closer around the desired DMPA. The mean errors are -2.2078 and -0.4124, and the mean squared errors are 5.565 and 0.739 for the P and the PI system, respectively.

Even though the above example relies on a highly simplified model of the airline operation, it can serve as a motivation to further investigate ways to apply control theory to airline operations. Modern control theory provides methods and tools to develop controllers for a wide range of plants, including nonlinear and/or time-variant plants with multiple input and/or output variables. To take advantage of these methods and tools in airline operations the properties of the operation need to be determined and input and output variables need to be chosen. The proposed simulation study presented in Chapter 5 and the proposed analysis described in Chapter 6 are initial steps in determining the properties and corresponding modeling parameters of airline operations.

Chapter 5

Simulation Setup

A simulation study should be conducted to provide data for an initial analysis of the KPIs described in Chapter 3. The simulation environment and a benchmark data set are described in this chapter. A proposed analysis of the data gained from the simulation is presented in Chapter 6.

The goal of this simulation experiment is to create a airline operation that is representative of real airline operations. Therefore, even though the simulation setup is based on the operation of Lufthansa, it does not model Lufthansa's operation in all details and cannot directly be utilized to predict future performance of Lufthansa. However, the simulation can serve as a source of data for fundamental research and as a testbed for airline scheduling strategies.

5.1 MIT Extensible Air Network Simulation

The flow of aircraft through the air network and the effects of weather on airport capacity are simulated with the MIT Extensible Air Network Simulation (MEANS). MEANS is an event-based simulation of traffic flows in the air transportation network [14, 18, 16]. It simulates the flow of aircraft, crew and passengers through a network of queues to model capacity constraints in the air transportation network. MEANS was initially designed in 2001 [17] and has been expanded in subsequent years. It is written in C++ and can be run on Linux and Unix systems. MEANS is used for the

research of new concepts in traffic flow management, airline schedule planning, and operations.

The modular structure of MEANS is shown in Figure 5-1. Aircraft flow through the different flight stage modules (gate, taxi, tower, en route) along the dotted lines. Two more modules (TFM, airline) make decisions about the schedule to be executed by the flight stage modules. An additional weather module provides weather information to other modules.

For each of these modules there exist a variety of implementations that can be used. The gate module determines the aircraft off-blocks time considering the airline

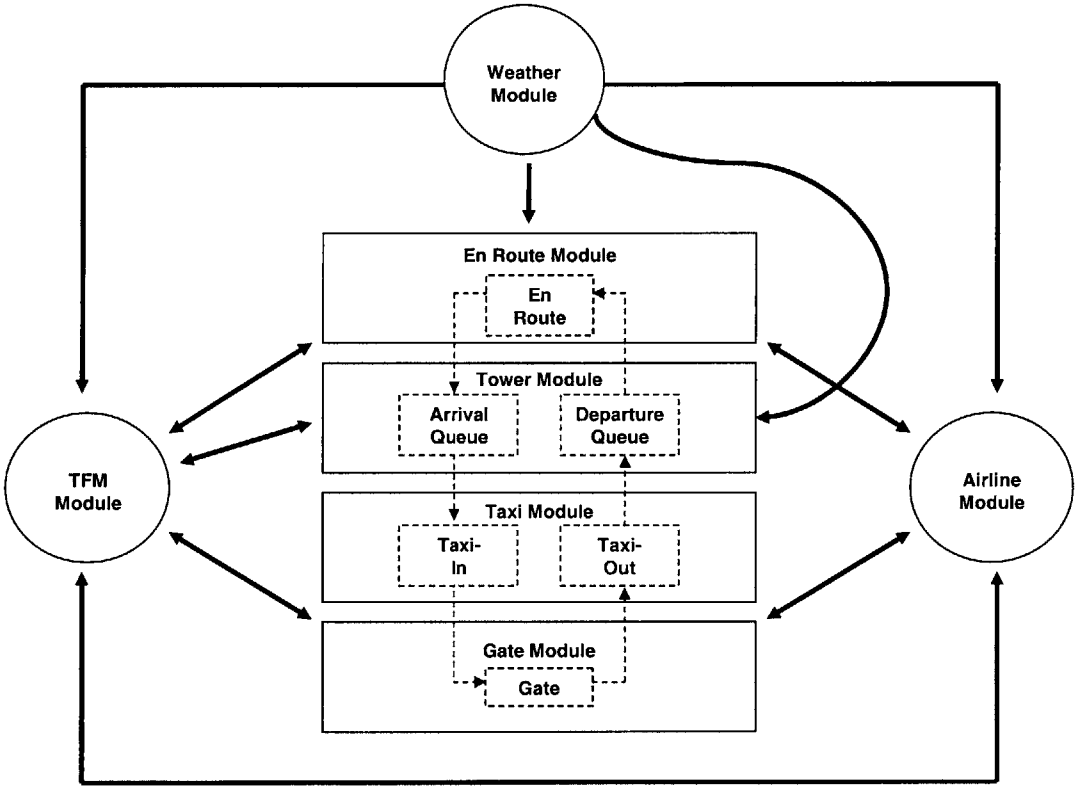

Figure 5-1: MEANS Modular Structure

schedule, the last on-blocks time of the aircraft, the MinGT for the aircraft, and traffic flow management programs. The implementation used in the simulation platform presented here utilized MinGTs provided by Lufthansa. No model for aircraft unavailability resulting from mechanical failures is used. The taxi module determines the taxi-out time (off-blocks to runway departure queue) and taxi-in time (landing to on-

blocks). The implementation of the taxi module used here determines the taxi times from taxi time distributions derived from historic data provided by Lufthansa (see 5.3.3). The tower module controls the departure and arrival queues at the airports. It also determines the available capacity at an airport using airport capacity profiles (see 5.3.5) and assigns take-off and landing times to flights. The en route model determines the flight time for each flight between take-off and the arrival queue of the destination airport. The implementation of the en route model used in the simulation platform described here does not model the details of en route flights such as actual aircraft routings, wind and weather effects, or en route congestion. The flight times are drawn from flight time distributions derived from historic data provided by Lufthansa (see Section 5.3.4). The Traffic Flow Management (TFM) module implements traffic flow management programs, such as ground delay programs (GDP). Here, no traffic flow management programs are used and the TFM module performs no function. The weather module determines the actual weather and forecasts at the airports and provides this information to the en route and tower modules. In the simulation platform presented here, en route weather is not simulated and only actual airport weather without forecasts is considered (see Section 5.3.6). The airline module tracks and recovers the airline schedule. It can make recovery decisions for aircraft, crew, and passengers. In the simulation platform presented here, crew and passengers are not tracked or recovered and the aircraft tracking and recovery is performed by the Integrated Operations Control System (IOCS).

5.2 Integrated Operations Control System

The Integrated Operations Control System (IOCS) is an airline operations control tool by Carmen Systems AB[1] [1]. It provides an interface for tracking and managing airline operation and decision support as well as automated recovery solvers. IOCS covers the aircraft, crew, and passenger domains and solves for optimal or close to optimum recovery options across the three areas focusing on passenger service and cost

[1]Carmen Systems AB, Odinsgatan 9, SE-411 03 Göteborg, Sweden

control. The scalable design of IOCS allows stepwise implementations and integration with existing systems. IOCS can be run on Unix or Linux servers in conjunction with an Oracle[2] database and takes standard and specialized Extensible Markup Language (XML) input streams.

