Insights into the potential of robust capacity planning at KLM Aircraft Services

Master Thesis Industrial Engineering & Management

Jeroen A.J. Harmsen BSc.

Exam Committee

University of Twente
dr. ir. M.R.K. Mes
dr. ir. L.L.M. van der Wegen

KLM Aircraft Services – SPL/K3
ir. M.H.G. Bovenkerk

Amsterdam Airport Schiphol
November 21, 2012

Summary

This report is written for KLM Aircraft Services, the department of KLM which is responsible for moving and preparing aircraft for flight at the platforms, and gives insight into the possibilities of creating and the potential of using a more robust capacity planning.

Using the current deterministic way of planning, it is not possible for Aircraft Services to create an explicit link between covering of workload (the number of operators needed to perform the amount of work which needs to be done at a certain moment in time) and on-time performance. Resulting from this, the planning gives just one option for the workload, independent on the variability in the arrival punctuality and/or process duration. When we take uncertainty and unforeseen events into account in the planning methods of Aircraft Services, we take deviations from the fixed schedules and timetables into consideration and cope with possible disturbances. This leads to a robust way of capacity planning. A robust capacity planning takes possible delays, which can occur at an operational level, into account in the planning phase. For this reason, it works well in many situations, instead of perfect in just one situation.

This project has the following goal:

"To provide insight into the effect of uncertainty and unforeseen events on the dynamics of personnel capacity planning, related to the performance of KLM Aircraft Services."

To achieve the goal of this project, it is only necessary to analyse one of the processes of Aircraft Services in depth. The insights of this analysis are then extended to the other processes. To choose between the different processes of Aircraft Services, we first characterised them and afterwards made an assessment, where we decided to choose for the refuelling service. The most important characteristics to characterise the services of Aircraft Services in the scope of this research are:

- The type of process: *strict* arrival/departure (executed *exactly at* arrival or *at* departure) or arrival/departure *oriented* (executed *after* arrival and *before* departure).
- The flexibility of scheduling the process (dependency on other processes and whether the process has flexible permissible time windows).
- The duration of the process.
- The variation in the duration of the process.

We analyse the different elements affected by uncertainty and unforeseen events in the planning of KLM. A part of the variability in these elements is relatively predictable and a part is truly uncertain. The most important elements in the planning are the arrival punctuality of aircraft, the duration of the ground processes and possible disturbances. To describe the behaviour of these elements, we use theoretical probability distributions.

We create four models to get from a predefined flight schedule to the corresponding performance of the refuelling process, while incorporating uncertainty and unforeseen events in the planning. Our first model uses Monte Carlo Simulation to generate workload lines based on the combination of predefined theoretical probability distributions and the flight schedule. We optimise the starting time per task using our second model: an optimisation heuristic. Our third model uses Linear Programming to fit a set of shifts to the optimised workload, on different levels of workload coverage. We determine the corresponding performance by using our fourth model: a Discrete Event Simulation, which we use to simulate a day of operation of the refuelling process.

When we create a workload planning incorporating uncertainty and unforeseen events, we obtain a wide range of possible values the workload can take. From this, we conclude that uncertainty and unforeseen events have a big influence on the planned capacity. The eight hour shifts used for scheduling have a large impact on the robustness of the planning, since they cover, due to their duration, more than just the peaks in the workload, which leads to a lot of excess capacity. This allows peaks to move within a shift without large consequences for performance. The shifts thereby neutralize a large part of the discrepancy between planning and reality, caused by the arrival punctuality of flights. Simulating a day of operation results in high performances, but due to the assumptions made to be able to model the refueling process with the limited scheduling guidelines KLM could supply, our simulation model overestimates the actual performance.

We cannot extent the simulated performance directly to reality, but we are able to extent the insights about dealing with uncertainty and unforeseen events in the workload planning to the other services of Aircraft Services. This is possible for the Refuelling service, Water service, Toilet service, Aircraft Handling Support and the Pushback and Towing service. For *strict* arrival/departure processes the variation in the arrival/departure punctuality of aircraft is most relevant and for arrival/departure *oriented* processes the variation in the task process time is most relevant.

To deal with uncertainty, we recommend implementing an approach where Aircraft Services manages the current amount of uncertainty and simultaneously start initiatives to reduce this uncertainty. Finally Aircraft Services needs to buffer against the amount of uncertainty which cannot be reduced. To manage uncertainty, Aircraft Services needs to make a good forecast for the parameters of the defined probability distributions. These probability distributions have to be used in the process of creating workload profiles (instead of the current planning norms for tasks) to create an estimate of the range of values the workload can possibly take. Then KLM needs an integrated planning approach to make well founded decisions about which parts of the workload to cover, which shifts to use and how to schedule other activities (like breaks). This integrated planning approach is a process with multiple cycles of design, testing and adapting, which has to be executed by an expert with a lot of knowledge of all the processes.

Reducing uncertainty starts with executing further research to define the impact of all sources of uncertainty and afterwards taking steps to decrease the most significant sources. We recommend reducing uncertainty in an iterative way, alongside managing and buffering against uncertainty.

KLM has to buffer against the amount of uncertainty and unforeseen events which cannot be reduced. Buffering against uncertainty is according to the Law of Buffering automatically being done by some combination of inventory (or cycle time), capacity and time. Arrival/departure *oriented* processes (Refuelling, Water, Toilet) can buffer with all three options. For *strictly* arrival/departure processes (Pushback, AHS) buffering with time is not desirable. KLM has to remember that introducing buffers in the planning always increases the robustness; however this goes at the expense of optimality.

Preface

During my study I got more and more interested in the large scale movement of people and goods. This led me to choosing the master track Production and Logistics Management, with a focus on the technical side. For my master thesis, I looked for companies involved in the world wide transport of people and goods, which resulted in contacting KLM.

I would like to thank Hans Heerkens for providing me with contact details for KLM Aircraft Services, giving me the opportunity to contact them directly. My supervisors from the University of Twente were Martijn Mes and Leo van der Wegen. I would like to thank them for helping me to focus my project and advising me in clearly structuring my ideas to translate them to an accessible report. The meetings at the university where always pleasant and really helped me further.

At KLM Aircraft Services, I first want to thank Erwin Thijssen and Mark Bovenkerk for giving me the opportunity to do my graduation project here. My supervisor at KLM Aircraft Services was Mark Bovenkerk, who I would like to thank for his support during this project. He provided me with a lot of useful ideas, suggestions and feedback, helping me take this project to a higher level. I also would like to thank my roommates at KLM: Niek, Marian, Merlijn, Tony, Ir Wei and Tonnie for always being available to answer my questions and the pleasant working atmosphere in the room.

I experienced Schiphol as a very inspiring working environment and I really enjoyed my time at KLM.

Jeroen Harmsen
Amsterdam Airport Schiphol, November 21, 2012

Contents

1 Introduction

This report provides insight into the effect of uncertainty and unforeseen events on the dynamics of personnel capacity planning, and relates this to the performance of KLM Aircraft Services. In this chapter, we first start with an introduction into KLM. Afterwards we formulate our problem and define the goal of this project. We finish with a description of the scope and an outline of the report.

1.1 Background

In this section we give a brief company description of KLM Royal Dutch Airlines and the division were this research takes place: KLM Aircraft Services (AS). We also give a brief overview of the different services delivered by KLM Aircraft Services. These services are explained in detail in Chapter 2.

1.1.1 KLM

KLM Royal Dutch Airlines is a worldwide company based in the Netherlands. It was founded on October 7, 1919 and is the oldest airline in the world operating under its original name. KLM Royal Dutch Airlines is the core of KLM Group, which further includes KLM Cityhopper, Transavia.com and Martinair. In 2004, KLM merged with Air France. After this merge, Air France KLM became the world's second largest airline in terms of financial turnover.

The KLM Group operates a fleet of 211 modern aircraft, supported by over 33,000 employees in 2011. In the financial year 2010-2011, the KLM Group carried 23.1 million passengers and 491,000 tonnes of freight. At this moment, KLM (and its partners) offer 77 intercontinental and 80 European destinations directly from Schiphol. All of this generated an income of 8,651 million euro's.

KLM also performs maintenance and technical modifications on aircraft, engines and components for many different airlines. This is done by KLM Engineering & Maintenance.

1.1.2 KLM Aircraft Services

Schiphol Amsterdam Airport is the global hub and home base of KLM. At Schiphol, KLM has her own organization taking care of ground handling: KLM Ground Services (GS). Ground Services provides all services related to flight operations to the passenger, luggage and the aircraft (excluding the technical part). This research takes place at the department of Ground Services taking care of moving and preparing the aircraft for flight at the platforms: KLM Aircraft Services (AS). Aircraft Services delivers the services as described in Table 1.1 as service provider to KLM, her partners and third parties.

Service	Description
Airside Handling Support	Connecting the aviobridge to the aircraft, crew transport and crew briefings
Board supply	Changing and distributing the non-food supplies of the aircraft (pillows, blankets, etc.)
Catering distribution	Changing the catering supplies of the aircraft (unloading the old and loading the new supplies)
Cleaning	Cleaning the interior of the aircraft
De-icing	Remove ice from the aircraft (de-icing) and applying a fluid that prevents freezing (anti-icing)

Flex tasks	Cooling and heating the cabin, giving jet starts and docking of mobile staircases on buffers
Pushback	Pushing an aircraft back from the gate, since aircraft cannot taxi backwards by themselves
Refuelling	Refuelling the aircraft with a specific amount of fuel, based on the flight destination
Security check	Checking the interior of the aircraft for unsecured objects or unsafe situations
Toilet service	Emptying and flushing the toilet tanks of the aircraft
Towing	Moving aircraft between hangars, gates and buffer positions
Water service	Filling or refreshing the water supplies of the aircraft

Table 1.1: Services that are delivered by KLM Aircraft Services.

1.2 Project formulation

Aircraft Services defines for all her processes targets for on-time performance. The different department managers are responsible to comply with these performance targets in a cost efficient way within predefined windows dictated by the flight schedule.

In the current situation, Aircraft Services uses for all her processes a planning which does not account for uncertainty and unforeseen events (a detailed overview of the current planning processes is given in Chapter 3). This leads to a completely deterministic planning due to the following aspects:

- Flights arrive and leave exactly according to the times defined in the flight schedule.
- The planning does not take deviations from the planned process time of the ground processes into account. The planning is created based on static planning norms which indicate an average duration of the processes per aircraft type.
- The planning does not account for disturbances.

Resulting from this, the planning gives an overview of the planned workload, independent of the variability in the different arrival and/or process times.

- *We define the workload as the number of (fuel, water, cleaning etc.) operators needed to perform the amount of work which needs to be done at a certain moment in time.*

- *We define a workload profile as an overview of the workload over the day.*

The current deterministic way of planning leads to just one line indicating the workload over a day. We show this in the example workload profile of Figure 1.1. KLM Aircraft Services uses this kind of workload profiles for planning her processes. In reality, a workload profile should not consist of just one line. The workload can vary a lot due to uncertainty and unforeseen events in the arrival and ground processes.

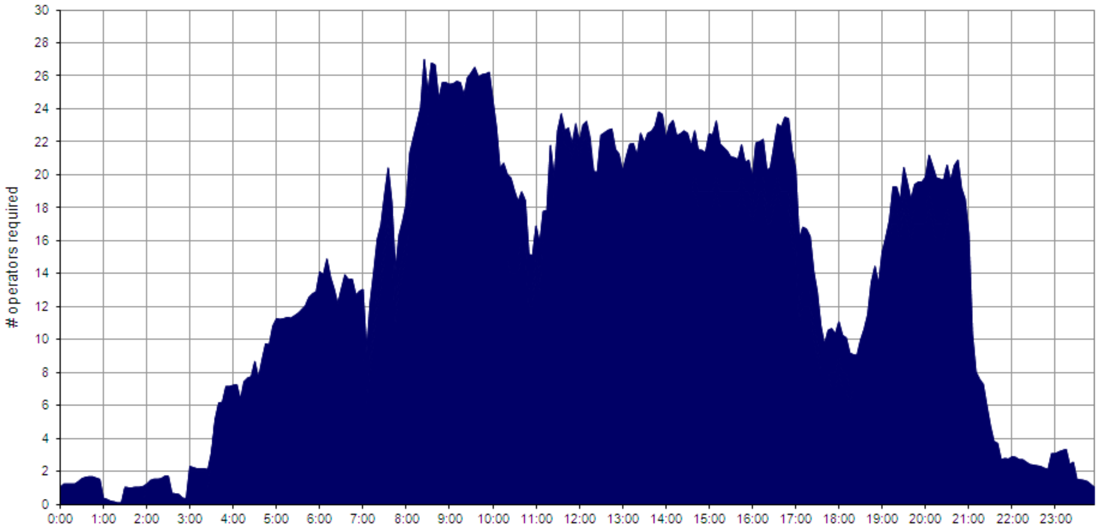

Figure 1.1: An example of a workload profile, indicating the required number of operators over the day.

Using the current deterministic way of planning, it is not possible to create an explicit link between covering of workload and on-time performance. KLM defines performance on the flight level and on the task level. On the flight level, this is the percentage of flights departing on time i.e., at the scheduled time of departure as defined in the flight schedule. On the task level, which is most relevant for this research, KLM defines performance as the percentage of tasks finished before their latest end times as defined in the norms (see Appendix C). These norms are defined per task type and per aircraft type. We explain the norms in depth in Chapter 3.

At the day of operation, the capacity planning of KLM is executed with the assumptions that no disturbances occur. According to Gao et al (2009), inflexible airline planning results in high operational costs, because in reality delays and disruptions occur on a regular basis. These are mostly caused by adverse weather, mechanical failures, air traffic control and other unavoidable reasons. This all leads to uncertainty in the planning and to the following goal for this research:

To provide insight into the effect of uncertainty and unforeseen events on the dynamics of personnel capacity planning, related to the performance of KLM Aircraft Services.

Galbraith (1973) defines uncertainty as: *"the difference between the amount of information required to perform a task and the amount of information already possessed."* Unforeseen events are according to the Oxford Dictionaries defined as: *"events which are not anticipated or predicted"*. When we take uncertainty and unforeseen events into account in the planning methods we take deviations from the fixed schedules and timetables into consideration and cope with possible disturbances.

Uncertainty and unforeseen events cause variability in the planning. The Law of Variability (Hopp and Spearman, 2008) says that increasing variability always degrades the performance of a production system. How performance degrades, is defined by Hopp and Spearman (2008) in the Law of Buffering, defining that variability in a production system will always be buffered by some combination of inventory (or cycle time), capacity and time (long lead times).

To determine the required capacity, taking the considerations mentioned into account, Cacchiani and Toth (2012) define a nominal and a robust way of capacity planning and make a clear distinction between these two based on a train timetabling case. The nominal version of the problem means optimising such that all the constraints are satisfied and an objective function is optimised (for example minimising the total costs). One can think of the traditional approaches of sensitivity analysis and stochastic linear programming (Mulvey and Vanderbei, 1995). These nominal approaches are based on maximising efficiency. However, when a delay occurs, it is common that an optimal timetable becomes sub-optimal or even infeasible. For this reason, it is important to create timetables that are able to absorb delays and avoid delay propagation (Cacchiani and Toth, 2012). This leads to a robust way of planning: a robust planning takes possible delays, which can occur at an operational level, into account in the planning phase. For this reason, it works well in many situations, instead of perfect in just one situation.

A robust optimisation model incorporates, according to Mulvey and Vanderbei (1995), two different objectives of robustness. They define a model to be “solution robust” when it gives a solution which remains close to the optimal solution for all scenarios of input data. A model can also be “model robust” which means that it remains feasible for all data scenarios. Robust optimization models are nowadays developed for a lot of applications, for example: power capacity expansion, image reconstruction, airline scheduling and scenario immunization for financial planning (Mulvey and Vanderbei, 1995).

We now focus on the capacity planning of Aircraft Services. Aircraft Services needs to strive for a solution robust capacity planning, to remain as close as possible to the optimal solution, even when using uncertain input data. It is important for Aircraft Services to work with a robust capacity planning, but on the other hand is it important to take efficiency (the nominal version of capacity planning) into account, because one should not be over-conservative in protecting against delays.

1.3 Research questions

To achieve the goal of this research, we formulate the following six research questions:

1. How can the different processes of KLM Aircraft Services be characterised?
2. What is the current way of planning at KLM Aircraft Services?

By answering question 1 and 2, we give an overview of the current situation in terms of processes and planning methods at Aircraft Services. We observe the processes at the aircraft on the platforms, perform interviews with process analysts, resource planners and managers at Aircraft Services and study earlier research executed at Aircraft Services. From this, we describe the processes and the current way of planning qualitatively.

3. What are the main elements affected by uncertainty and unforeseen events in the planning of KLM and KLM Aircraft Services?

To answer question 3, we perform quantitative analyses where we analyse historical data to get insight into the dynamics of the arrival process of aircraft and the refuelling process. Based on these analyses, we give an indication of the importance of different sources of uncertainty and unforeseen events and how they affect the main elements the planning is based on. Since KLM keeps an extensive database which contains a lot of data about for example arrival times and process starting and ending times, the magnitude of the mentioned uncertainty and unforeseen events can be determined. When we determine this magnitude, the uncertainty can be predicted and thereby potentially reduced.

Aircraft Services believes that a planning methodology, which explicitly uncertainty and unforeseen events takes into account, can give better insights in the required capacity and thereby lead to more effective steering on performance. We investigate this by answering question 4.

4. What performance can be expected when a robust capacity planning is used at the aircraft refuelling process?

To answer question 4, we create a model which creates workload profiles with arrival and process times based on probability distributions instead of deterministic norms. This model uses Monte-Carlo simulation (see Halton, 1970) to create workload profiles based on these probability distributions and the flight schedule. The Monte-Carlo simulation than gives a whole range of workload lines, indicating the effects of uncertainty and unforeseen events. The management of Aircraft Services can now choose which line to follow, based on the expected covering of workload.