The implementation of IOCS used in the simulation platform described here acts as the airline module within MEANS. It tracks aircraft and provides aircraft recovery decision. Crew and passengers are not tracked or recovered. IOCS reevaluates and resolves the aircraft situation every 30 minutes in simulation time.

5.3 Simulation Input

5.3.1 Flight Schedules

The benchmark input flight schedules used in the simulation study are based on Lufthansa's continental flight schedule of the months November 2004, December 2004 and January 2005. For each simulation run, one benchmark flight schedule is created by extracting three consecutive days from the historical flight schedule. Only the first of these three days is simulated. The remaining two days are included in the simulation to enable IOCS to make recovery decisions that consider the next two days.

5.3.2 Aircraft Assignments

The benchmark aircraft assignments are based on the aircraft assignments provided by Lufthansa along with the flight schedules. For each day, the chains (or strings) of flights that are serviced by a single aircraft in the original schedules are kept intact. Between days, each of these chains of one day is matched with a chain of the next day on a first come first serve basis at the overnight station while ensuring that matched chains have the same fleet type. Specific aircraft (tailnumbers) are then assigned to the resulting, jointed chains. For some simulation runs the global fleet composition

[2]Oracle Corporation, 500 Oracle Parkway Redwood Shores, CA, 94065, USA

is changed from the one provided with the schedules. In these cases, one or multiple chains that are originally operated by the Boeing 737 fleet are changed to be operated by the Airbus A320 fleet.

5.3.3 Taxi Times

In the benchmark input data sets, all aircraft types have the same distribution of taxi times. For FRA and MUC, taxi time distributions are derived from historic data provided by Lufthansa for April, May, August, and September 2005. These warmer months of the year were chosen to avoid expected effects of deicing in the winter months. The taxi times available from historic data include the actual time required to taxi to/from the runway, as well as waiting time during the taxi process and at the runway before take-off. To derive the actual time required to taxi to the runway as an input for the simulation (the waiting time on the ground is added by the tower module of MEANS), the number of historic samples was significantly truncated to exclude flights that were likely to have waiting time.

FRA and MUC have night curfews which prohibit aircraft from taking off or landing during the night. One exception to these curfews are domestic mail flights. Since almost no other traffic exists at night, it is assumed that mail flights during curfew hours never have to wait on the ground before departure due to traffic. Therefore, the taxi times recorded for these flights are assumed to not include any holding time on the ground and are used to derive the distributions of taxi times. The derived distribution for taxi times is shown in Figure 5-2. The cumulative distribution function (CDFs) are smoothed by a linear interpolation between integer minutes because the granularity of time in MEANS is seconds.

For all airports other than FRA and MUC deterministic taxi-in and time-out times based on historical data are used in the benchmark input data sets.

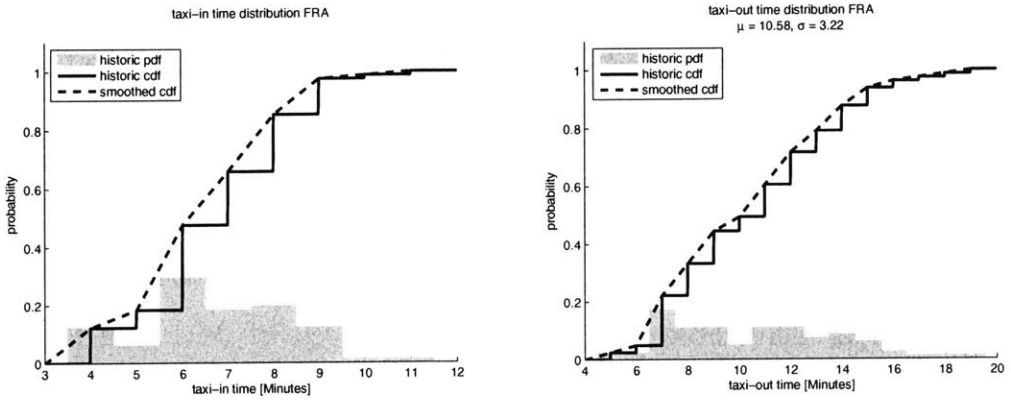

(a) FRA taxi-in

(b) FRA taxi-out

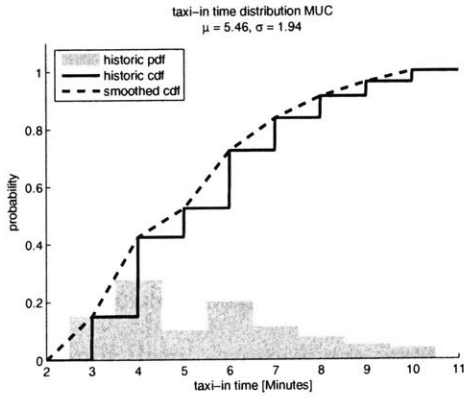

(c) MUC taxi-in

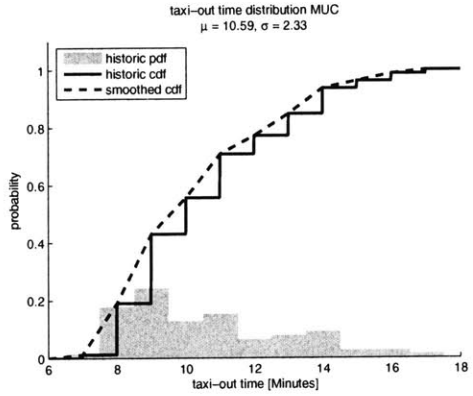

(d) MUC taxi-out

Figure 5-2: Taxi Time Distributions

5.3.4 Flight Times

Because airport traffic other than that of the airline in focus is not simulated at any airport other than FRA and MUC (see 5.3.7), it is not possible for any other airport to reach capacity. Therefore, only the simulation arrival queue in FRA and MUC will reach a notable size. The simulation arrival queue at the other airports will remain very small and, therefore, no significant simulated in-flight waiting time is expected at these airports. To account for this phenomenon, two different methods of determining the benchmark input flight times to MEANS are used, dependent on the arrival airport of the flight. Historic flight time records show the actual flight time, including in-flight waiting. For flights arriving at any airport other than FRA or MUC, this overall flight time is used as an input to MEANS. For flights arriving in FRA and MUC, the in-flight waiting time is determined by the MEANS tower module. Therefore, the flight time input to MEANS for FRA and MUC needs to be unimpeded (i.e., it cannot include the waiting time).

To derive unimpeded flight time distributions for flights arriving at FRA or MUC and overall time distributions for all other airports, an assumption is made that the total flight time consists of two components. The first component is the direct flight time without any traffic-induced in-flight waiting. The main influence on this component is the wind the aircraft encounters during its flight. It is assumed that the wind effect is distributed as a Gaussian distribution around a mean wind effect. As a result we assume that the direct flight time is also distributed as a Gaussian distribution. The second component is the in-flight waiting time. Here we assume that short in-flight delays occur much more frequently than long in-flight delays and are, therefore, assumed to be Exponentially distributed. The total flight time is a random variable $T = D + W$ where D and W are independent random variables that have Gaussian and exponential distributions, respectively. The derivation of the probability density function of T, $f_T(t)$, is given in Appendix B. The distribution of T is dependent on three input parameters: the mean (μ) and the standard deviation (σ) of the Gaussian distribution and the Poisson arrival rate (λ) of the Exponential

distribution. For each flight leg, these three parameters are determined for a best fit of $f_Z(z)$ to the historic data[3]. The mean direct flight time is σ, and the mean in-flight waiting time is $\frac{1}{\lambda}$.

Four examples of the best-fit Gaussian-exponential distribution for flights arriving at airports other than FRA and MUC are shown in Figure 5-3. The frequency of occurrence in the historical data is shown in the gray bars and the probability density function of the best-fit Gaussian-exponential distribution is shown as a solid line. The parameters of the best-fit Gaussian-exponential distribution as well as the R^2 values of the best-fit are shown below the graphs.

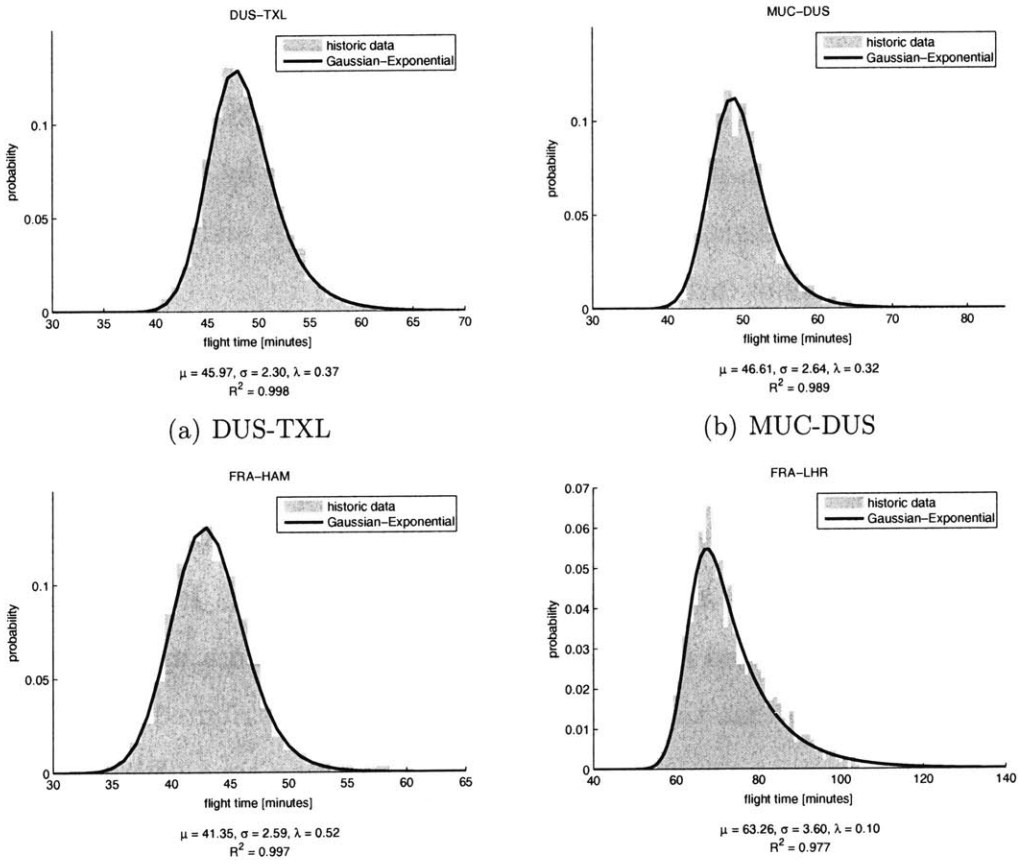

(a) DUS-TXL

(b) MUC-DUS

(c) FRA-HAM

(d) FRA-LHR

Figure 5-3: Example Flight Time Distributions

Four examples of the best-fit Gaussian-exponential distribution for flights arriving

[3]using the *fit* function provided with the Curve Fitting Toolbox of Matlab, Version 7.0.1.24704 (R14) SP1 by The MathWorks, Inc, 3 Apple Hill Drive Natick, MA, USA

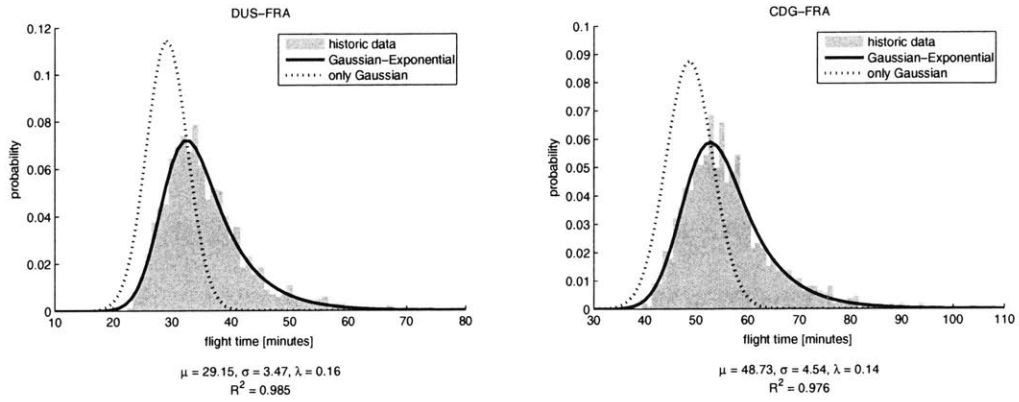

(a) DUS-FRA

(b) CDG-FRA

(c) FRA-MUC

(d) MAN-MUC

Figure 5-4: Example Unimpeded Flight Time Distributions

at FRA and MUC are shown in Figure 5-4. In addition to the historical data (gray bars) and the best fit Gaussian-exponential distribution (solid line), a Gaussian using the parameters of the Gaussian part of the Gaussian-exponential distribution is shown as a dashed line. The Gaussian distributions shown in the plots peak at a lower flight time value than the Gaussian-exponential distributions and the historic data. The addition of nonnegative in-flight waiting times will shift the peak to the right in the graphs. The Gaussian distributions are also narrower than the Gaussian-exponential distributions. The variance of a sum of two random variables, D and W is given by

$$\begin{aligned} \mathrm{Var(D+W)} &= E\left[(D+W-E[D]-E[W])^2\right] \\ &= E\left[(D-E[D])^2+(W-E[W])^2+2(D-E[D])(W-E[W])\right] \\ &= \mathrm{Var}(D)+\mathrm{Var(W)}+2\mathrm{Cov}(D,W) \end{aligned}$$

Since here D and W are independent,

$$2\mathrm{Cov}(D,W)=0$$

and, thus,

$$\mathrm{Var(D+W)} = \mathrm{Var}(D)+\mathrm{Var(W)}$$

and because the variance of a distribution of real numbers is always nonnegative

$$\mathrm{Var(D+W)} \geq \mathrm{Var}(D)$$

Here $\mathrm{Var}(W)=1/t^2$, which will in all practical cases be a positive number. Therefore, the Gaussian distribution will always be narrower than the Gaussian-exponential distribution.