To determine the expected performance based on certain coverages of workload, we create a discrete-event simulation model using Siemens Technomatix Plant Simulation. We model the arrival and departure process of aircraft and the refuelling process. Using this model, we analyse the difference in performance between the current situation and a situation where we do account for uncertainty and unforeseen events in the capacity planning.

The reason that we use a simulation can be explained using the book of Law (2007). Law states that when one can analyse a system analytically, then this should be done because it gives exact results. Simulation should be used for analysing systems that are too complex to be evaluated analytically. Simulation gives an estimate of the true characteristics of a system, not an exact solution. Law mentions that designing and operating transportation systems, such as airports, is a good example of a system for which simulation has been found to be a useful and powerful tool. Because the process of aircraft refuelling is highly complex due to the uncertain nature of events and the dependency on many factors outside the influence of KLM, we cannot analyse this system analytically and thereby simulation is a good choice.

5. How and to what extent can robust capacity planning be extended to the other ground processes of KLM Aircraft Services?

By answering question 5, we extent the insights of analysing the refuelling process to the other processes of KLM Aircraft Services.

6. What is the best way to implement this robust way of capacity planning at KLM Aircraft Services?

By answering question 6, we give recommendations for implementing the results of this research into the planning process of KLM Aircraft Services to make it more effective and efficient and to give management the ability to make well-founded choices in covering the expected workload.

1.4 Scope

The planning of the services executed by Aircraft Services is quite complex. Thereby, to make this project executable within the set time span of six months, we define a scope to set boundaries to the extent of this research. By setting this scope correctly, we ensure that the goals of this project are preserved.

An important aspect to notice is that we do not strive to create a complete new planning method in this research. We use the current planning methods and adapt the input variables by adding uncertainty and unforeseen events to provide insights on how this influences the performance of KLM Aircraft Services.

1.4.1 Planning level

For positioning this research, we define hierarchical levels of control according to the classic definition of management control from the framework of Anthony (1965). Hans (2001) gives an overview of this framework from which we can distinguish three levels of management control, namely: strategic, tactical and operational.

The strategic decision level involves decisions on the long-term planning horizon made by senior management. The basic function is setting long-term company goals to establish an environment capable to meet the overall goals of the company. The tactical decision level involves decisions on the medium-range planning horizon concerned with allocating resources such as workforce and equipment to meet the targets set on the strategic level. The operational decision level is concerned with the short-time scheduling. On this level, the operator and equipment capacity levels are fixed, only task are sequenced and coupled to operators and equipment (Hans, 2001).

In this research, we consider decision-making at the tactical decision level, because we focus on the allocation of resources to meet the strategic targets. The planning processes of Aircraft Services at the tactical decision level can be expressed in time from six months before the day of execution till one day before the day of execution. Because of the fixed horizon, we are automatically talking about offline planning.

In this research, we account for calculations made for a short horizon (about six months). Everything with a longer horizon is out of scope, because this is covered at the strategic decision level and thereby outside our direct influence. We do also not account for the operational decision level because the tasks at this level are executed by lower level managers at the different executing departments of Aircraft Services. One should keep in mind that the focus on the tactical decision level does not imply that we can ignore the other two decision levels completely. These other decision levels are considered fixed and non-adjustable.

1.4.2 Services

To reach the goals of this project and to keep the project within the predefined time limit of six months, we first investigate the effect of uncertainty and unforeseen events on the planning of just one service. We than translate the lessons learnt from studying this process to the other services by using the characteristics as we describe in Chapter 2. In Chapter 2, we conclude that aircraft refuelling is in this case the best service to choose. An argumentation for this can be found at the end of Chapter 2.

1.5 Outline

In this section we link the research questions of this project to specific chapters in the report.

To answer research question 1, we describe the different processes of KLM Aircraft Services in Chapter 2, where we start with describing the arrival and departure process from the view of Aircraft Services and afterwards continue with characterizing the different processes.

We answer research question 2 in Chapter 3, where we analyse the current planning process and give a description of the planning norms and methods.

In Chapter 4 we describe the sources of uncertainty and unforeseen events and the current situation in numbers, to answer research question 3. Finally we fit probability distributions to the arrival punctuality of aircraft, to the duration of disturbances and to the process times of the refuelling process.

In Chapter 5, we create models which we use to incorporate the effects of uncertainty and unforeseen events in the planning process to answer research question 4. We then continue with determining the effects of a more robust capacity planning on the performance of the refuelling process by simulating a day of operation.

We answer research question 5 in Chapter 6, where we generalise the results of the analysis from Chapter 5 to the other processes of KLM Aircraft Services. We make an assessment, based on the process characteristics as defined in Chapter 2, for which processes we can apply the insights of Chapter 5.

Research question 6 concerns the implementation of the results of this research and will be treated in Chapter 7. This chapter gives recommendation for what KLM should change in their (planning) processes to deal with uncertainty and unforeseen events and to get to a more robust capacity planning.

Finally, we end with conclusions and recommendations for further research in Chapter 8.

2 Characteristics of processes and services

This chapter gives a description of the aircraft turnaround process and then zooms in on the characteristics of services performed by AS. At the end of this chapter, we draw conclusions about why we choose the refuelling process as the primary process to investigate.

2.1 The arrival and departure process

All services of AS need to be executed between the arrival and the departure of an aircraft. The time between these two processes is called the ground time. The ground time depends on the flight schedule, which defines the time between an inbound and an outbound flight executed by the same aircraft.

Since KLM uses Schiphol as her hub-airport, KLM has to deal with arriving and departing passengers, but also with transfer passengers who only use Schiphol to transfer from one aircraft to another. This makes the processes at the airport a lot more complex, since passenger (PAX) and baggage flows (BAX) from local check-ins and different incoming aircraft have to be matched and merged in time at a departing aircraft (Schiebaan, 2002). This process is illustrated in Figure 2.1.

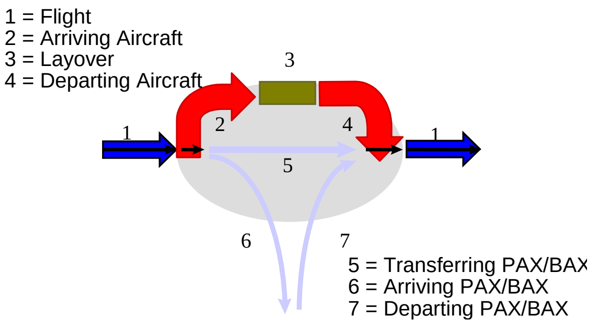

Figure 2.1: The main hub processes (Schiebaan, 2002).

Process 1 is the incoming and departing flight. This process ends when the doors of the aircraft open en starts again when they are closed for departure. The processes 2, 3 and 4 are together called a turnaround: the time during which an aircraft stands at the airport and is being made ready for its next departure. These processes stand for arriving (2), departing (4) and the layover (3). During the lay-over, an aircraft is available for maintenance. Process 5, 6 and 7 refer to loading and unloading the aircraft with passengers (PAX) and baggage and cargo (BAX). PAX and BAX in Process 5 are transferring and PAX and BAX in Process 6 and 7 have Schiphol Amsterdam Airport as origin or destination (Schiebaan, 2002). Most services of AS are executed during process 2 and 4, except de-icing and pushback, because these services are executed at departure and thereby part of the building block flight.

Multiple forms of turnarounds are possible, depending on the duration, the work that needs to be done and the place of processing. Below, the different forms of turnaround are illustrated using some schematics and a brief description, partially based on the research of Schiebaan (2002). These are an extension to the building blocks of Figure 2.1.

- *Short turnaround*

During a short turnaround, the arrival and departure processes are combined into one process, where Process 2 and 4 are executed in parallel. The permissible ground time of an aircraft is very close to the sum of the norm-times per process as defined in the Ground Operations Manual Schiphol (GOMS). We explain these norms in depth in Chapter 3. Sometimes a quick turnaround is possible; then the ground time is shorter than defined in the norms.

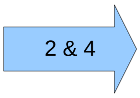

Figure 2.2: Short turnaround.

- *Long turnaround on gate position*

When a turnaround lasts significantly longer than the minimum norm ground times, this is a long turnaround (mostly several hours). This means that the services of AS can be planned at a suitable moment, within the permissible ground time.

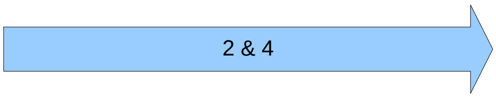

Figure 2.3: Long turnaround on gate position.

- *Long turnaround partly on buffer position, partly on gate position*

When an aircraft has a long ground time, it cannot always stand at the gate for its entire layover. When the gate needs to be cleared for the next aircraft, the first aircraft needs to be towed to a buffer. The services of AS can also be performed at a buffer, however not during towing.

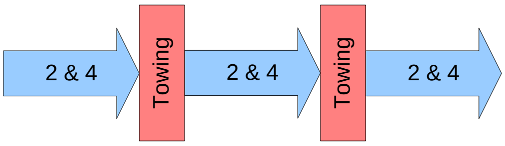

Figure 2.4: Long turnaround on buffer.

- *Long turnaround at hangar*

In some cases, the Engineering & Maintenance department (E&M) requests an aircraft at the hangar for technical service. During a long turnaround, the aircraft is then towed to the hangar and is not available for any of the services of AS during towing and technical service.

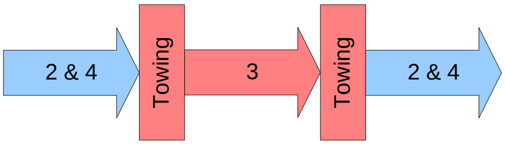

Figure 2.5: Long turnaround at hangar.

2.2 Services provided by Aircraft Services

AS provides twelve services to prepare aircraft for flight at the platforms and to tow them to other locations at Schiphol. Figure 2.6 zooms in on Figure 2.1 to show the predecessor/successor relations for the different processes. This figure shows that there exists a strict sequence in some processes, that some processes can be executed in parallel and that some processes cannot be executed simultaneously (indicated by two arrows between those processes). In this figure, towing is indicated as an optional process, because this is only done when an aircraft needs to be moved to a different gate, a buffer, or a hangar. We illustrate the different services provided by AS in blue.

Two types of turnaround processes are distinguished, namely the **strictly** arrival or departure processes and the arrival or departure **oriented** processes.

- *Strictly* arrival or departure processes are executed exactly *at* arrival or *at* departure. These processes are part of the flight process (Figure 2.1, Process 1). Examples are boarding, de-boarding, pushback and de-icing as shown in the 'Flight' boxes in Figure 2.6.
- Arrival or departure *oriented* processes are executed *after* arrival and *before* departure. This are processes to prepare the aircraft for flight and executed in Process 2 and 4 in Figure 2.1. Examples can be found in the 'Arriving aircraft' and 'Departing aircraft' boxes in Figure 2.6.

Most of the services provided by AS are arrival or departure *oriented* processes. Arrival or departure *oriented* processes can level their workload by using longer ground times (within the permissible time windows) to advance or delay processes. *Strict* arrival or departure processes are unable to level their workload, because they have to be executed at a fixed moment in time (Schiebaan, 2002).

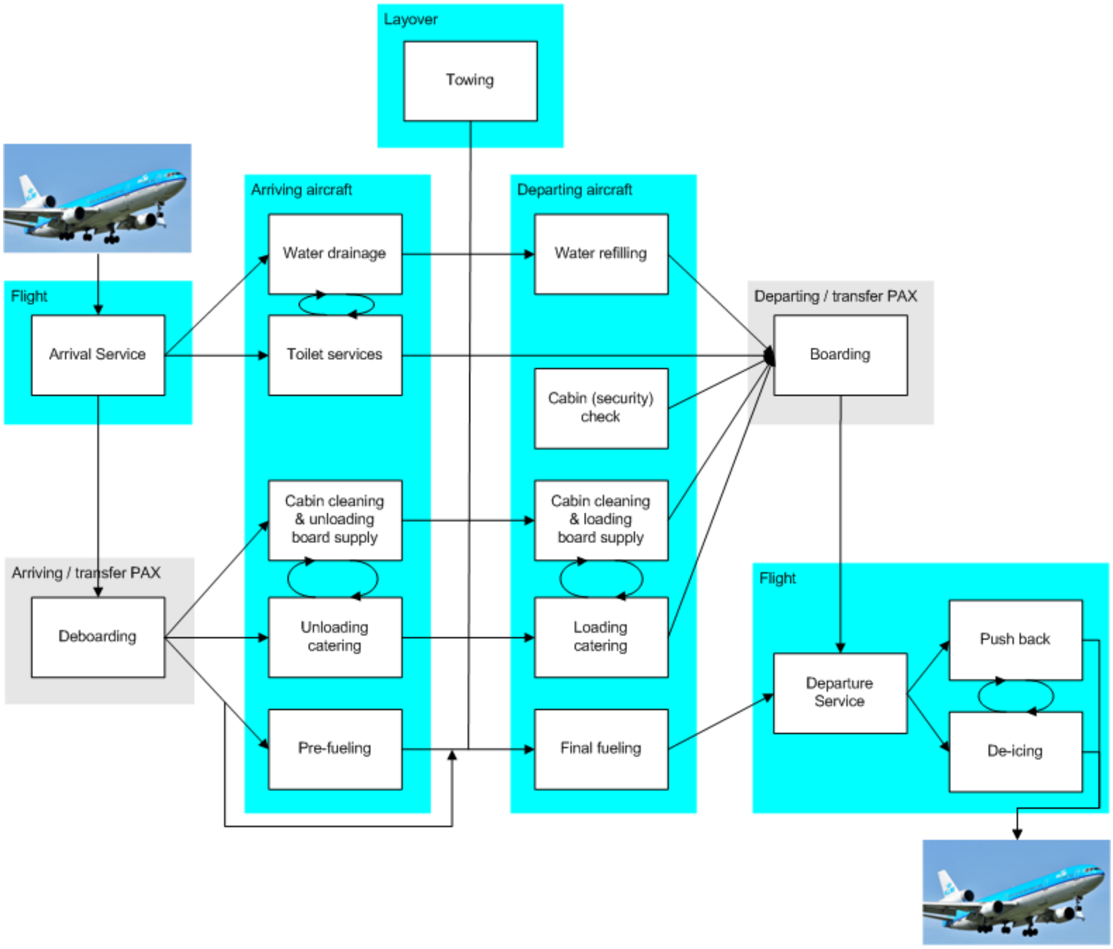

Figure 2.6: The services of AS (blue) positioned between the flight processes

The flex tasks are missing in Figure 2.6, because these are relatively small, and less important for the total picture, so they are kept out to keep the diagram readable.

2.2.1 Description of services

We give in this section a detailed description of the relevant parts of the twelve different services provided by AS.

Airside handling support
Airside handling support (AHS) connects the aviobridge to the aircraft. AHS also transports the cockpit and cabin crew to the aircraft (if the aircraft is placed at a buffer) and facilitates crew briefings. This is the arrival and departure service in Figure 2.6.

Board supply
Board supply is the service which distributes the non-food supplies to the aircraft (pillows, blankets, etc.). This service is executed by the cleaning companies (Asito and Klüh) after they finished cleaning. The supplies are being brought on board by KLM Catering Services (KCS). This means that the service is almost entirely outsourced.

Catering distribution
The catering service is outsourced to a daughter company of KLM: KLM Catering Services (KCS). This is a different company, however it is fully owned by KLM. KCS takes care of the supply, preparation and distribution of food and non-food supplies to the aircraft. It also takes care of loading the board supplies to the aircraft.

Cleaning
The cleaning service is outsourced to cleaning companies Asito and Klüh. They work autonomously, but are supervised by KLM contract managers for quality control. After cleaning is finished, the cleaning personnel takes care of the distribution of board supplies over the aircraft.

De-icing
KLM De-Icing Services takes care of removing ice from the aircraft (de-icing) and applies a fluid on the aircraft to prevent freezing (anti-icing). De-icing can be done at the gate or at remote de-icing positions. Remote de-icing is the safest, fastest and most efficient option. An aircraft arrives here with the engines running and is after de-icing immediately ready to depart. De-icing is only relevant in the winter period: from the second half of October till the end of March.

Flex tasks
The flex tasks service takes care of a lot of different processes like: cooling (in summer) and heating (in winter) the cabins by connection mobile coolers/heaters. This service also gives jet starts when a jet engine is unable to start on its own. The last task executed is the docking of mobile staircases to aircraft on buffers or on gates without an aviobridge.

Pushback
The pushback service is delivered by KLM Aircraft Towing & Push-Back Services and takes care of pushing an aircraft back from the gate. This service is needed because an aircraft cannot taxi backwards by itself. Pushback is done by tugs, the same ones as used for towing. The pushback tug places the aircraft on the taxiway in the right direction, from where the aircraft drives to the runway under its own power.

Refuelling
The refuelling service, executed by KLM Aircraft Refuelling, takes care of refuelling the aircraft with a specific amount of fuel which is based on the flight destination, the expected

weather underway and the total weight of the aircraft. At peak times, mostly it is too busy to serve all aircraft for the required amount of fuel. For this reason, several aircraft are pre-fuelled, to take in the minimum amount of fuel needed given the conditions from above. Shortly before departure, when the exact amount of fuel is known, a final fuel is executed (which is usually a small amount). In normal situations, only aircraft used for intercontinental flights are being pre-fuelled, because the ground time for European flight is, taking into account the additional set-up time, considered too short.

At a safe distance from the main terminal, the aircraft refuelling department has an airport tank farm from where the fleet of modern bowsers (refuelling trucks) and dispenser trucks are being deployed. The fleet consists of 3 large bowsers (80 m^3) and 15 smaller bowsers (40 m^3). These bowsers are used to refuel aircraft on remote stands and gates without a hydrant fuelling system.