5.3.5 Airport Capacity Profiles

An airport capacity profile defines a maximum number of operations (departures and arrivals) per given time frame envelope for an airport under certain conditions [7]. It provides the maximum number of possible departures per given time frame as a monotonically decreasing function of the number of arrivals and vice versa. An airport's capacity depends on many conditions: the number of runways, their layout, and whether they are used for take-offs, landings or both; the wind conditions which specify which of the available runways can be used (certain cross- and tail-wind limitations apply to aircraft taking off or landing); and the separation requirements between the aircraft, which are governed by the ceiling and visibility conditions.

In the benchmark input data sets, four different airport capacity profiles for different weather scenarios are used for both FRA and MUC: (1) a Visual Meteorological Conditions (VMC) profile represents the best weather case at which the maximum airport throughput is achieved; (2) a marginal VMC (MVMC) profile represents a case of poorer weather which, however, still enables at least some of the aircraft to operate at the airport with visual separation of other traffic; (3) an Instrument Meteorological Conditions (IMC) profile represents the scenario under which all aircraft and, therefore, need to maintain radar separation from each other (which is greater than visual separation); (4) a Low Visibility (LowVis) scenario under which special approach and departure procedures are in use and the throughput is lowest. The choice matrix for the airport capacity profiles is given in Table 5.1.

The airport capacity profiles in the benchmark input data sets are based on data provided by the Airport Coordination of the Federal Republic of Germany[4] and German Air Traffic Control [5] [31]. The derived airport capacity profiles for FRA and MUC are shown in Figure 5-5.

[4]Airport Coordination Federal Republic of Germany, Terminal 2-E, FAG-POB 37, D-60 549 Frankfurt/Main, Germany

[5]DFS Deutsche Flugsicherung GmbH, Am DFS-Campus 10, 63225 Langen

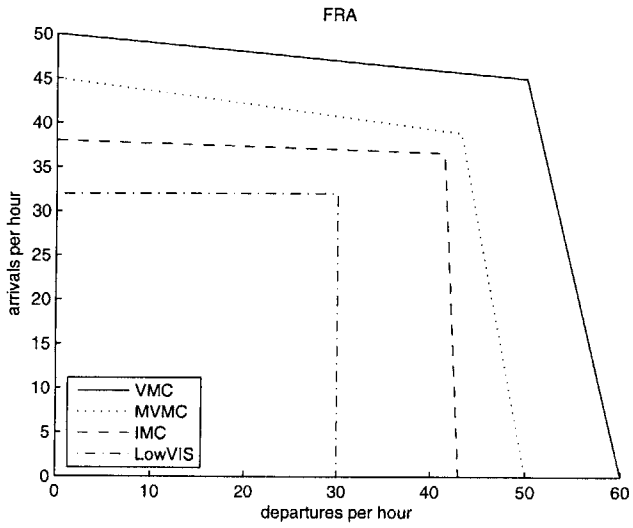

(a) FRA

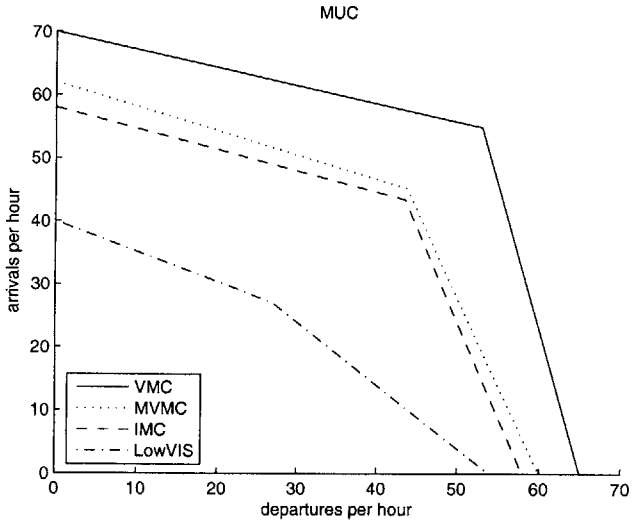

(b) MUC

Figure 5-5: Airport Capacity Profiles

5.3.6 Weather Scenarios

The weather inputs to MEANS in the benchmark input data sets are Meteorological Terminal Area Reports (METARs) provided by the Aviation Digital Data Service (ADDS)[6] of the National Weather Service (NWS)[7]. For FRA and MUC, the weather given by the METARs is assigned to a weather category according to the choice matrix given in Table 5.1. These weather categories determine the airport capacity profile used (see 5.3.5). The benchmark weather conditions at FRA and MUC for each simulation day are shown in Appendix C.

Table 5.1: Weather Categories [27]

Category	**Ceiling**		**Visibility**
VMC	greater than 3,000 feet AGL	and	greater than 5 miles
MVMC	1,000 to 3,000 feet AGL	and/or	3 to 5 miles
IMC	500 to below 1,000 feet AGL	and/or	1 mile to less than 3 miles
LowVis	below 500 feet AGL	and/or	less than 1 mile

5.3.7 Airport Traffic

Even though Lufthansa is the dominant carrier at both of its major hubs, it only accounts for 52% and 43% of the total operations at FRA and MUC, respectively [15]. Therefore, in any weather scenarios other than LowVis, neither of those two airports will be utilized up to its capacity by flights of Lufthansa only. Therefore, traffic other than the Lufthansa continental flights, which are included in detail in the benchmark flight schedules, (i.e., Lufthansa regional and intercontinental flights, flights of other airlines, the military, and general aviation) have to be modeled in the simulation to project realistic holding times in the air and on the ground at FRA and MUC.

In the benchmark input data sets, one such traffic scenario for each FRA and MUC

[6]http://adds.aviationweather.noaa.gov/

[7]US Dept of Commerce, National Oceanic and Atmospheric Administration, National Weather Service, 1325 East West Highway Silver Spring, MD, 20910, USA

is used. They are derived from airline schedule data from [15] and are presented in Figure 5-6. The traffic in Munich occurs mainly in two waves, a wave occurring in the morning and the afternoon. In between the two waves there is a two hour period of low traffic density. During the strict night curfew between 0:00 and 6:00 there is no scheduled traffic in MUC.