Most gates on Schiphol are connected to a hydrant fuelling system. This is an underground network of fuel types. For refuelling aircraft at these gates only a dispenser is necessary, connecting the aircraft to the hydrant system. This is much faster, because no bowsers are needed who need to refill their tanks at the tank farm (for intercontinental flights this is needed multiple times per aircraft).The aircraft refuelling department has 21 dispensers available.

The refuelling department pumps no less than 2.5 million m^3 of jet fuel into more than 120.000 aircraft each year. Since 2010, the aircraft refuelling is done in close cooperation with Shell Aviation. This means that the equipment of KLM Aircraft Refuelling also services more than a dozen Shell customers.

Security check
The security check is the service which checks the interior of aircraft for unsecured object or unsafe situations.

Toilet service
The toilet service is part of KLM Aqua Services and is responsible for emptying and flushing the toilet tank of the aircraft. To ensure maximum hygiene, toilet services are kept strictly separate from the water services (except for the biggest planes).

Towing
The towing service is responsible for moving aircraft between gates, to and from buffers and to and from hangars. KLM Aircraft Towing & Push-Back Services delivers the towing service using tugs, the same ones as used for pushback. KLM Aircraft Towing & Push-Back Services equips 39 one-man tugs and 8 conventional vehicles. The drivers are fully qualified to perform part of the pre-flight inspection, which means airlines do not have to send in a ground engineer for this work. Pushback has in most cases priority over towing, but AS has to keep in mind that towing needs to deliver aircraft on-time at the gate and has contracts with the Engineering & Maintenance department for on-time delivery.

When an aircraft needs to be towed, three situations can occur. In the first situation, an aircraft arrives and departs from a different gate. When the ground time is very short, there is no need to use a buffer so it is possible to tow the aircraft directly from the arrival gate to the departure gate: gate-to-gate towing.

The second situation is called gate-buffer-gate towing and is used when an aircraft has a longer ground time then the time it can be kept at the gate (because the gate needs to be used by other aircraft). When this case occurs, the aircraft is towed to a buffer. When it has to get ready for departure, the aircraft is towed back to the same or a different gate.

In the third situation the aircraft needs to be towed to a hangar at the KLM Engineering & Maintenance (E&M) department, situated at Schiphol East. This is quite a long drive, and thereby these operations are planned outside the peak hours.

Water service
The water service is part of KLM Aqua Services and is responsible for filling and/or refreshing the drinking water supplies of the aircraft. To ensure maximum hygiene, the water services are kept strictly separate from the toilet services (except for the biggest planes).

2.2.2 Characteristics of services

The twelve services provided by AS all have different characteristics which we display in Table 2.1. The list of characteristics is composed based on the work of Schiebaan (2002) and Dekker (2010) and checked and finalised based on discussions with process analysts of AS.

	Airside handling support	Board supply	Catering distribution	Cleaning	De-icing	Flex tasks	Pushback	Refuelling	Security check	Toilet service	Towing	Water service
Arrival/departure oriented process		x	x	x		x		x	x	x	x	x
Strict arrival/departure process	x				x		x					
Aircraft type dependent	x			x			x	x	x	x	x	x
Destination dependent		x	x									
Gate allocation dependent	x				x	x	x	x			x	
Weather dependent					x	x		x				
Flexibility in permissible time windows								x	x	x		x
No flexibility in permissible time windows	x	x	x	x	x	x	x				x	
Sequence dependent	x	x	x	x	x	x	x		x		x	
(relative) Sequence independent								x		x		x
Quantity dependent (# passengers)		x	x									
Quantity dependent (distance)			x					x				
No other processes can be executed simultaneously							x				x	
Heavily connected with board supply and catering distribution				x								

Heavily connected with board supply and cleaning			x									
Heavily connected with catering distribution and cleaning		x										
Is not allowed to be executed at the same moment as toilet service (except at the B747 and the MD11)												x
Is not allowed to be executed at the same moment as water service (except at the B747 and the MD11)										x		

Table 2.1: Characteristics of the services provided by AS.

2.2.3 Process times of services

To give an indication of the duration of the different services, compared to the available ground times, we create Figure 2.7. In this figure, we divide, for all aircraft operated by KLM, the norm duration of all services by the norm ground times. Then we take the average over all aircraft types, per service. All these norm times are documented in GOMS. We give an example of the planning norms for the refuelling process in Appendix C.

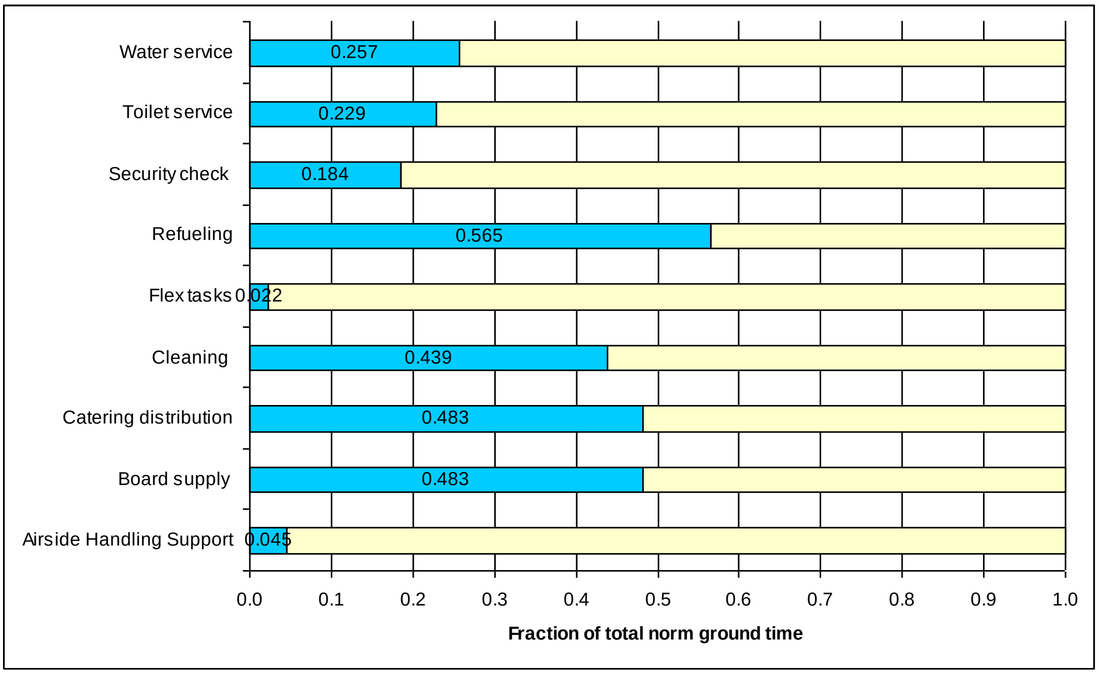

Figure 2.7: The average norm process times as fraction from the total norm ground times.

A few remarks have to be made to Figure 2.7. First, the fractions do not add up to one. This is because the mentioned processes are executed for a major part of the time in parallel. Second, we use only the arrival or departure oriented processes, because these processes are executed within the window set by the norm ground times. Third, the norm times we take to create this figure are from a short turnaround. A short turnaround is in Section 2.1 defined as a turnaround entirely on one gate, so there is no towing. For this reason, we do not include the towing process.

From Figure 2.7 we conclude that the refuelling service has on average the longest norm process time compared to the norm ground times. The fractions of the catering distribution and the board supply are the same, because they are heavily interconnected. This also accounts partially for the cleaning process, which has also a relatively strong link with catering distribution and board supply.

2.2.4 Variation in process times of services

To be thorough, we have to analyse the norm duration of the different services and the variation within the process times. The Oxford Dictionaries define variation as*: "a change or slight difference in condition, amount, or level, typically within certain limits"*. The variation from the average process times is defined by the standard deviation. A low standard deviation indicates that there is almost no variation and a high standard deviation indicates high variation. The standard deviations of the different services are illustrated in Figure 2.8.

The catering distribution and the board supply are not in Figure 2.8, because they are executed by KCS, which is a different company and does not use CHIP (the communication and scheduling program which also logs all data; in Dutch: Communicate & Hub Indelings Programma). This means that the actual data is not directly available to us for analysing.

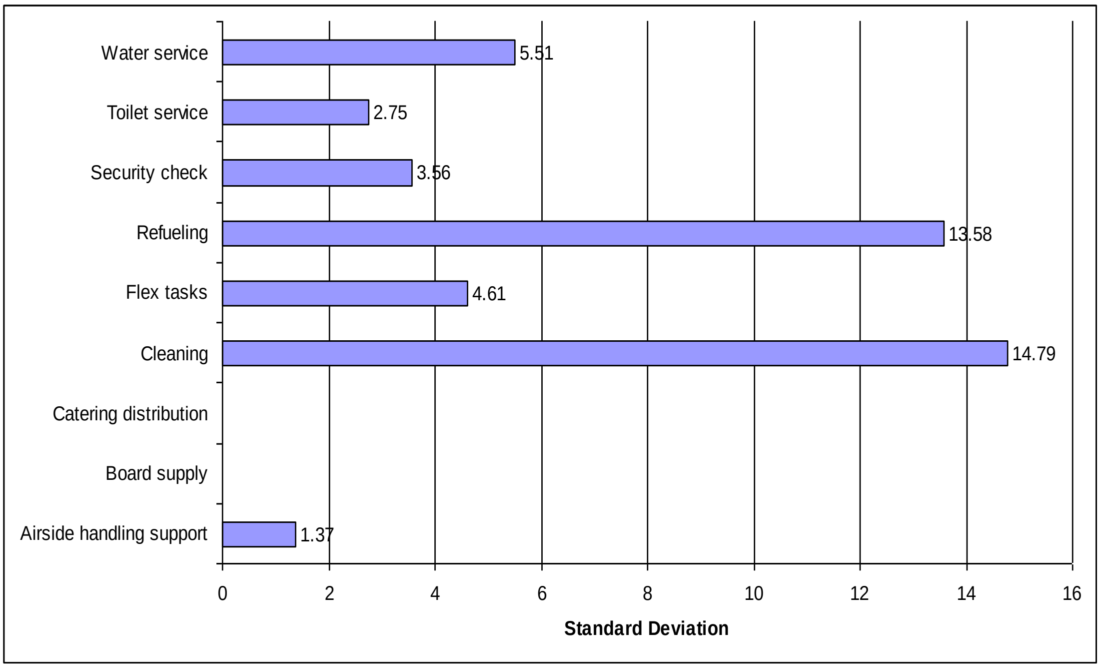

Figure 2.8: The standard deviation in the duration of the different services.

We conclude from Figure 2.8 that the refuelling and the cleaning service have the biggest standard deviations and thereby they have the biggest variation in their process times. This means that these services are the most likely services to cause delays. One should keep in mind that the cleaning service is being outsourced, so this is not under direct control of AS and thereby not suitable for us to analyse is depth.

To give an indication of the magnitude of the standard deviation, we give in Figure 2.9 an overview of the coefficient of variation per service. The coefficient of variation is a scale free variable and thereby in this case independent of the duration of a process. In statistics, the coefficient of variation is defined as the standard deviation divided by the mean (Bedeian and Mossholder, 2000).

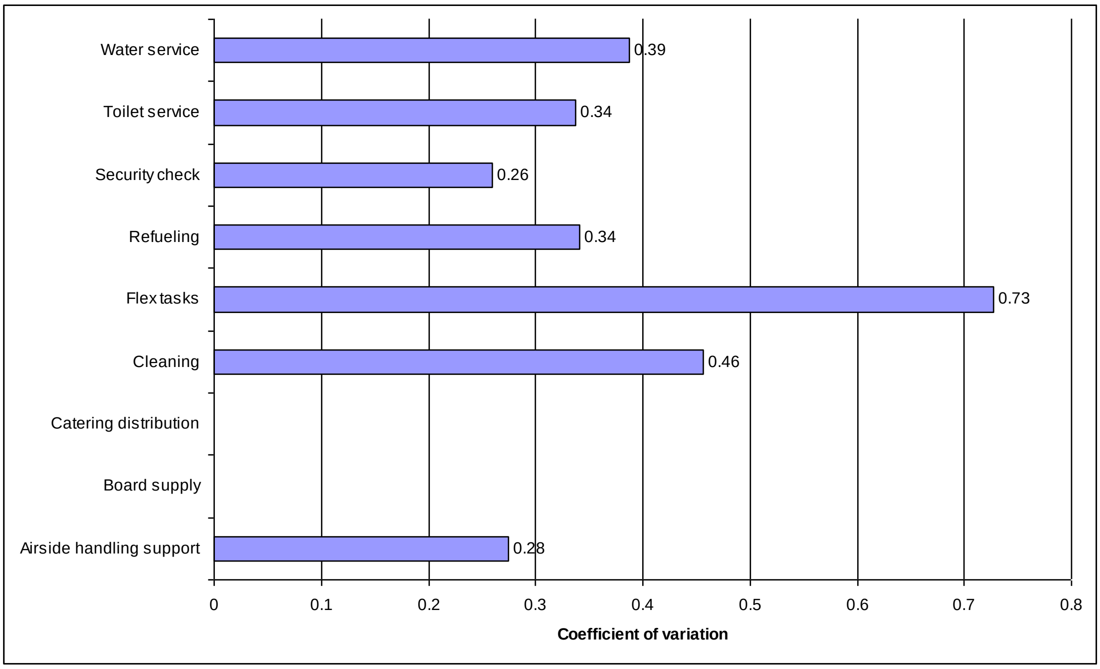

Figure 2.9: The coefficient of variation for the different services.

Figure 2.9 shows that the coefficient of variation is the largest for the Flex tasks. This indicates a relatively high variable process (according to the other processes), which can be logically declared by the characteristics defined in Table 2.1.

2.3 Choice for a service to analyse

As defined in Chapter 1, the goal of this research can be reached by analysing only one of the services of AS in dept. We can extent the insights of this analysis to the other services to give a good indication of the possibilities of implementing a more robust planning at AS.

We choose the refuelling service as a good service to use for a detailed analysis in the coming chapters, based on the following reasons:

- The refuelling service is arrival/departure *oriented*, so it is possible to level the workload and use longer ground times (within the permissible time windows) to advance or delay processes (see Section 2.2).
- The refuelling service works with flexible permissible time windows, which gives the controllers some space to plan the job more flexible within this time window.
- The refuelling service is a relative independent process, so by choosing this service we do not need to account for strict relations with other services of AS.
- The refuelling service has according to Figure 2.7 on average the longest norm process time compared to the norm ground times. For this reason a delay in the refuelling process can easily disturb the flight punctuality.
- According to Figure 2.8, the refuelling service has the second largest variation in process times compared to the other services. Since the cleaning service is outsourced to other companies, the standard deviation of the refuelling process is the largest under KLM’s own control.
- The refuelling service has, according to Figure 2.9, an average coefficient of variation. One should notice though, that the services with the two largest coefficients of variation (Flex and Cleaning) are not suitable for this research. The Flex Tasks have a very short duration and thereby they do not have a big impact on

the total performance of AS. As earlier mentioned, the cleaning service is being outsourced and thereby not controlled entirely by AS.

There are also some reasons not to choose the other services:

- The catering and cleaning services are outsourced to respectively KCS and Asito and Klüh so this is not being controlled entirely by AS. For this reason we do not analyse these services in depth.
- The flex tasks, water and toilet services have a quite short duration relative to the norm ground times, so planning these services more efficient will not lead to huge improvements in the prevention of delays. Furthermore, they have a significantly smaller standard deviation than the refuelling service.
- Towing is heavily dependent on the gate allocation, which is not under the influence of KLM. This makes the service not useful for this project.
- Pushback, de-icing and aircraft handling support are strictly arrival or departure processes so it is not possible to level the workload and use longer ground times (within the permissible time windows) to advance or delay processes.
- The de-icing process is furthermore a winter-process which is heavily dependent on the weather. This makes it also a less applicable service.

2.4 Conclusion

To achieve the goal of this project, we only have to analyse one of the processes of KLM Aircraft Services in depth and then extend the insights of this analysis to the other processes. To choose between the different processes, we first characterised them. The most important characteristics in the scope of this research are:

- The type of process (*strict* arrival/departure or arrival/departure *oriented*)
- The flexibility of scheduling the process (dependency on other processes and whether the process has flexible permissible time windows).
- The duration of the process
- The variation in the duration of the process

We conclude that the refuelling process is the best process to analyse in depth, based on that this process is arrival/departure *oriented*, relatively flexible to schedule and has a relatively long duration and high variation. Based on the characteristics mentioned in this chapter, the cleaning process would also be a good process to analyse. However, since KLM outsources this service, we chose not to.

3 Current way of planning

To be able to perform her tasks, AS requires sufficient employees and equipment at the right time in the right place. These requirements are captured in a capacity planning. The number of employees and equipment required on the day of operation depends on a lot of factors. First this is dependent on the number of scheduled flights from the flight schedule. Second, the duration of a task is important, since this varies per aircraft type, destination of the flight, equipment type, etc.

In this chapter, we describe the current way of planning. First, we describe the planning process of AS step-by-step. Second, we describe the standard times to perform a task and standard ground times: the norms.

3.1 Planning

The tactical planning department (ST) of GS runs the complete planning chain twice a year, because the planning makes a distinction between the winter (November – March) and the summer (April – October) period. There are also calculations made further ahead in time (2 till 5 years), but these calculations are out of scope for this project. Every step in the planning chain uses assumptions with respect to the workload (planning principles, norms and flight schedule) and with respect to the available resources (equipment and attendances of employees).