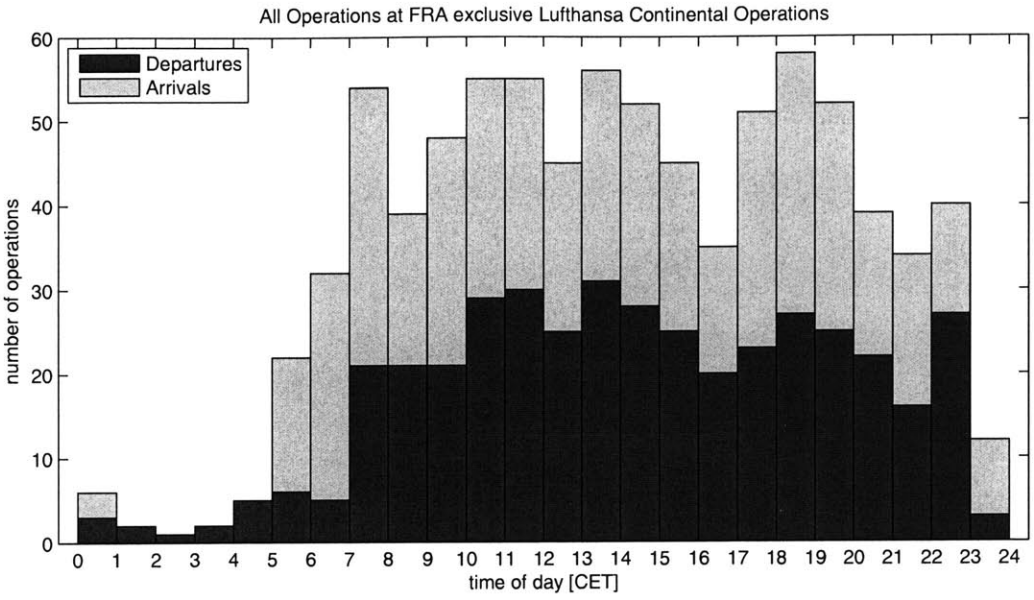

(a) FRA

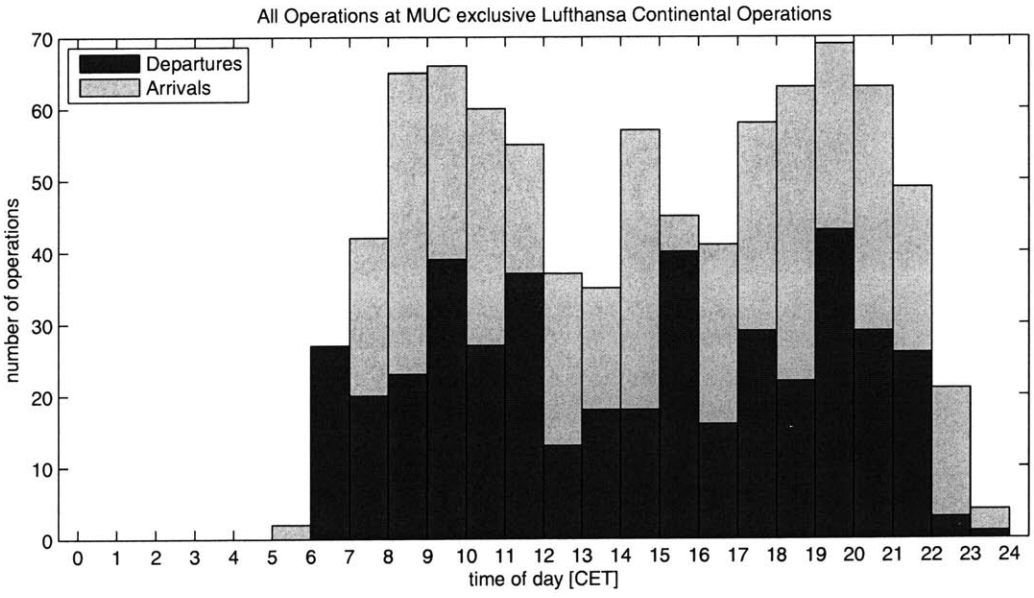

(b) MUC

Figure 5-6: Airport Traffic

Chapter 6

Proposed Analysis

The simulation platform described in Chapter 5 can be used to gather large sets of data for analysis. The benchmark input data prepared as part of this thesis as described in Chapter 5 will provide such a data set. This chapter proposes a methodology for analysis.

6.1 Initial Analysis

To analyze the correlations between the KPIs and the output measures, first calculate the KPIs defined in Chapter 3 for each input schedule in the input data set for a set of weighting coefficients and functions. Then compute the the output measures defined in chapter 3 for each simulation output. This provides a set of points, one for each KPI and output measure combination. The points lie on a two-dimensional coordinate system with the KPI values on one axis and the output measure values on the other axis.

Each of these sets of points is tested for correlation. A variety of fit-types (e.g., linear, exponential, polynomial) can be examined. A hypothetical plot of the values of a KPI (here GSD) vs. the corresponding values of an output measure (here 15 minute OTP) is shown in Figure 6-1(a). Each input schedule, which determines the GSD value, is simulated in a set of different weather scenarios and with different taxi and flight times, which are drawn randomly from the distributions described in

Sections 5.3.3 and 5.3.4, respectively. Therefore, identical schedules will produce a range of OTP values. As a result, there exist multiple OTP values for each GSD value.

The curve can be fitted to all available data points (see Figure 6-1(a)). To avoid strong influences of outlying data points, the curve fit can also be applied to a set of points that are derived from the complete set of data points. One such approach is to use the mean value of the set of output measure values for each KPI value (see Figure 6-1(b)). Another approach is to use the median value of the set of output measure values for each KPI value (see Figure 6-1(c)). The choice of approach for the analysis depends on the properties of the original complete set of points. Here, both approaches result in significantly different best fit curves, because the mean and median values differ. A concave curve is used for the mean and a convex curve is used for the median.

6.2 Optimization of Weighting Coefficients and Functions

In the initial analysis described in Section 6.1, the weighting coefficients and weighting functions for the KPIs were given. In addition to the optimization of the curve fit (e.g., minimizing the mean squared error) by adjusting the curve parameters, the weighting parameters of the KPI can be optimized to yield a better curve fit. For each distinct set of weighting parameters of the KPI, one best fit curve can be determined with standard curve fitting tools. The goodness of fit of these curves (e.g., R^2 value) will vary for different weighting parameters of the KPI. By varying the weighting parameters, the KPI values can be adjusted while the corresponding output measure values remain unchanged. Thus, when including the weighting parameters in the optimization, a better fitting curve can be should be found for the correlation between KPI and output measure. The process of optimizing the weighting parameters of a KPI for a set of data can be computationally costly, however, the procedure will only

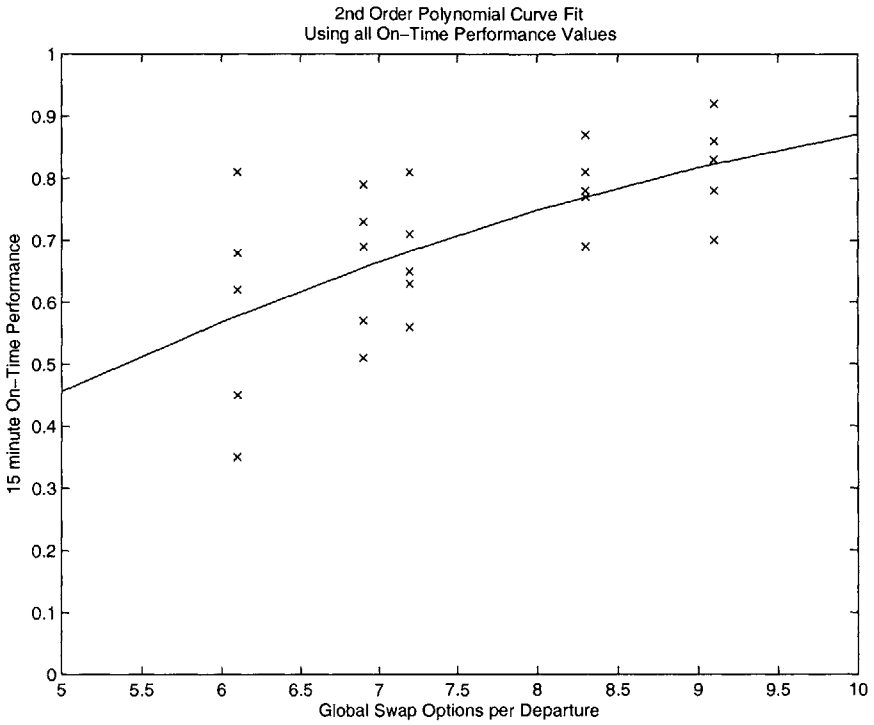

(a) curve fit using all OTP values

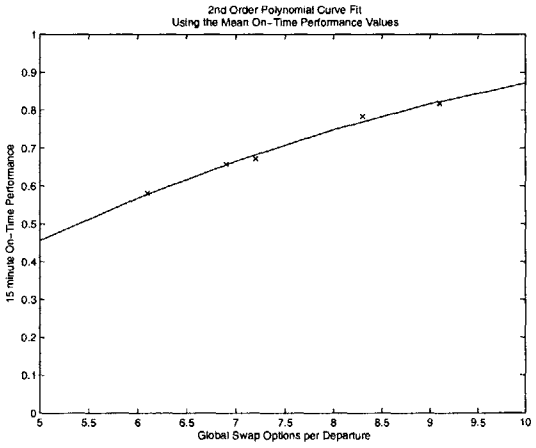

(b) curve fit using the OTP mean values

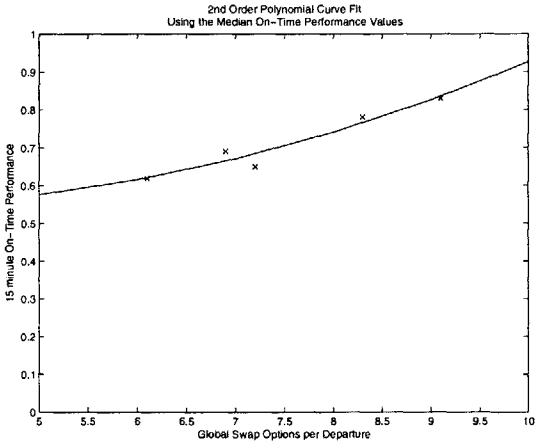

(c) curve fit using the OTP median values

Figure 6-1: Plot of Hypothetical Correlation between GSD and OTP

be required infrequently. Once a set of optimized weighting parameters is found, it can be used for the computation of the KPI until significant changes occur in the environment in which the schedules are operated. In practice, a re-optimization of the weighting parameters every month should be sufficient, because the environment of the operation does usually not change significantly in a shorter time frame.

Chapter 7

Directions for Future Research

Multiple extensions to the proposed simulation study can be carried out with additional data sources. Airport capacity profiles, weather, and other airline traffic are currently only modeled for the airports FRA and MUC. Additional airports can be modeled without great difficulty given the required data. While multiple weather scenarios are considered in the study presented in this thesis, only one airport traffic scenario for FRA and MUC is considered. Additional weather and airport traffic scenarios can increase the sample space for analysis.

Expansions to the simulation platform can increase its fidelity. Currently only aircraft are tracked and recovered. A valuable addition to the simulation would be the tracking and recovery of crew and passengers. A more detailed en route model considering en route winds and weather can capture correlation between the flight times of different flights that operate along similar flight paths at a similar time. These correlations are currently not considered in the simulation. Models for reduced aircraft dispatch reliability due to technical failure and for a maintenance system can increase the fidelity of aircraft availability in the simulation.

With the integration of crew and passenger tracking and recovery into the simulation, the analysis of KPIs that include crew and passenger data can be conducted. These KPIs can potentially cover the situation of the airline operation more completely than those analyzed in this thesis.

An analysis of different time frames for which the KPIs are evaluated can help

to identify the most meaningful KPI and time frame combinations. The KPIs and output measures analyzed in this thesis are aggregated over a day, however, the length of the time frame for the computation of a KPI can be varied.

Appendix A

Acronyms and Initialisms

ADDS	Aviation Digital Data Service
AGL	above ground level
ATC	Air Traffic Control
BOS	three letter IATA code for Boston Logan International airport
CDG	three letter IATA code for Paris Charles De Gaulle airport
DUS	three letter IATA code for Duesseldorf International airport
DFS	Deutsche Flugsicherung (German Air Traffic Control)
FIFO	First-In First-Out
FRA	three letter IATA code for Frankfurt Rhein-Main airport
GFCI	Global Fleet Composition Indicator
GSD	Global Single Swap Options per Departure
GTSD	Global Time Averaged Swap Options per Departure
HAM	three letter IATA code for Hamburg Fuhlsbuettel airport
IATA	International Air Transport Association
IFR	Instrument Flight Rules
IMC	Instrument Meteorological Conditions
IOCS	Integrated Operations Control System
KPI	Key Performance Indicator
LHR	three letter IATA code for London Heathrow airport
LowVis	Low Visibility conditions

MAN	three letter IATA code for Manchester Ringway airport
MEANS	MIT Extensible Air Network Simulation
METAR	Meteorological Terminal Area Reports
MinGT	Minimum Ground Time
MIT	Massachusetts Institute of Technology
MUC	three letter IATA code for Munich Franz Josef Strauss airport
NWS	National Weather Service
OCC	Operations Control Center
OTP	On-Time Performance
STR	three letter IATA code for Stuttgart Echterdingen airport
TDM	Total Delay Minutes
TFM	Traffic Flow Management
TXL	three letter IATA code for Berlin Tegel airport
USA	United States of America
VFR	Visual Flight Rules
VMC	Visual Meteorological Conditions
WBD	Average Weighted Buffers per Departure
WSD	Weighted Single Swap Options per Departure
WTSD	Weighted Time Averaged Swap Options per Departure

Appendix B

Derivation of Probability Density Function of Flight Time

$$\begin{aligned} D &\sim \text{Normal}(\mu, \sigma) \\ W &\sim \text{Exponential}(\lambda) \\ T &= D + W \end{aligned}$$

$$f_T(t) = \int_{x=-\infty}^{x=\infty} f_D(t-x) \times f_W(x) dx$$

Since $f_W(w) = 0$ for $w < 0$ the integration limits can be reduced to 0 and ∞.