Moment of execution	**Goal**	**Instrument**
2 times a year (3 months before the start of the season)	- Check for feasibility (in terms of gate availability and equipment). - Create rosters - Budgeting schedule (estimate the needed workforce and equipment).	OPC: Operational Plan Check for the busiest week of the season.
1 to 2 months before the start of every single month	Adjustments to the schedule based on insights in advance.	RP Basic: Basic rolling planning made by ST based on the OPC. This is done for one representative week for every month of the season.
Every 4 weeks	Exact adjustments to the schedule based on the latest information.	RP Actual: The basic rolling planning is update every 4 weeks based on the latest information.
Daily (out of scope)	Optimize the schedule on the day of execution.	Day planning: Planning in CHIP with actual demand and actual capacity.

Table 3.1: The planning steps of GS.

In Table 3.1, the different steps of the planning chain are explained. These steps are integral for all of GS. For this project, the OPC and the rolling planning are relevant. The daily planning (scheduling the different operators to their tasks) is out of scope, because this project is executed on the tactical plan level. We explain the relevant part of Table 3.1 in the following sections.

3.1.1 OPC

The first phase in the planning for a new period is the Operational Plan Check (OPC). According to the work of Dekkers (2010), an OPC has the following goals:

- Determine whether the flight schedule determined by the network department is feasible.
- Give an indication of the corresponding costs.
- Indicate the required changes in capacity.
- Create rosters.

An OPC is executed twice per year, once for the flight schedule of the summer period and once for the flight schedule of the winter period. During an OPC, calculations are made for the workload of all departments for the busiest week of the coming period. Based on these calculations, business managers make choices for covering the workload. This coverage serves as a base for the creation of timetables for the different departments.

Flight schedule

Flight schedules are made by the Network-department and serve as a base for the activities of GS. The flight schedule is continuously updated by changed flights, aircraft, etc. For this reason the flight schedule used in the OPC is a snapshot. An example of such a snapshot can be found in Appendix D. Since AS handles aircraft from KLM, her partners and third parties, the flights of these airlines are also included in the flight schedule. ST checks this flight schedule for strange ground times and for that all airlines are included. ST also executes a check for the required and available equipment. These are not allowed to deviate much. If the flight schedule is feasible, then it is used in the OPC. When ST notices some errors, the flight schedule is sent back to Network.

Workload profile

To indicate the required capacity, a workload profile is made by ST for every day of the OPC-week (for an example see Appendix E). The combination between the number of tasks on a day and the duration of the single tasks based on the pre-defined norms leads, after an optimisation step, to a workload profile for every service provided by AS. We illustrate the creation of a workload profile by ST in Figure 3.1. The workload depends on a number of variables. We treat these variables step by step, following the numbers in the figure.

When the calculated workload is too high to be covered by the available equipment, ST can object to the given flight schedule and send it back to Network. Buying new equipment is unfeasible since it has a lead time of about one year. Furthermore this is a strategic decision and thereby out of scope for this project.

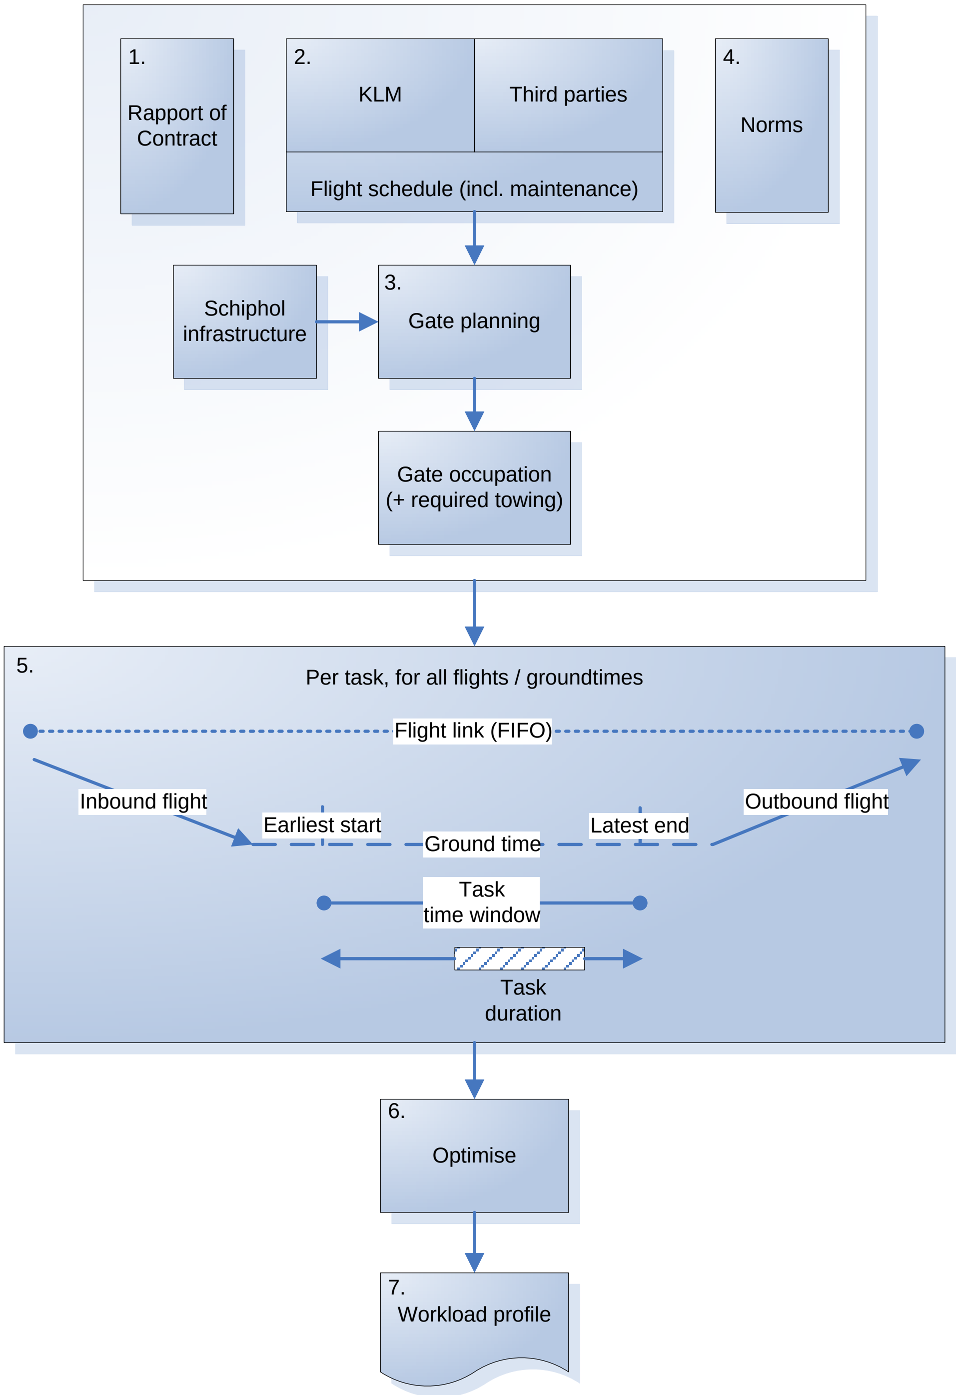

Figure 3.1: The creation of a workload profile.

- The first item in Figure 3.1 is the Rapport of Contract. This contains the made ground servicing agreements with KLM and contracted third party airliners.
- In item 2, the flight schedule is created. The flight schedule contains all flights from KLM and third party airliners including their aircraft maintenance schedule. We give an example in Appendix D.
- From the flight schedule follows the gate planning (item 3), depending on the available infrastructure at Schiphol. The gate planning determines whether an aircraft needs to be towed to a buffer and whether an aircraft can be refuelled on a gate where a hydrant fuelling system is present. Unfortunately, KLM is not able to

influence this gate planning in a large extent since it is under control of the Schiphol Airport Authority.
- In item 4, the norms are included. The norms are explained in Section 3.2.
- From the combination of item 1 till 4, item 5 is created. Here an inbound flight is linked to an outbound flight, which results in a flight link. This link defines which aircraft to use for which flight and thereby the ground time of an aircraft. The linking of flights is basically determined by Network, which creates these links based on the last-in-first-out (LIFO) principle for short ground times and on the first-in-first-out (FIFO) principle for long ground times. When network creates these links, it takes the minimum ground times from the norms into account.
- When the flight links are determined, an optimisation follows (item 6). ST generates for all task durations (defined in the norms) an earliest start and a latest end time, based on the flight links. The task needs to be executed within this window. The planning of these tasks is currently optimised in such a way that the workload profile is as fluent as possible, so ST strives to eliminate sharp edges.
- The result from this optimisation is a workload profile (item 7). We give an example of a workload profile in Appendix E.

Scheduling
Based on the workload profile provided by ST, the schedules for the different departments of AS are created by the resource planners. In an ideal situation the capacity planning of AS follows the workload profile very strictly. However, due to Dutch law (CAO and Arbeidstijdenwet), it is not possible to schedule on such a flexible level. Most employees have contracts dictating to work 40 hours a week in 5 days. By allowing this kind of contracts, AS constraints herself to only let employees work in shifts of about eight hours.

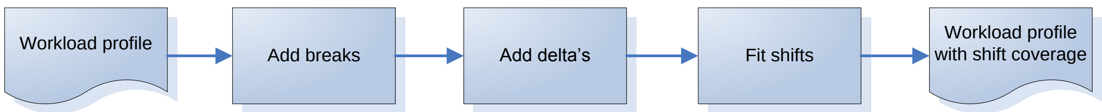

Figure 3.2: The transformation from a workload profile to a realistic coverage of workload.

The resource planners of AS edit the workload profile as given in Figure 3.2. They add two breaks, one of 30 minutes and one of 15 minutes, which occur preferably during the off-peak hours. However, the 30 minutes break should be planned between the start of a shift +2 hours and the end of a shift -2 hours.

The resource planners also add delta's to the workload profile. A delta is a certain value computed from historical data, which is extra work where ST is not able to account for in her planning. These delta's are thereby created manually by the process analysts of AS. For the refuelling process this extra workload consists of filling the small bowsers, maintenance fuelling for aircraft and unexpected marginal fuelling for aircraft.

After adding the extra workload of breaks and delta's, the resource planners fit a set of shifts on the resulting workload profile to cover the workload. The blue line in the workload profile displayed in Appendix E indicates the required shifts. Also the complete build-up of the workload profile is displayed in this appendix.

Based on the required shifts a basic schedule is created. The required shifts specify the required number of employees. This number is multiplied with the specific AWF (in Dutch:

aanwezigheidsfactor), the presence factor which is a factor which is the inverse of the absenteeism rate, for the coming period. After this correction a schedule is created which must first be approved by the Roosterwerkgroep and the Groepscommissie. When approved, specific operators are linked to the shifts, so every operator gets a personal roster.

3.1.2 Rolling Planning

The rolling planning is a monthly repeated re-run of the OPC. This is necessary because the airline industry is a very dynamic environment, so a workload profile made three months ago is never completely in line with the workload at the day of execution. There also exists variation in the flight schedule within a summer or winter period and there are other adjustments which need to be made. To cover all of this, a rolling planning is used which is updated every month. Were the OPC is based on the one busiest week of the period; the rolling planning is based on one representative week for every month to keep the planning specific to that certain month.

The rolling planning gives as output an overview of the daily shift including the required number of operators. The types of shifts used are the same as in the basic schedule, but the rolling planning can adjust the number of employees per shift by making use of flexible workforce.

3.1.3 Day of operation

The day of operation is out of scope for this project, but we discuss it briefly for completeness. On the day of operation, each department should have a sufficient amount of operators available: the amount the rolling planning dictates. Every day there can be last minute adjustments to the flight schedule, operators can be ill, go to training or be absent for other reasons. In this case, it is the job of the Shiftleader to make sure that he employs the required amount of operators. He can do this by varying the number of employees with a day off, are sent to training or by adjusting the number of employees hired from the employment agency. The day of operation is being coordinated from the Hub Control Centre (HCC), where controllers are monitoring the processes, setting priorities and dispatching operators to the different tasks.

3.2 Norms

When treating norms, we have to make a distinction between minimum ground times from GOMS (Ground Operations Manual Schiphol) and task norms. The norms from GOMS define the minimum ground time and the moment in time all task should be finished to allow the plane to leave on schedule. KLM measures her performance based on these times. The task norms are the standard times needed to execute a task at an airplane. We give an example of the task norms in Appendix C.

3.2.1 GOMS

Airliners define minimum ground times for serving aircraft. KLM stores these times in GOMS. These minimum ground times are the minimal time needed for preparing an aircraft for the next flight and vary per aircraft type and per turnaround type. If an aircraft needs a short turnaround (as defined in Chapter 2), the arrival and departure processes are executed subsequently and the norm ground time of the aircraft is very close to the sum of the norm times for performing the single tasks. In this case the latest completion time for processes is related to both the arrival and the departure time of an aircraft (arrival time + x = departure time – y). When an aircraft stays at Schiphol for a longer time (a long turnaround), the arrival and departure processes are executed independently of another. The latest completion time for arrival processes is in this case related to the arrival time (arrival time + x) and the latest completion time for departure processes is related to the departure time of an aircraft (departure time - y).

3.2.2 Task norms

The second type of norm is a standard time to perform a certain task. These times vary per aircraft type, destination and airliner. AS handles aircraft from KLM, her partners and third parties with all different wishes. This gives a wide array of varying time windows to complete the services. The norms are checked before each OPC for completeness and actuality. According to the work of Dekkers (2010), norms are used for the following purposes:

- Determine the demand for capacity.
- Make a planning for assigning available resources to tasks.
- Check whether the operation is executed according to the agreements made with the airliners, following from the rapport of contract.

A norm gives the time a certain task should take (which is mostly an average per aircraft type), the moment in time a task is allowed to start and the moment in time the task should be finished. These times are defined relative to the time of arrival or departure as described in Section 3.2.1.

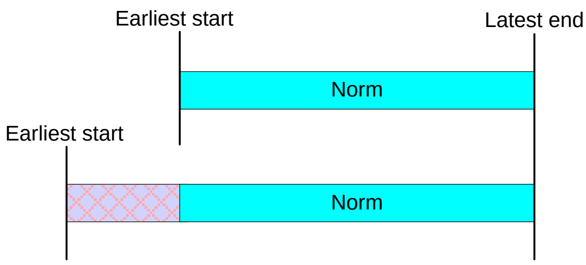

Figure 3.3: The difference between a normal norm and a norm with a flexible permissible time window.

Some processes work with flexible permissible time windows (in Dutch: schuifnorm). Figure 3.3 illustrates the difference between a normal norm and a norm with a flexible permissible time window. In the case of a flexible permissible time window the norm process time of the task is less than the difference between the earliest start and the latest end. This gives the controllers some slack in the time window to plan the job more flexible. According to the Oxford Dictionaries slack is defined as: “not taut or held tightly in position; loose”. So in capacity planning we define slack in the process times as the amount of time a task can be delayed without causing a delay to subsequent tasks.

Refuelling norms

AS uses three different types of norms for the refuelling process, because this process can be executed in two steps and this process takes a relative large percentage of the available ground time. We illustrate these three types of norms in Figure 3.4.

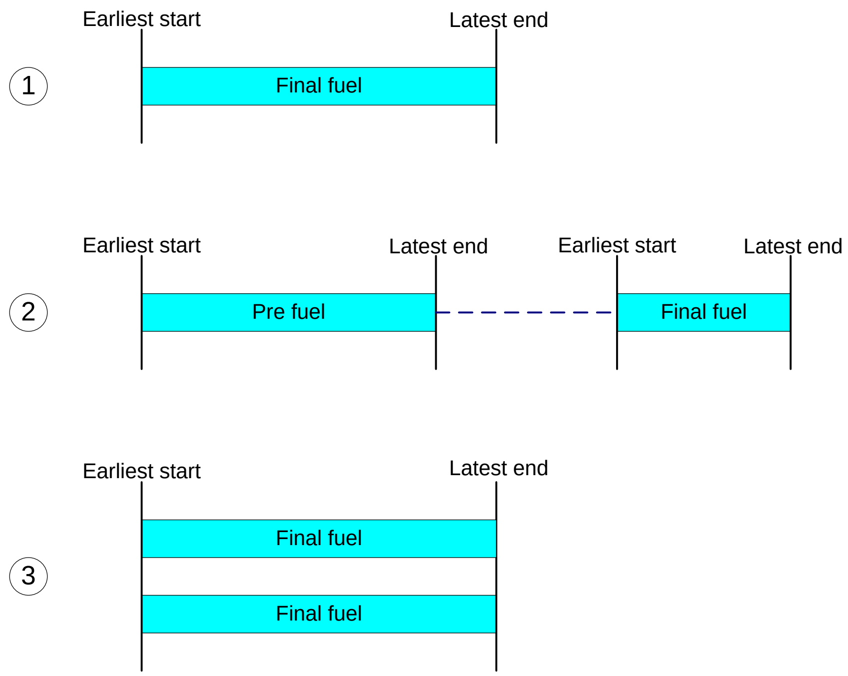

Figure 3.4: The three different types of norms for the refuelling process.

The first type is a single refuel, which is used for the smaller aircraft (European flights) and the larger aircraft (international flights) if they have a relative short turnaround. We describe the specific criteria for this in Appendix C, where we give the refuelling norms, broken down per type of aircraft, for the summer of 2011.

AS uses the second type of refuelling norm only for the large aircraft (intercontinental flights) and only if they have a long turnaround (in any case more than 3 hours, see Appendix C). In this situation an aircraft is pre-fuelled after arrival and the final fuel is executed just before departure, when the required amount of fuel is known. The total process time of the refuelling process is in this case longer than when an aircraft is refuelled all at once, because we have to account twice for setup and administration time.