$$\begin{aligned} f_T(t) &= \int_{x=0}^{x=\infty} f_D(z-x) \times f_W(x) dx \\ &= \int_{x=0}^{x=\infty} \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-(z-x-\mu)^2}{2\sigma^2}} \times \lambda e^{-\lambda x} dx \\ &= \frac{-\lambda}{2} e^{\mu\lambda + .5\sigma^2\lambda^2 - \lambda z} \left(-1 + erf\left(\frac{(-z+\mu+\lambda\sigma^2)\sqrt{2}}{2\sigma} \right) \right) \end{aligned}$$

Appendix C

Weather Scenarios

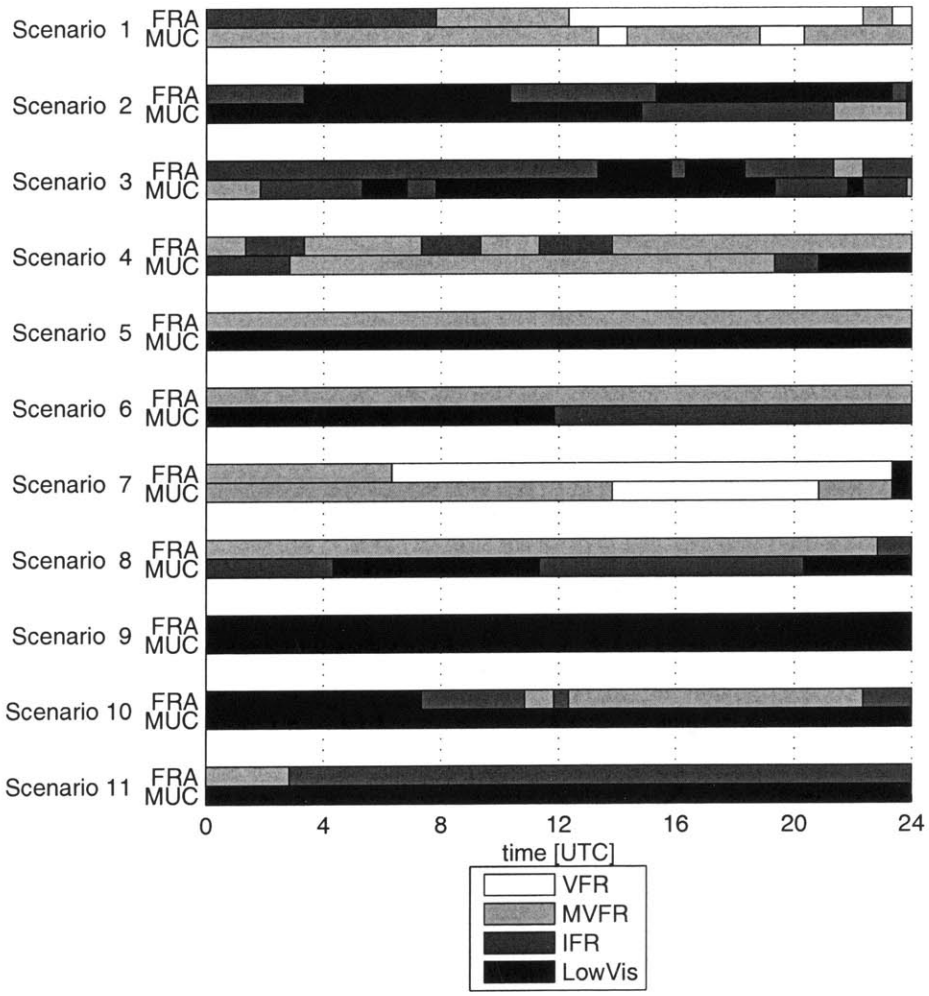

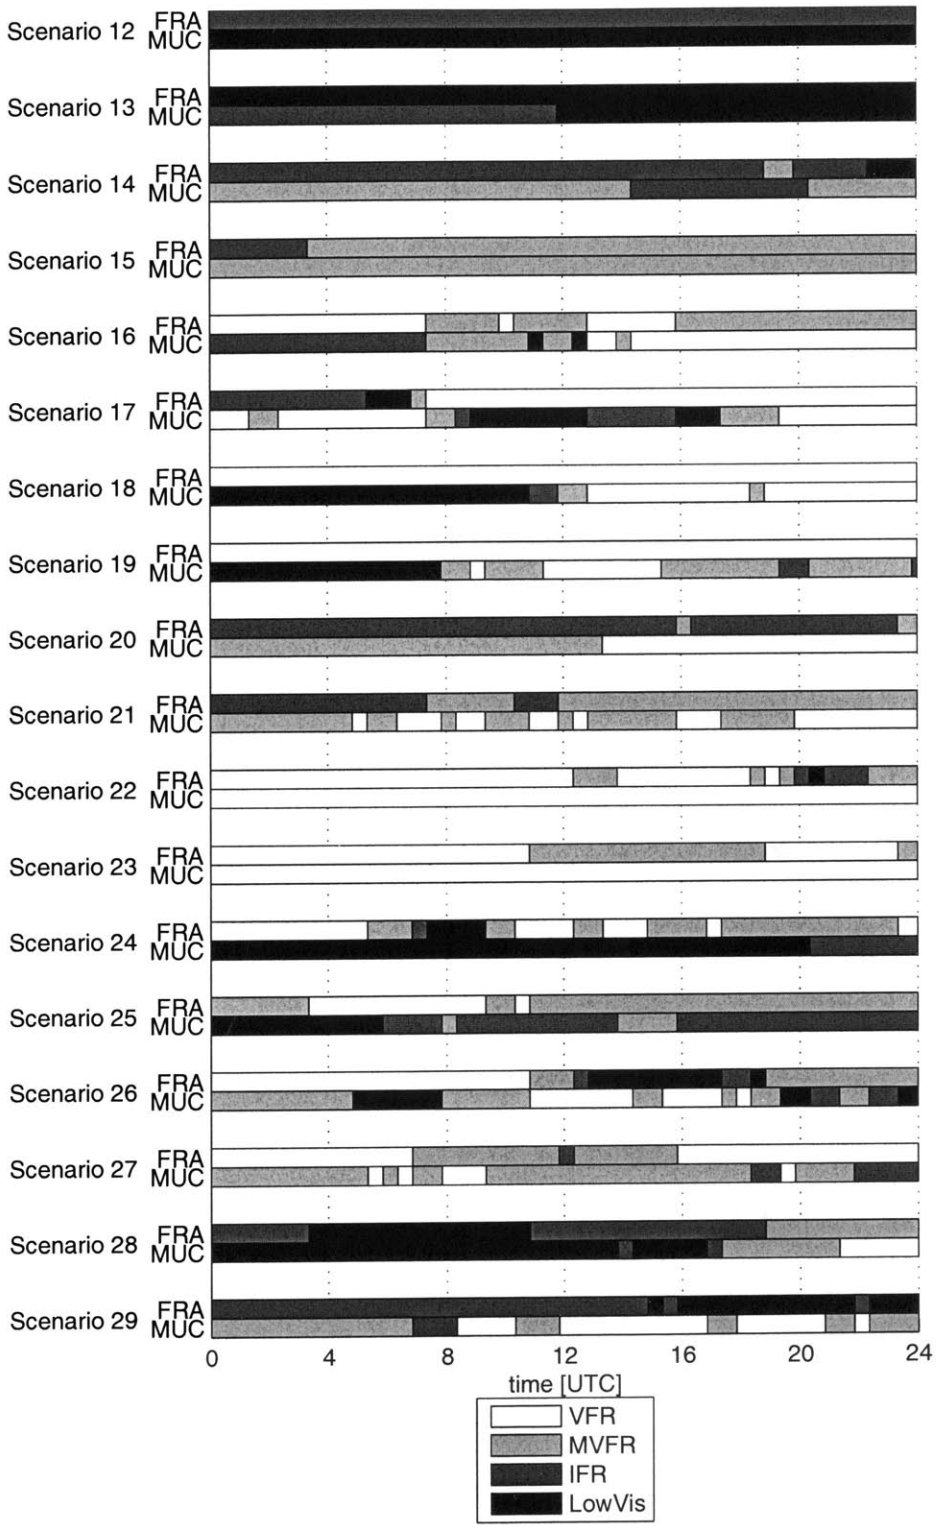
Scenario 12
Scenario 13
Scenario 14
Scenario 15
Scenario 16
Scenario 17
Scenario 18
Scenario 19
Scenario 20
Scenario 21
Scenario 22
Scenario 23
Scenario 24
Scenario 25
Scenario 26
Scenario 27
Scenario 28
Scenario 29
FRA
MUC
0
4
8
12
16
20
24
time [UTC]
VFR
MVFR
IFR
LowVis