When using the third type of refuelling norm, an aircraft is filled by two dispensers or bowsers in parallel. This makes the process a lot faster, so AS uses this when a large aircraft needs a very short turnaround. AS uses this type of norm also at freighter, because there is no hydrant fuelling system present at the cargo platforms. This means that multiple bowsers need to refuel the aircraft, since a single bowser does not have enough capacity for refuelling a freighter all at once.

3.3 Conclusion

The capacity planning of Aircraft Services is primary based on the workload following from the flight schedule and static norms. We conclude that Aircraft Services does, in the current situation, not account for uncertainty and unforeseen events in the arrival process of aircraft and in the process times of the different services. All process durations used for planning are directly derived from the norms and the arrival and departure time of aircraft is defined exactly according to the flight schedule without incorporating any slack for possible disturbances. This leads to a deterministic coverage of the total workload.

4 Quantitative analysis

In this chapter, we use data from airliners which are served by KLM Aircraft Refuelling in the month June 2011 for a quantitative analysis of the arrival punctuality and the refuelling process. We choose this month because it is a relative average month if we look to the amount of flights and the weather circumstances. The month lies just before the crowded holiday high-season and there are no extreme weather events like snow or ice which are highly complex to model within the scope of this project.

The arrival punctuality is defined by the actual time of arrival minus the planned time of arrival. By using this definition we determine the number of minutes an aircraft is late (a negative number means to early). We define the process time of the refuelling process as the time between the arrival of an operator at the aircraft and the time he finishes the complete task.

We start this chapter with Section 4.1 about data editing, where we describe the steps we take to obtain a clean dataset. Second, we describe in Section 4.2 the sources of uncertainty and unforeseen events affecting some more predictable influences on the processes of AS. Third, we zoom in on the arrival punctuality in Section 4.3 and the process times of the refuelling process in Section 4.4. We finish with a fit of probability distributions to the arrival punctuality, the refuelling process times and the duration of disturbances in Section 4.5. With these probability distributions, we are able to facilitate data generation outside the range of historical data, as input for the models from Chapter 5.

4.1 Data editing

When executing a quantitative analysis, the first thing to do is check the data for strange or incorrect entries. These entries should be investigated to get a valid representation of reality. It is possible that we need to consider these values as outliers, but for this we need further analysis.

A common way to identify outliers is based on the interquartile range (the range between the first and third quartile of a distribution). This method is based on the median and thereby more robust than a method based on the mean and standard deviation, because these are sensitive to extreme values. A method based on the median uses rank ordering instead of value influences. In this case: when a point falls more than 1.5 times the interquartile range below the first quartile or above the third quartile, this point is indicated as a possible outlier (Walfish, 2006). However, one should not discard possible outliers solely based on statistical methods. Outliers which cannot be readily explained should be investigated further. There should always be a good reason to discard outliers.

When we identify outliers based on the interquartile range in the arrival punctuality, this leads to a lot of outliers. When we investigate these values, we conclude that they cannot be discarded. Even when we use 3 times the interquartile range, we should discard all values that deviate more than one hour from the planned time of arrival. This is not feasible, because it happens regularly that flights are delayed with more than one hour. For this reason, we need to work with a more pragmatic approach, were we identify outliers based on the experience of the KLM management. They have to decide what kind of deviation from the planned time of arrival should be incorporate in the planning and what kind of deviation should be seen as extreme events.

After consulting management, we decided to couple the refuelling tasks to their corresponding arriving flight. We discard only complete couples (arriving flight + refuelling tasks) to minimise the risk of data pollution. We have to make the remark here that especially the refuelling data is highly polluted. In total we have to discard about 5% off our dataset.

We discard all couples with flights arriving more than 1400 minutes (about a day) too early or too late. KLM should not try to plan for these extreme values. We also discard all couples with 100% cargo flights, because they have other ground servicing protocols.

Based on the refuelling data, we discard couples for the following reasons:

- Flights without a corresponding refuelling task.
  - These flights are not serviced by KLM Aircraft Refuelling.
- Fuel tasks without a corresponding flight.
  - Probably a flight registration change.
- Pre-fuel tasks without a final fuel task.
  - An aircraft receives always a final fuel task.
- All couples with maintenance fuelling, storm fuelling or de-fuelling tasks.
  - These tasks pollute the process times of the final or pre-fuelling.
- The amount of fuel delivered is more than the theoretic maximum amount an aircraft can take.
  - This is an incorrect entry.
- No fuel or a negative amount of fuel is delivered at a final or pre-fuel task.
  - This is probably a task change or an incorrect entry.
- The fuel operator did not clock an arrival time.
  - This is an incorrect entry and thereby we cannot determine the total process time.
- Negative process times.
  - This is an incorrect entry.
- Process times of less than 7 minutes.
  - Management decided that this is too short to execute a task.
- Process times of less than 10 minutes for all aircraft larger than a Boeing 737.
  - Management decided that this is too short to execute a task on this type of aircraft, mainly due to the needed set-up time.

4.2 Factors influencing the processes

When creating a capacity planning, KLM has to take a lot of factors into account which are partly predictable and partly uncertain. We make a split between partly predictable and uncertain factors and treat these in the following sections.

4.2.1 Uncertainty and unforeseen events

Uncertainty and unforeseen events cause disturbances in the arrival punctuality of aircraft and the task duration of the different processes. This uncertainty and unforeseen events has different sources, which we describe in this section.

The weather
The weather is the factor with one of the biggest influences on the arrival punctuality. An aircraft can for example be early when its trip goes downwind of late when it goes upwind. Other bad weather influences like thunderstorms or snow can also cause big disturbances in the arrival punctuality.

The weather is also highly influential on the refuelling process. For example:

- When the next flight goes upwind, the aircraft consumes more fuel then downwind, so the refuelling takes a longer time.
- When the wind is very strong at the airport, AS needs to storm-fuel the smaller aircraft to add weight and thereby keep them stable on the ground.

Last minute changes
Last minute changes in the load factor of an aircraft in terms of cargo or the amount of passengers can cause an additional fuel task. This is due to the higher fuel consumption of heavier aircraft.

Gate planning
The gate planning influences the way of fuelling. When an aircraft stands at a gate, in most cases hydrant fuelling is possible. When an aircraft is placed on a buffer, hydrant fuelling is not possible so a bowser (with limited capacity) needs to do the refuelling.

The gate planning is uncertain because this is done by Schiphol and not under the influence of KLM.

Fuel-operator behaviour
The behaviour of the fuel-operators influences the process times of the refuelling process. The process times depend partly on how fast an operator works. The quality of our data depends also on the operator, because we assume that all operators press the confirm button on their handheld on the same moment, however in reality this is different.

Unforeseen events and process disturbances
There is always the possibility of unforeseen events disturbing the arrival punctuality or ground processes. A big example is the eruption of the Icelandic volcano in 2010 which forced the closure of some parts of the airspace. A less extensive example is a breakdown or technical failure at an aircraft, creating a last-minute change of aircraft and thereby delaying the flight.

Before or during the refuelling process, a lot of unforeseen events can occur. For example, an aircraft can be too late on the gate where it needs to be serviced, ground handling equipment can break down, cargo can stand in the way or an operator has to wait due to other reasons. We investigate the duration of disturbances in detail in Section 4.4 and 4.5.

4.2.2 Predictable influences on the planning

There are also some factors which are relatively predictable and have a high influence on the distribution of flights over the day, the arrival punctuality and the process times of the refuelling process. We describe these factors in this section.

Moment of the day

At Schiphol, there are certain peak moments, when most flights arrive or depart to ensure a short transfer for transfer passenger. This means that flights do not arrive and depart at Schiphol uniformly distributed throughout the day. We give an illustration of the distribution of arriving and departing flights over the day in Figure 4.1, split to European (EUR) and intercontinental (ICA) flights.

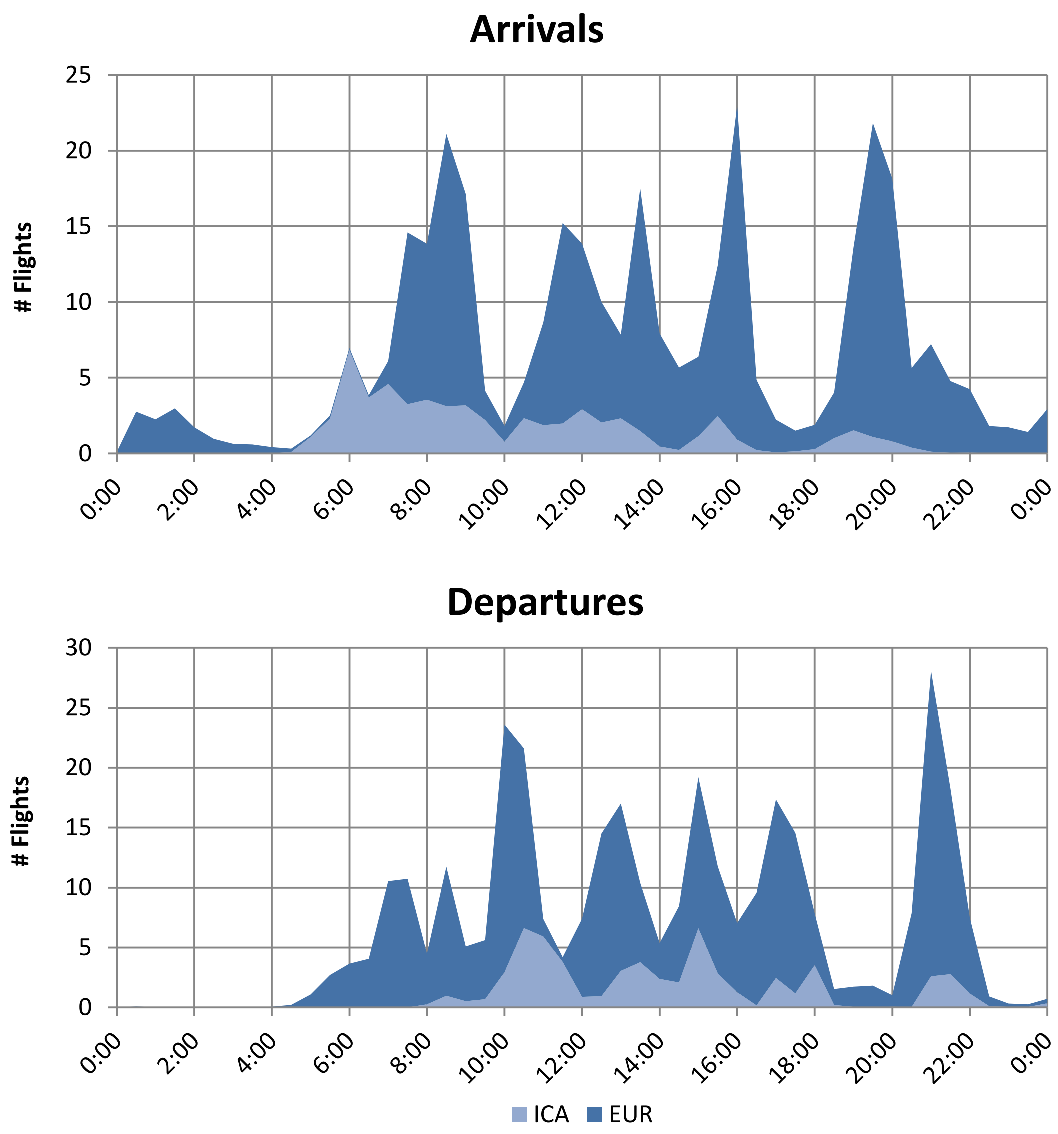

Figure 4.1: The distribution of the number of flight arrivals and departures over the day at Schiphol.

We check whether the moment of the day influences the arrival punctuality. The means as displayed in Table 4.1 gives an indication that there exists a difference.

	Mean	StDev	Median
Morning	-3.08	29.32	-7
Evening	0.418	31.985	-4
Afternoon	-0.427	36.27	-5

Table 4.1: Summary statistics of the arrival punctuality per part of the day (defined as minutes late).

To check this hypothesis we execute a 2-Sample T-Test to compare the difference in means. The results are displayed in Table 4.2.

T-test	Estimate for difference	P-value
Morning =Afternoon	-2.649	0.000
Morning = Evening	-3.494	0.000
Afternoon = Evening	-0.845	0.101

Table 4.2: The results of the 2-Sample T-Test for the difference in mean of the arrival punctuality.

We conclude from Table 4.2 that there exists a significant difference in arrival punctuality between an arrival in the morning and an arrival in the afternoon and evening at the 95% confidence level, since the p-value is smaller than 0.05. There exists no significant difference between an arrival in the afternoon and an arrival in the evening. We conclude from this that the arrival punctuality of flights arriving in the morning differs significantly from flights arriving in the afternoon or evening.

Seasonal influences

The airline industry is heavily exposed to seasonal influences. The months July and August are the busiest because this is the tourist high-season. During these months, aircraft are easily delayed because of crowded skies or a shortage on airport capacity. Figure 4.2 gives an overview of the distribution of flights over the year, divided over EUR and ICA flights. We conclude from this figure that there are significantly more EUR flights than ICA flights and that the deviation in the number of flights per month is larger at EUR flights than at ICA flights.

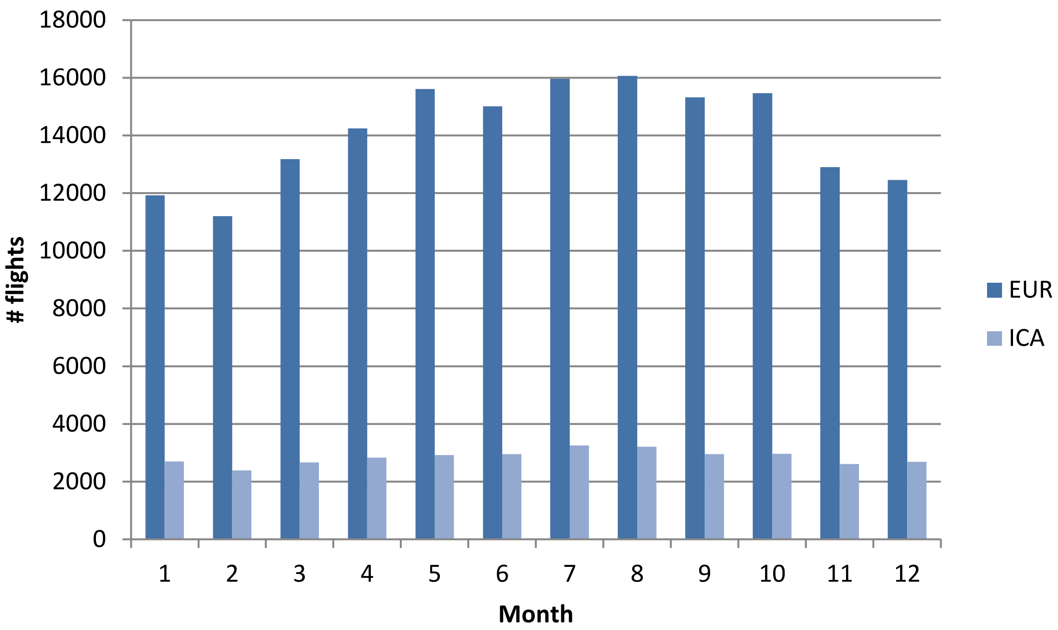

Figure 4.2: The number of flights per month in 2011, split to EUR and ICA flights.

We illustrate the distribution of flights over the different planning periods (summer/winter) in Figure 4.3. From this figure, we can conclude that the deviation in the number of flights between the summer period and the winter period is significantly bigger at EUR flights than at ICA flights.

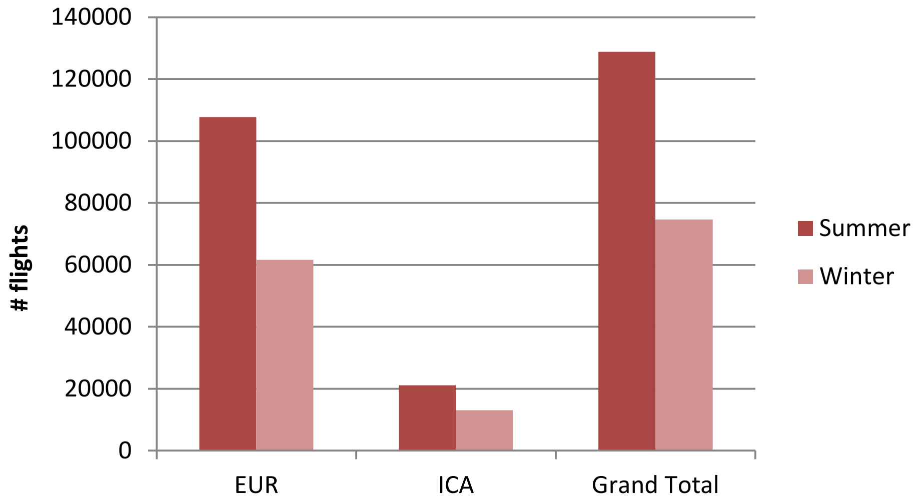

Figure 4.3: The number of flights per period (summer/winter) in 2011, split to EUR, ICA and the grand total.

The arrival punctuality deviates a lot per month. However, creating a figure indicating the arrival punctuality would give a disturbed image, since this would give the average arrival punctuality where a flight arriving early can compensate for a flight arriving late. For KLM, it mostly matters when a flight arrives late, because this shortens the available ground time. Figure 4.4 gives an indication of the lateness of aircraft split per month, where the lateness of an aircraft arriving early is set to zero. Again we see differences over the months, which are influenced by the factors as described earlier in this section.