Bibliography

[1] Carmen Systems AB. Carmen Integrated Operations Control - Product Sheet. http://www.carmen.se/air_products/pdf/air_integrated_operations_control.pdf, 2004. accessed on Feb. 02, 2006.

[2] J. Abara. Applying Integer Linear Programming to the Fleet Assignment Problem. *Interfaces*, 19(4):20–28, 1989.

[3] Deutsche Lufthansa AG. Zahlen, Daten, Fakten 2000/2001. Lufthansa Basis, 60546 Frankfurt am Main, Germany, 2001.

[4] Claudine Biova Agbokou. Robust Airline Schedule Planning: Review and Development of Optimization Approaches. Master's thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, September 2001.

[5] Yana Ageeva. Approaches to Incorporating Robustness into Airline Scheduling. Master's thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, August 2000.

[6] Massachusetts Port Authority. Airport Programs - Environmental - How Logan Operates. http://www.massport.com/logan/airpo_noise_howlo.html.

[7] Elisabeth Bly. Effects of Reduced IFR Arrival-Arrival Wake Vortex Separation Minima and Improved Runway Operations Sequencing on Flight Delay. Master's thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, February 2005.

[8] John Van de Vegte. *Feedback Control Systems.* Prentice Hall, Upper Saddle River, NY, 07458, USA, third edition, 1993. ISBN 0130163791.

[9] Institut du Transport Aerien. Cost of Air Transport Delay in Europe - Final Report, November 2000.

[10] Bob Rae et al. Ontario - a Leader in Learning, February 2005. ISBN 077947564X.

[11] Eurocontrol Central Office for Delay Analysis. Enhanced coda. `http://www.eurocontrol.int/eCoda/`.

[12] Steering Committee for the Review of Government Service Provision. *Report on Government Services 2005.* Commonwealth of Australia, LB 2 Collins Street East Post Office, Melbourne VIC 8003, Australia, 2005. ISBN 1740371607.

[13] Waitakere City Council Strategic Group. Key performance indicator council services survey. `http://www.waitakere.govt.nz/AbtCnl/kpisurveys.asp`. accessed on Jan. 15, 2006.

[14] John-Paul Clarke, Terran Melconian, Elizabeth Bly, and Fabi Rabbani. MEANS - MIT Extensible Air Network Simulation. *Simulation*, 81(12), December 2005.

[15] OAG Worldwide Limited. OAG MAX - System Version 6.1. Data CD-ROM.

[16] Terram Melconian. *MEANS User Documentation.* Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA, 02139, USA.

[17] Terran Melconian. Effects of Increased Nonstop Routing on Airline Cost and Profit. Master's thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, September 2001.

[18] Terran Melconian and Elizabeth Bly. *MEANS Programming Documentation.* Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA, 02139, USA.

[19] United Parcel Service of America Inc. Operating in Unison - 2004 UPS Corporate Sustainability Report. 55 Glenlake Parkway, Atlanta, GA, 30328, USA, June 2005.

[20] The Court of Justice of the European Communities. Judgment of the Court (Grand Chamber) in case c-344/04, the Queen on the application of International Air Transport Association, European Low Fares Airline Association v Department for Transport, January 2006.

[21] Curtin University of Technology. Annual Report 2004. GPO Box U1987 Perth, WA, 6845, Australia, 2005.

[22] Bureau of Transportation Statistics. Airline On-Time Statistics and Delay Causes. `http://www.transtats.bts.gov/OT_Delay/ot_delaycause1.asp`, October 2005.

[23] Katsuhiko Ogata. *Modern Control Engineering.* Prentice-Hall International, Inc, second edition, 1990. ISBN 0135987318.

[24] The European Parliament and the Council of the European Union. Regulation (EC) No 261/2004 of the European Parliament and of the Council of 11 February 2004 establishing common rules on compensation and assistance to passengers in the event of denied boarding and of cancellation or long delay of flights, and repealing Regulation (EEC) no 295/91 - commission statement, February 2004.

[25] Eric T. Peterson. *Web Site Measurement Hacks.* OReilly, 1005 Gravenstein Highway North, Sebastopol, CA, USA, 1 edition, 2005. ISBN 0596009887.

[26] Oxford University Press. *Oxford English Dictionary.* Oxford University Press, Oxford, United Kingdom, draft revision edition, December 2001.

[27] National Weather Service. ADDS - METARs Help Page. `http://adds.aviationweather.noaa.gov/metars/description_ifr.php`. accessed on Feb. 05, 2006.

[28] Georgetown University University Information Services. Data warehouse: Glossary. http://uis.georgetown.edu/departments/eets/dw/GLOSSARY0816.html. accessed on Feb. 05, 2006.

[29] Barry Craig Smith. Robust Airline Fleet Assignment. PhD thesis, Georgia Institute of Technology, Industrial and Systems Engineering, August 2004.

[30] Ashish Tewari. *Modern Control Design.* John Wiley & Sons Ltd, Baffins Lane, Chichester, West Sussex, PO19 1UD, United Kingdom, 2002. ISBN 0471496790.

[31] DFS Deutsche Flugsicherung GmbH Geschftbereich Tower (TWR/FP). *Studie zur Kapazität an den Flughäfen mit DFS-Flugplatzkontrolle.* Am DFS-Campus 10, 63225 Langen, Germany, 2003.

[32] Edith Cowan University. Annaual Report 2004. Joondalup Drive, Joondalup, WA, 6027, Australia, 2005.