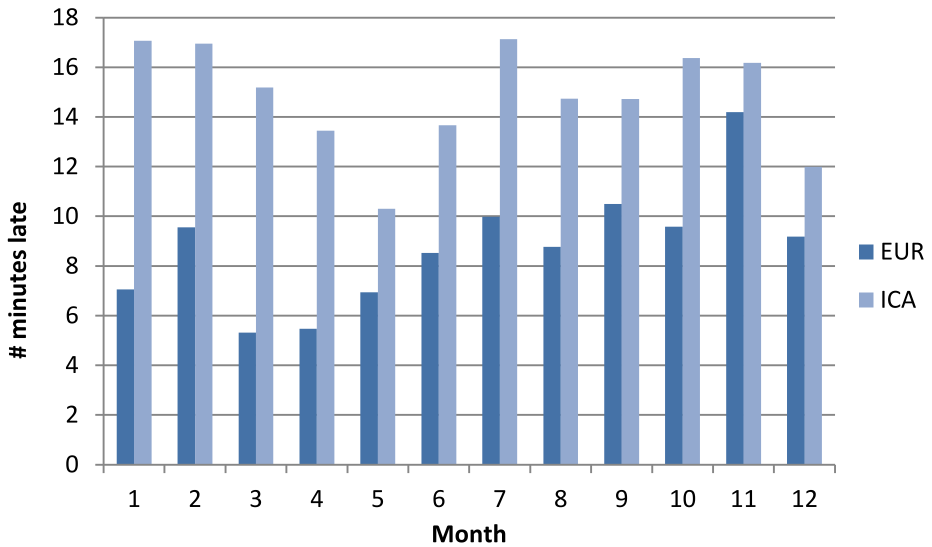

Figure 4.4: The average lateness in the arrival of aircraft per month in 2011, split to EUR and ICA flights.

We also analyse the arrival punctuality per period (summer/winter). An overview of this is given in Figure 4.5. This figure illustrates that there is a relative large difference between the summer and the winter period at the EUR flights. There is almost no difference in arrival punctuality between the two periods at the ICA flights.

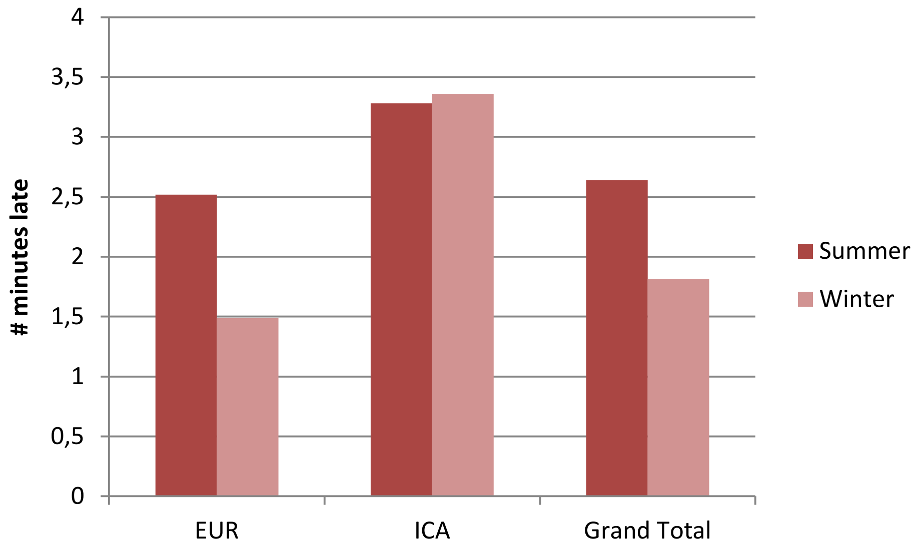

Figure 4.5: The average arrival punctuality per period (summer/winter), split to EUR, ICA and the grand total.

We conclude from Figure 4.5 that the arrival punctuality of ICA flights does not differ a lot between the two periods. The difference is larger in the general arrival punctuality and at the EUR flights.

Flight link

The ground time (and thereby the time available to service an aircraft) is highly dependent on the flight link. The flight link is, as described in Chapter 3, the link between an inbound flight and an outbound flight, executed by a single aircraft. KLM can choose to change the aircraft which executes a flight when an aircraft breaks down or for various other reasons. This influences the departure time of an aircraft and thereby the latest end time for the refuelling process. For this reason, it is important to set the latest end time on the flight level and not for single aircraft.

In the scope of this research, we want to know if a change in flight link also affects the workload of the refuelling department and thereby whether we need to create new flight links when we build a model which generates workload using stochastic arrival punctuality. We look to the distribution of ground time of an individual aircraft type on a single day in Figure 4.6. This figure illustrates clearly that the flight link between two flights has been changed, namely flight KL0428 and KL0692. Notice that a high peak does not have to indicate a delay. In most cases this means that another aircraft which stood in reserve at a buffer is used for executing this flight. The question is now whether this gives large differences in the workload for that day. For the total amount of workload, this does not, because the permissible flexible time window for refuelling one flight is shortened and for another it is made larger. A change in the flight link does however influence the time window during which a task can be executed. This means that a change in flight link can influence the location and height of a peak. In this research we keep the existing flight links to keep our analysis manageable, but in reality, KLM should keep in mind that a change in flight link can affect the impact of peaks in the workload.

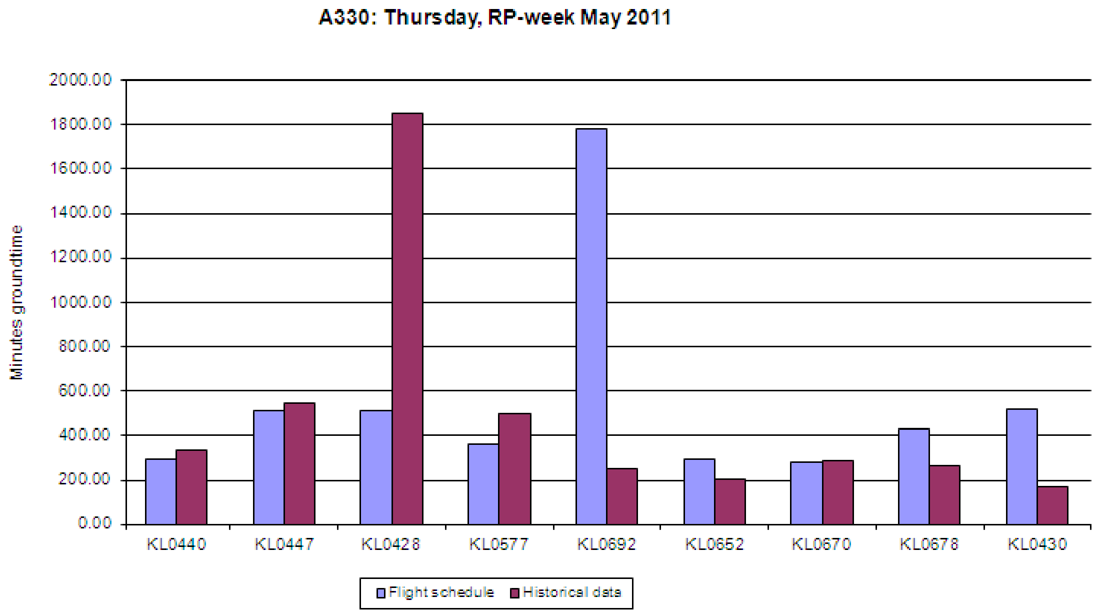

Figure 4.6: The planned ground time according to the flight schedule compared to the real ground time from the historical data split per flight.

4.2.3 Dependencies

An important aspect to mention is that there exist dependencies between the arrival punctuality of aircraft. When one aircraft is arriving late, it is likely that other aircraft are also delayed that day, for example due to the weather circumstances. We investigate this dependency by calculating the autocorrelation between arriving flights at ten lag values. This means that we calculate the autocorrelation in the arrival punctuality of an arriving flight and a flight arriving till a maximum of ten flights later. To indicate possible autocorrelation in the arrival punctuality we analyse a sample day in Figure 4.7. We conclude from Figure 4.7 that there exists positive correlation, but this is not significant on the 95% confidence level.

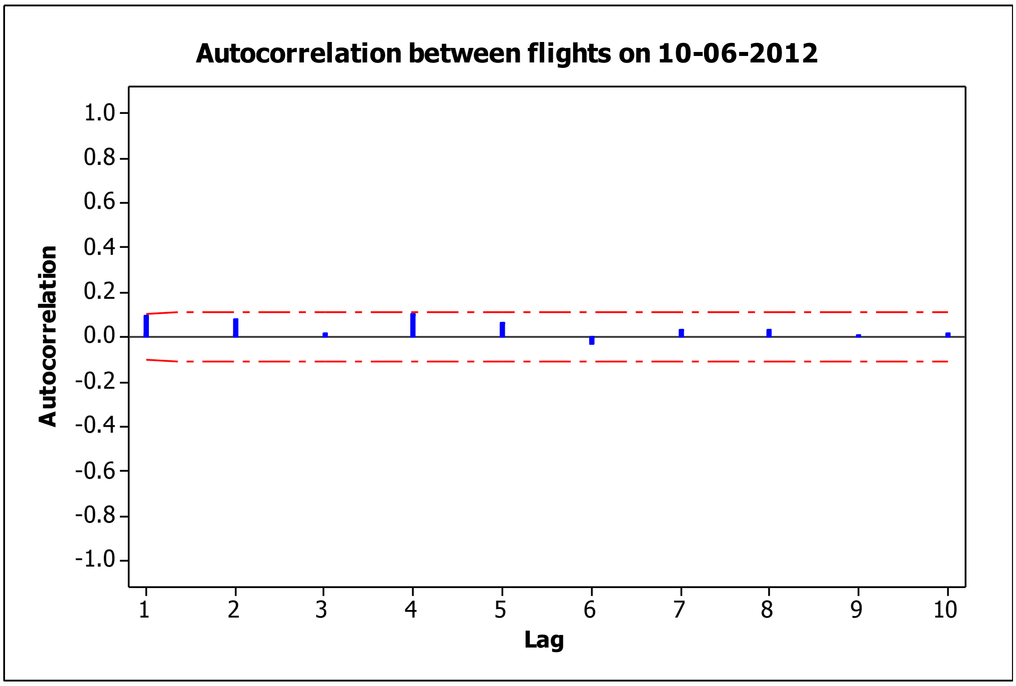

Figure 4.7: The autocorrelation between flights in the arrival punctuality (the red line is the 95% confidence level).

Some other dependencies are described earlier in this section: the moment of the day and the weather circumstances. For example: when aircraft are delayed in the morning, they can delay the schedule of KLM and Amsterdam Airport Schiphol for the whole day.

Destination / Aircraft type / Passenger load factor
The process time of the refuelling process follows mainly from the amount of fuel a fuel-operator needs to deliver. This amount depends on the type of aircraft (consumption per flight minute), the destination (the amount of miles an aircraft needs to fly) and the passenger load factor (an empty aircraft consumes less fuel than a fully loaded aircraft). We analyse this further in Section 4.4.

4.3 Arrival punctuality

In this section, we address the arrival punctuality of the aircraft handled by KLM Aircraft Refuelling. We can make a distinction in the data on various levels of detail; for example between European or Intercontinental flights, per airliner, per aircraft type or even per flight destination. This all affects the dynamics of the arrival punctuality. In the following sections, we give an indication of the magnitude of these dynamics with a focus on the month June 2011.

4.3.1 Nature of flights

The flights from which we determine the arrival punctuality are all from airlines served by KLM Aircraft Refuelling. We give an overview of the distribution of these flights in Table 4.3. The main part of the flights is coded KL which stands for KLM and KLM Cityhopper. Another relatively big part is from Transavia (HV) and Delta (DL).

Airline	# Flights	% of total	Cum %
KL	8242	81.13%	81.13%
HV	1247	12.27%	93.40%
DL	505	4.97%	98.38%
SU	44	0.43%	98.81%
MP	17	0.17%	98.98%
AAN	34	0.33%	99.31%
PY	19	0.19%	99.50%
CZ	20	0.20%	99.69%
KQ	12	0.12%	99.81%
SQ	12	0.12%	99.93%
A9	7	0.07%	100.00%

Table 4.3: The distribution of flights over the airlines served by AS in the month June 2011.

We categorise all flights in European, intercontinental flights and other flights (e.g., test flights). Figure 4.8 gives an overview of this, where we see that the main part of the flights, about 82%, comes from a European airport.

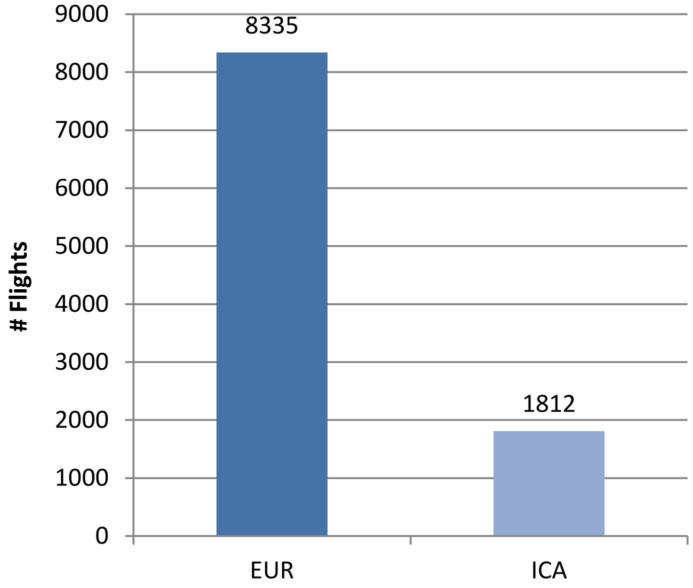

Figure 4.8: The total number of flights in the month June 2011 split to European (EUR), and intercontinental flights (ICA).

We describe the properties of EUR and ICA flights by a statistical summary in Table 4.4. From this summary we conclude that the mean is relative close to zero, because on average all flights arrive relatively on-time (it would be bad if this were different). We conclude that the ICA flights are more likely to be disturbed than the EUR flights because of the higher mean and standard deviation.

An important aspect to notice is the difference between the mean and the trimmed-mean (TrMean). The trimmed-mean is the mean, calculated after discarding 5% of the values on both the high and low end of the distribution. Since the difference between the mean and the trimmed-mean is relative large, we conclude that the 5% in the tails is quite significant for the distribution of the data and the location of the mean. This is also stated by the high kurtosis. This indicates that the data has a large peak, and that thereby a large part of the variance is caused by the extreme values.

NatureOfFlight	Count	Mean	TrMean	StDev	Minimum	Maximum	Range
EUR	8335	0.904	-2.01	30.09	-144	1314	1458
ICA	1812	2.226	-2.607	40.18	-67	443	510

NatureOfFlight	Q1	Median	Q3	IQR	Skewness	Kurtosis
EUR	-10	-3	5	15	15.09	493.45
ICA	-17	-5	8	25	4.23	28.76

Table 4.4: The properties of the arrival punctuality of arriving flights split to EUR and ICA flights in the month June 2011.

Table 4.4 displays that the data is positively skewed. Figure 4.9 gives the same indication, since the indication for the mean does not lie on the same spot as the median line. This indicates a tail on the right side of the distribution which is longer than the left tail, so the distribution is asymmetric. Due to this skewness, we cannot assume a normal distribution for this data.

With the boxplots from Figure 4.9, we compare the distribution of EUR flights with the distribution of ICA flights based on the median and the different quartiles. This measure is more robust to extreme values then measures based on the mean and standard deviation, since the median is not influenced by extreme values. We conclude from Figure 4.9 that the arrival punctuality of the ICA flights is more widely distributed than the arrival punctuality of the EUR flights. In general, we conclude that an ICA flight is more likely to be disturbed due to the broader distribution of data.

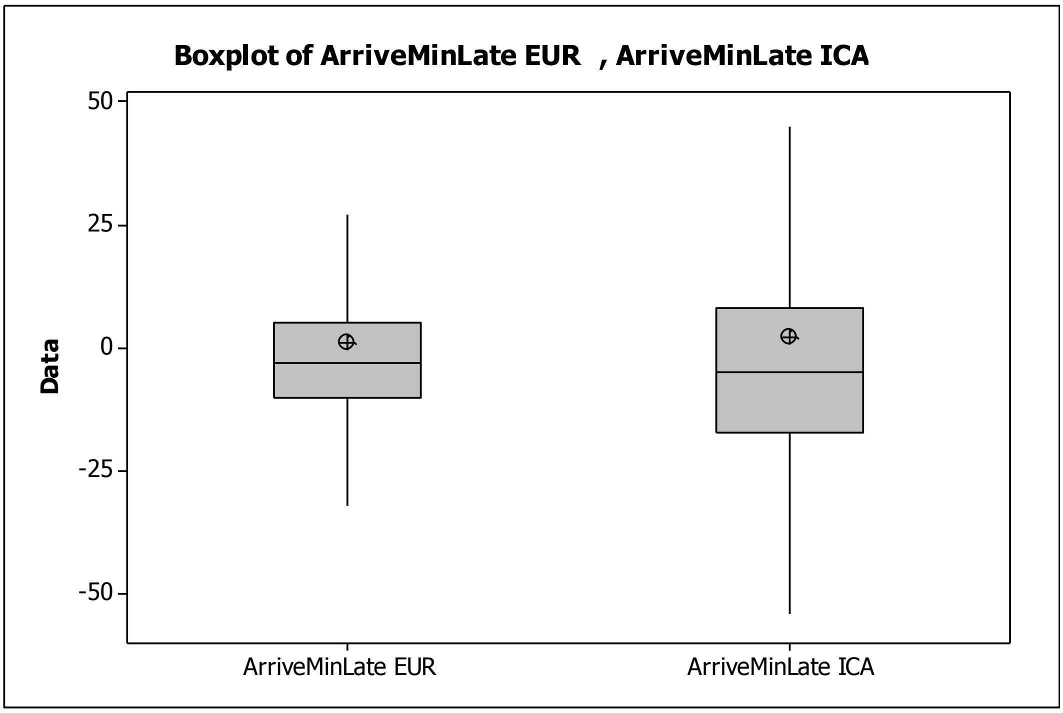

Figure 4.9: **Boxplots of the arrival punctuality split to the nature of flight in the month June 2011 (including the cross as mean symbol).**

To test whether we need to make a distinction between the arrival punctuality of EUR and ICA flights, we used a 2-sample T-test to compare the means of the two types of flights. We conclude from this test that the mean of EUR flights is, on the 95% confidence level, not significantly different from the mean of ICA flights so we can consider them similar. The results of the tests are given in Table 4.5.

T-test	Estimate for difference	P-value
μ EUR6 – μ ICA6 = 0	-1.32	0.186

Table 4.5: The results of the 2-Sample T-Test for the difference in mean of the arrival punctuality between EUR and ICA flights in the month June 2011.

4.3.2 Aircraft types

The airliners served by AS operate a wide range of aircraft types, with a lot of subtypes. We generalise these subtypes to their general aircraft type and display the aircraft types which are served most by KLM Aircraft Refuelling in Figure 4.10. We create, based on this figure, a clear link between general aircraft type and the nature of flight. There are some exceptions, but mostly an aircraft is linked to either EUR or ICA flights. The exceptions come mostly from holiday flights to countries outside Europe around the Mediterranean Sea, which are officially ICA, but are served by smaller airplanes which fly mostly in Europe.

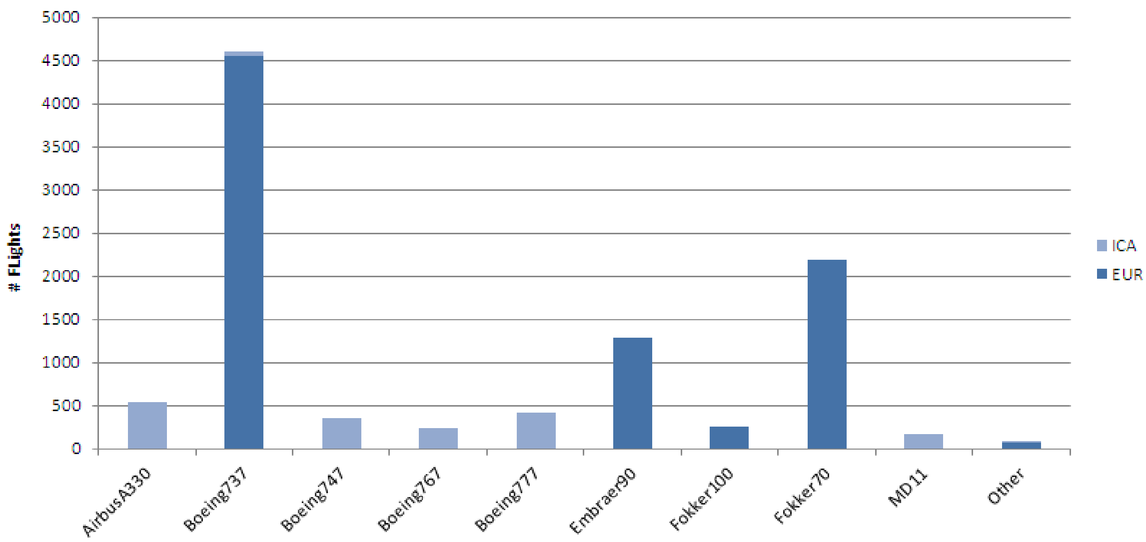

Figure 4.10: The nature of flight of different aircraft types in the month June 2011.

We now check whether the arrival punctuality of the nature of flight (EUR and ICA) is a good indication for the arrival punctuality of all aircraft with the same nature of flight. To check this, we create the boxplots from Figure 4.11, where we again see skewed distributions. We see differences in the distribution of the arrival punctuality between different aircraft types.

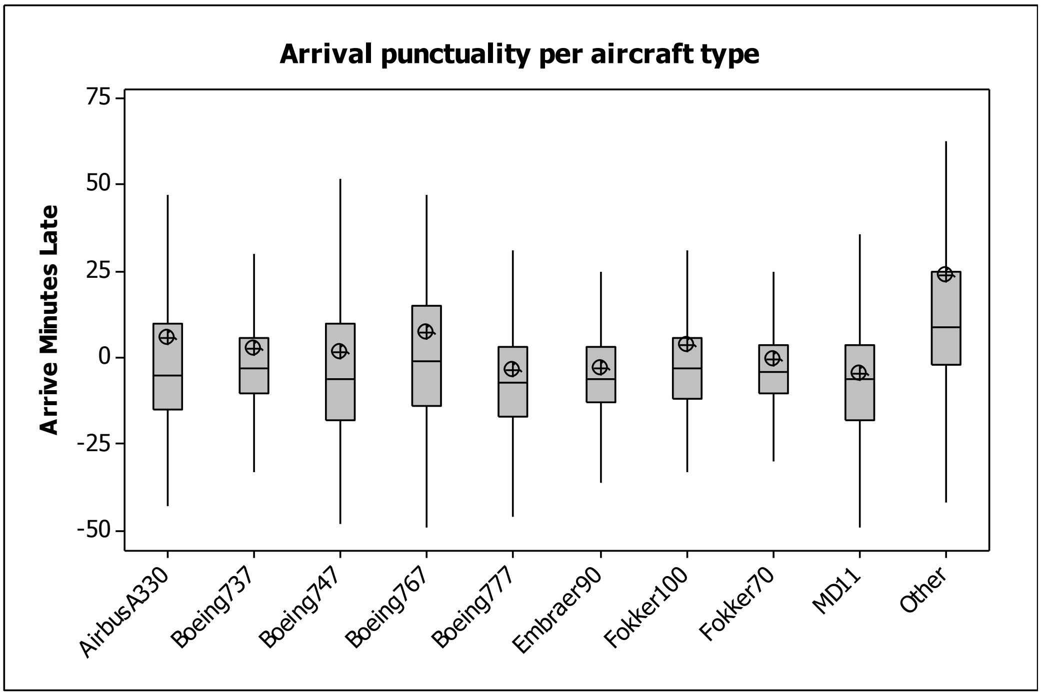

Figure 4.11: Boxplots of the arrival punctuality per aircraft type in the month June 2011 (including the cross as mean symbol).

To test the hypotheses that the arrival punctuality of a single aircraft type is similar to the arrival punctuality of its nature of flight, we use a 2-sample T-test to compare the means of the different aircraft types to the mean of their nature of flight. We conclude from this test that, on a 95% confidence level, for some aircraft the arrival punctuality is similar to the arrival punctuality of their nature of flight and for some not. The results of the tests are displayed in Table 4.6. We will come back to these results in Section 4.5.1, where we fit theoretical probability distributions to the historical data.

NatureOfFlight	GeneralAircraftType	Mean of aircraft type similar to nature of flights (EUR/ICA)	P-Value
EUR	Boeing737	No	0.004
	E90	No	0.000
	F100	Yes	0.310
	F70	Yes	0.187
ICA	AirbusA330	Yes	0.103
	Boeing747	Yes	0.773
	Boeing767	Yes	0.081
	Boeing777	No	0.002
	M11	No	0.003

Table 4.6: A summary of the results from the 2-sample T-tests comparing the mean of the arrival punctuality of the different aircraft types to their nature of flight.

4.4 Refuelling process

We address the refuelling process and its dynamics in this section. Hereby we focus on the month June 2011, just as in the previous sections. The total time that refuelling an aircraft takes depends on the process time of refuelling and the duration of possible disturbances. First we describe possible disturbances in the refuelling process and the way this affects the process times. Second we describe the process times itself and how this is influenced by various factors.

As previously mentioned, we define the process time of a refuelling task as the time an operator finishes this refuelling task minus the time he arrives at the aircraft. This looks like a solid measurement, but in reality the time clocked as the arrival time differs per operator. Most operators follow the rules and clock their arrival at the moment they arrive at the aircraft. However, some operators clock the arrival time at the moment they exit the vehicle (after some administrative work) to start executing the fuel task. This gives several problems. First the most important one: the measurement is less reliable. Second, when a disturbance occurs, an operator can clock this disturbance before he has clocked his arrival. However, it is physically impossible for most types of disturbances to see the disturbance before arriving at the aircraft. The only exception is an obligatory working stop in case of a thunderstorm, this can occur before an operator arrives at the aircraft. We will treat these problems in the following sections.

4.4.1 Disturbances

There were 1667 disturbances in the month June 2011, caused by different reasons as mentioned earlier. In this month, there were 10525 fuel tasks. This means that a disturbance occurs in 15.84% of all refuelling tasks. The duration of these disturbances varies from 1 to 67 minutes, with an average of 8.76 minutes. We give an overview of the dynamics of the disturbance times in Table 4.7.

	Count	Mean	TrMean	StDev	Minimum	Maximum	Range
TotalDisturbanceTime	1667	8.759	7.784	7.909	1	67	66

	Q1	Median	Q3	IQR	Skewness	Kurtosis
TotalDisturbanceTime	4	6	11	7	2.8	12.02

Table 4.7: Summary statistics for the disturbance times in the month June 2011.

Almost all disturbances occur between the time a fuel operator arrives at an aircraft and the time he finishes his task at this aircraft. This means that, to obtain the real refuelling process time, the disturbance time needs to be subtracted from the actual refuelling process times. Executing this is complicated, due to the different behaviour of fuel operators as mentioned earlier in this chapter. After consulting management, we decided to partially neglect this behaviour and give the recommendation to try to change it. If the new process time complies with the following rules (obtained from Section 4.1), then we do not subtract the disturbance time:

- Process times of less than 7 minutes.
- Process times of less than 10 minutes for all aircraft larger than a Boeing 737.

In all other cases we define, from this point on, the real refuelling process time as: finish time actual refuelling task minus arriving time fuel operator minus possible disturbance time.

4.4.2 Process times

As described briefly in Section 4.2, the amount of fuel an aircraft needs depends among others on: the type of aircraft, the destination of the flight and the amount of passengers and cargo on-board. The quantity that needs to be delivered by a fuel operator gives a lot of variation in the process times of the refuelling process. Figure 4.12 gives an illustration of this. KLM deals with this variation by splitting the historical data to aircraft type and thereby also to EUR and ICA.

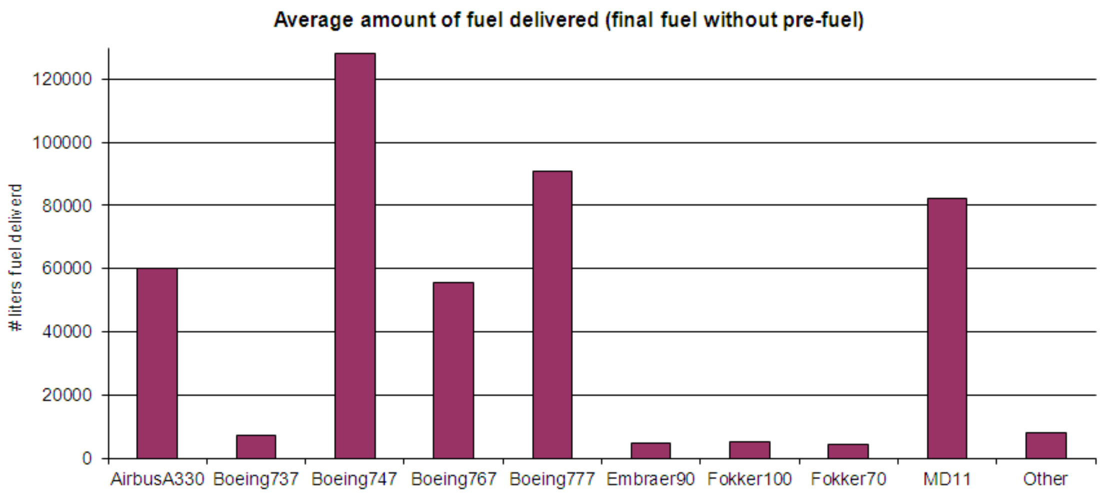

Figure 4.12: The average amount of fuel delivered in a final fuel (without a pre-fuel), split per aircraft type.

KLM splits the refuelling tasks to final fuel and pre-fuel tasks. Only the large aircraft (flying ICA flights) are in some cases pre-fuelled. The norms defining when to pre-fuel an aircraft are given in Appendix C. We give, per aircraft type, the percentage of aircraft which are pre-fuelled in Table 4.8.

AircraftType	% Pre-fueled
AirbusA330	0.37%
Boeing737	0.00%
Boeing747	21.57%
Boeing767	0.42%
Boeing777	28.99%
Embraer90	0.00%
Fokker100	0.00%
Fokker70	0.00%
MD11	31.28%
Other	0.00%

Table 4.8: The percentage of aircraft which are pre-fuelled split per aircraft type in the month June 2011.

The process time of a refuelling task depends on whether an aircraft is pre-fuelled or not. When an aircraft receives a pre-fuel and a final fuel, the total process time is longer than when an aircraft is refuelled at once: with only a final fuel. This is due to the set-up times (connecting the dispenser to the aircraft and administration). An advantage of pre-fuelling is that the process time can be split, so a major part of the refuelling can be done at an off peak moment. We give an overview of the average process times of the refuelling process in Figure 4.13. The process times are split in pre-fuelling tasks, final fuelling after a pre-fuelling task and fuelling at once (with only a final fuel task).

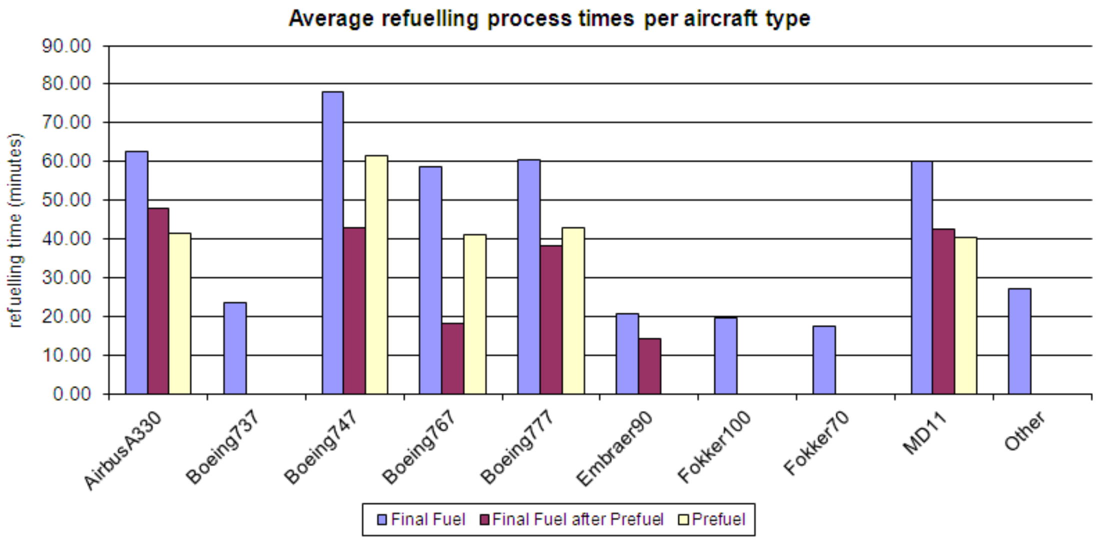

Figure 4.13: The average refuelling process times in the month June 2011, split per aircraft type and type of fuel task.

To give an indication of the distribution of the refuelling process times, we create boxplots of the final fuel tasks (without pre-fuel) in Figure 4.14. This figure gives a clear distinction between aircraft flying EUR flights (blue) and aircraft flying ICA flights (red). The EUR aircraft are much smaller and thereby refuelling takes less time than refuelling ICA aircraft. We also see relatively larger differences within these natures of flight. We conclude from Figure 4.14 that the refuelling process time of ICA aircraft is more widely distributed than the refuelling process time of EUR aircraft, indicating a higher variability.

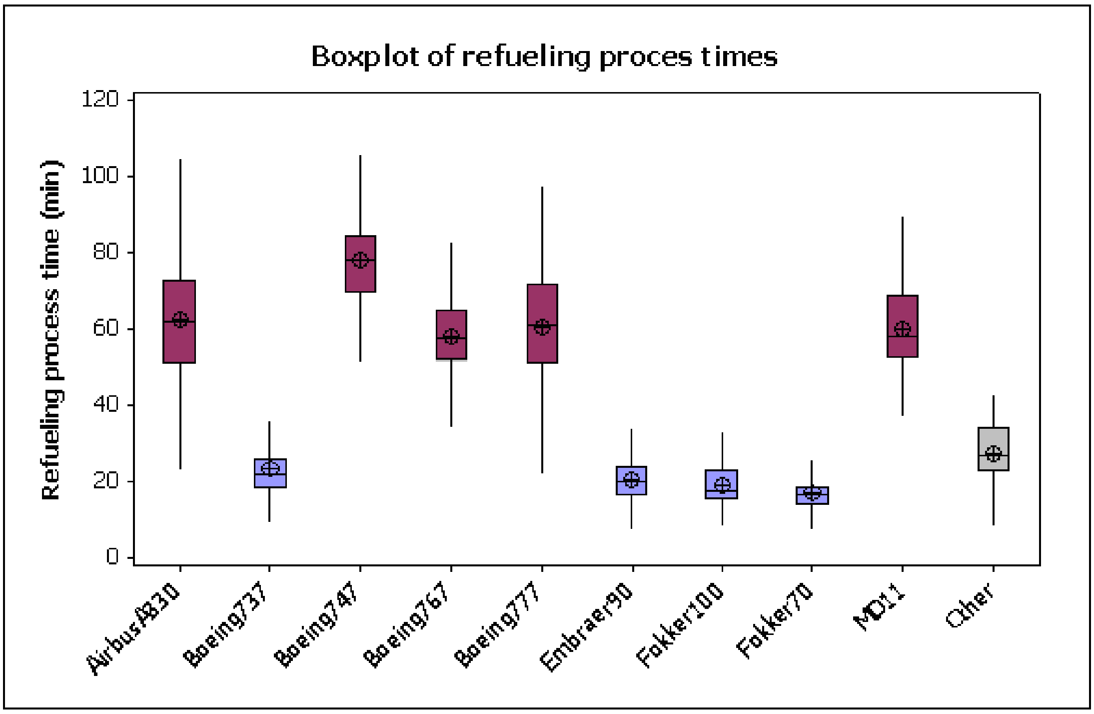

Figure 4.14: Boxplots of the process times of the final fuel tasks (without pre-fuel), split per aircraft type.

4.5 Probability distributions for the different processes

We use probability distributions to specify the dynamics of a process in terms of variability and the location of the mean. By using this probability distribution, we are able to model random input in our process and give a realistic representation of reality. According to Law (2007), there are three main approaches to specify random input data:

1. Use the data themselves directly.
2. Use data to define an empirical distribution function.
3. Fit a theoretical distribution function to the data.

These approaches all have advantages and disadvantages, which we describe in Table 4.9. According to Law (2007) one should first try to fit a theoretical distribution function, but if this is not possible due to complexity of the data structure or due to some other good reasons relative to Table 4.9, then one should take an empirical distribution function or use the data directly.

	Advantages	Disadvantages
Direct data usage	• Includes intrinsic data issues, such as possible correlations and time-varying parameters. • Useful for model validation.	• Data describes what happened in history, not what could have happened using similar data. • Usually insufficient data to make all simulation runs.
Empirical distribution function	• Unlimited input data generation. • Facilitates data patterns, different from, but comparable to historical data.	• Does not include intrinsic data issues, such as possible correlations and time-varying parameters. • No smoothing of possible irregularities in the data resulting from limited observations.
Theoretical distribution function	• Smooths possible irregularities in data resulting from limited observations. • Easier modifications (e.g. variation in order size, setup times), useful for sensitivity analysis. • See the advantages of an empirical distribution function.	• Does not include intrinsic data issues, such as possible correlations and time-varying parameters.

Table 4.9: Advantage and disadvantages of the three main approaches to specify random input data.

4.5.1 Arrival punctuality

When analysing the arrival punctuality, we have to account for negative values in the data. Aircraft arrive on time, too late or too early. Due to this last property, the negative values are also valid and we should incorporate them when we specify the random input data. Negative values give some trouble when fitting a theoretical distribution, since most of the theoretical probability distributions are only defined for positive values. The most common used exception on this is the normal distribution, but since this distribution is symmetric, it does not fit to our data, which is positively skewed. We solve the problem of negativity in the data by adding or subtracting a constant to the distribution, resulting in (depending on the type of distribution) a 2- or 3-parameter distribution. This constant is an extra parameter, called the shift or threshold parameter.

We conclude from Figure 4.4 that the arrival punctuality of aircraft differs significantly over the different months of the year. For this reason, we need to create a specific probability distribution for the arrival punctuality of every month. When there is not enough data available to create a distribution, we need to standardize data from other months to the month we evaluate to increase the data availability. In the case of arrival punctuality one month of data is sufficient, because there were 10147 flights in the month June 2011. For this month, we investigate the distribution of the arrival punctuality. Since the T-test from Table 4.5 indicates no significant difference between the arrival punctuality of ICA and EUR flights, we consider them statistically equal. This means that we are allowed to fit one distribution over the entire dataset from the month June. When this gives a good fit, no further differentiation in the dataset is needed.

From Section 4.3, we already learned that the distribution of the arrival punctuality has a large tail to the right and that the distribution is positively skewed. We create the histogram of Figure 4.15 to illustrate this. We also plot the 3-parameter-lognormal distribution in this figure, which is the best fitted distribution. The parameters for this distribution can be found in the legend of the figure.

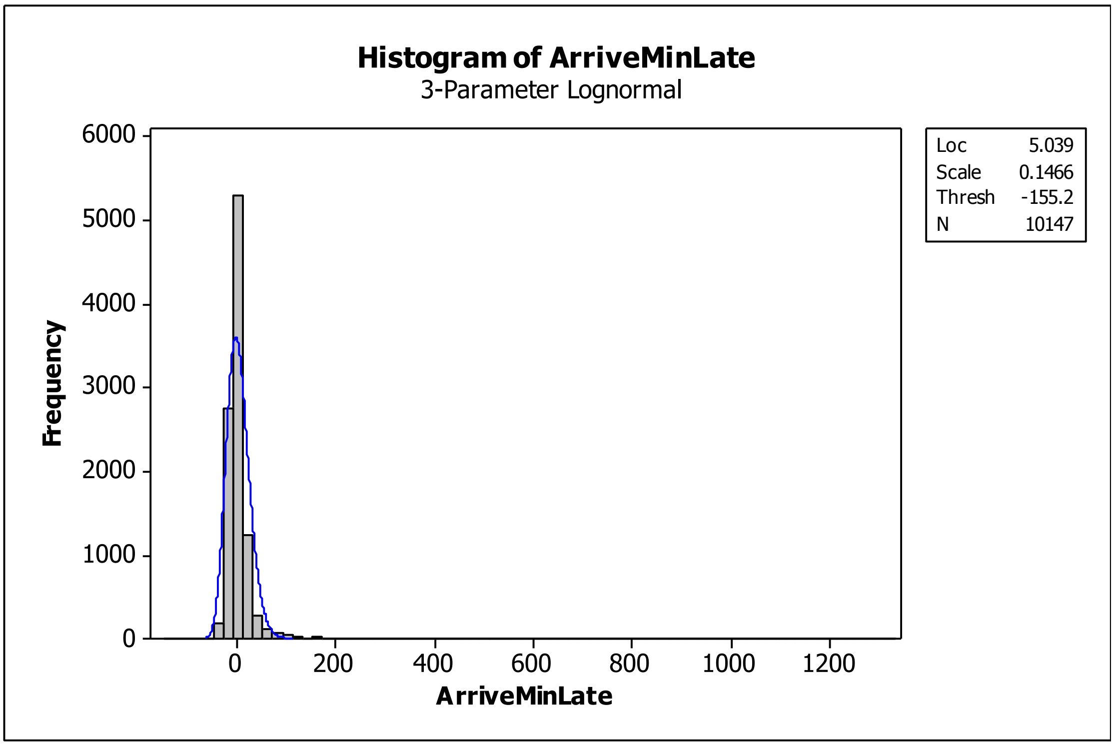

Figure 4.15: Histogram of the arrival punctuality in the month June 2011.

Figure 4.15 illustrates that the shape of the 3-parameter-lognormal distributions fits good to our data. To check whether this fit is good enough, we execute Goodness of Fit tests on the applicable probability distributions which results in the Anderson-Darling statistic (AD) and the p-value. The Anderson-Darling statistic indicates a good fit when the value of the statistic is low (preferably below one for statistical significance), and a p-value higher than $\alpha=0.05$ indicates a statistical significant fit to the data on the 95% confidence level. The p-value cannot be determined for 3-parameter lognormal distributions, due to restrictions in the software we use, however we can be relatively certain that we do not have a statistically significant fit on this data since the Anderson-Darling statistic gives a high value of 471.97.

Since we cannot find a perfect fit to the data and since Table 4.6 indicates significant differences in the arrival punctuality between some types of aircraft and their nature of flight, we expect that the difference in aircraft types (which automatically leads to differences in the origin of flights) significantly influences our distribution fit.

When we make the differentiation to the type of aircraft, we see that the 3-parameter lognormal distribution fits best for all aircraft. We illustrate this with the histograms for the arrival punctuality of different aircraft types (including a distribution fit) in Appendix F. The parameters of the distributions vary between the different types of aircraft. This gives for some aircraft types a reasonable good fit, and for some a poor fit. We display the distribution parameters and the results of the Anderson-Darling statistic in Table 4.10.

		Parameters				*Test values*
Type	**Theoretic Distribution**	**Location**	**Shape**	**Scale**	**Threshold**	**AD**
Airbus330	3-parameter Lognormal	4.210		0.420	-68.940	9.958
Boeing737	3-parameter Lognormal	5.098		0.142	-163.100	281.916
Boeing747	3-parameter Lognormal	4.878		0.233	-133.700	6.838
Boeing767	3-parameter Lognormal	4.081		0.488	-60.040	3.364
Boeing777	3-parameter Lognormal	4.378		0.300	-86.850	8.528
Embraer90	3-parameter Lognormal	4.067		0.275	-63.940	14.731
Fokker100	3-parameter Lognormal	3.869		0.428	-50.220	6.368
Fokker70	3-parameter Lognormal	5.250		0.104	-192.000	127.056
MD11	3-parameter Lognormal	4.330		0.294	-84.020	2.941
Other	3-parameter Lognormal	5.250		0.253	-173.000	5.024

Table 4.10: Overview of the best fitted theoretic distribution on the arrival punctuality, per type of aircraft.

We conclude from Table 4.10 that we obtain a better distribution fit when we split the dataset to different types of aircraft since the AD-values are lower than the first obtained 471.97. However, the type of distribution stays the same (3-parameter Lognormal) and the parameters do not vary in a very large extent.

The fit of the 3-parameter lognormal distribution to the arrival punctuality per aircraft type is reasonable good, but due to the high AD-values we can reasonably assume that the fit is not statistically significant. This can be partially explained by the large amount of data. When the dataset gets larger, the size of the confidence interval decreases. This leads to an increasing probability of the test value falling outside the confidence interval, and thereby indicating a difference (which can be very small) between the data and the theoretical probability distribution.

We recommend using theoretical probability distributions, since we concluded from the visual comparison in Figure 4.15 and Appendix F that it is possible to obtain a good fit. When management does not want to use this theoretical probability distribution, then it is also possible to choose for an empirical distribution function or direct data usage. The empirical distribution function is the most appropriate function (based on the arguments of Table 4.9), since we want to be able to generate an unlimited amount of input data to be able to perform many simulation runs. Irregularities in the data resulting from limited observations are not a big issue, since our data set contains many values over the entire range.

4.5.2 Refuelling process

The distribution of the total process times of the refuelling of aircraft depends on the duration of the refuelling task plus the duration of possible disturbances. In this section, we fit theoretic probability distributions to the historical data of June 2011 for both the duration of disturbances in the refuelling process and the process times of the refuelling tasks. For the process times of the refuelling tasks we make a differentiation on the type of aircraft and on the type of refuelling task.

Disturbances
When we analyse the data of Table 4.7 and a histogram of the duration of disturbances, we conclude that the duration of disturbances (when one occurs) is best described using a lognormal distribution. The histogram of Figure 4.16, with a Lognormal fit plotted in the figure, states this conclusion. We give the parameters of the distribution in the legend of the figure.

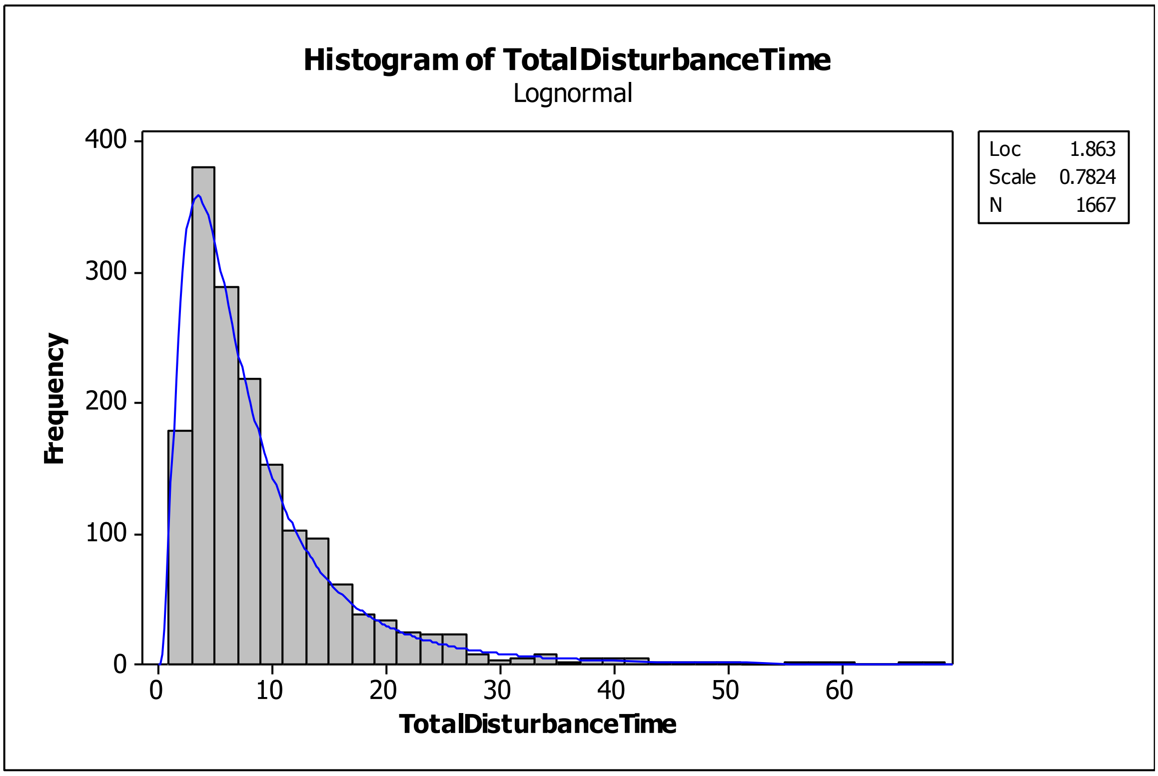

Figure 4.16: Histogram of the disturbance times in the month June 2011, including a plot of the lognormal distribution.

The distribution fit we made in Figure 4.16, is the best possible based on this data. However, when we execute the Goodness of Fit tests we obtain a p-value smaller than α=0.05 which means that the fit is not statistically significant. The AD-value is 5.708, which is relatively low. This means that the deviation from the optimum is small. We conclude that the distribution fit is good and usable, because when we make a visual comparison, the theoretical probability distribution fits well to the data.

Refuelling process times
To fit a theoretic probability distribution to the refuelling process times, we split the dataset of June 2011 to the different general aircraft types and the three types of refuelling task as defined in Figure 4.13. It is necessary to split to the different types of aircraft because of the difference in fuel intake capacity per aircraft. We give the best fitted theoretic probability distributions, including the parameters, in Table 4.11. This table also gives an overview of the p-value for statistical significance and the outcome of the Anderson-Darling Goodness of Fit Test.

Prefuel		*Parameters*				*Test values*	
	Theoretic Distribution	**Location**	**Shape**	**Scale**	**Threshold**	**AD**	**P-value**
Boeing747	3-parameter Lognormal	5.031		0.0814	-92.280	3.641	-
Boeing767	*						
Boeing777	Gamma		15.44	2.7790		0.538	0.187
MD11	3-parameter Lognormal	3.477		0.2193	7.294	1.132	-

Final Fuel (after pre-fuel)		*Parameters*			*Test values*	
	Theoretic Distribution	**Location**	**Shape**	**Scale**	**AD**	**P-value**
Boeing747	Normal	42.92		9.871	0.384	0.389
Boeing767	*					
Boeing777	Normal	38.13		10.31	0.319	0.532
MD11	Normal	42.39		12.24	0.410	0.335

** We were unable to fit a theoretic probability distribution, because there were only two pre-fuel tasks for the Boeing 767 in the month June 2011.*

Final Fuel (no pre-fuel)		*Parameters*				*Test values*	
	Theoretic Distribution	**Location**	**Shape**	**Scale**	**Threshold**	**AD**	**P-value**
AirbusA330	Gamma		16.030	3.912		0.837	0.033
Boeing737	3-parameter Lognormal	2.938		0.317	3.547	15.695	-
Boeing747	3-parameter Lognormal	5.431		0.056	-150.800	1.315	-
Boeing767	3-parameter Lognormal	6.111		0.026	-392.700	1.291	-
Boeing777	3-parameter Weibull		5.329	76.820	-10.280	0.163	>0.500
Embraer90	3-parameter Lognormal	3.343		0.194	-8.043	4.638	-
Fokker100	3-parameter Lognormal	3.124		0.237	-3.942	0.686	-
Fokker70	Lognormal	2.807		0.260		11.104	<0.005
MD11	Normal	60.170		12.550		0.304	0.565
Other	3-parameter Weibull		4.878	38.790	-8.263	0.311	0.487

Table 4.11: Overview of the best fitted theoretic distribution per type of fuel task and per type of aircraft.

We give the histograms of the refuelling process times in Appendix F to give an illustration of the fitted theoretic probability distributions compared to the historical data. A statistical significant fit could not be found for every type of aircraft. However, since the visual comparisons from Appendix F indicate a good fit, the theoretical probability distributions can very well be used.

When KLM wants to implement this research, it is good to also differentiate on the destination of flight or the distance to fly (in addition to the type of aircraft), to make a better distinction between the amount of fuel needed and thereby the duration of refuelling. In this project we do not execute these differentiations due to time limitations.

4.6 Conclusion

We showed in this chapter a lot of different sources of uncertainty and unforeseen events. The magnitude of the uncertainty varies, for the arrival punctuality of aircraft, over the time of the day and over different periods of the year. For the refuelling service, the process time differs a lot over different aircraft types and per type of refuelling task. For further research, we recommend to also differentiate to the destination of flight. The refuelling process is disturbed in 15.84% of all tasks.

We conclude that it is possible to fit theoretic probability distribution to the arrival punctuality of aircraft, the refuelling process times and the duration of disturbances. The visual comparison between the historical data and the used theoretical probability distributions, indicate a good fit. Due to the large amount of data we use, obtaining statistical significant fits is hard, but this does not undermine the usability of theoretical probability distributions for the goal of this research.

Aircraft Services has to investigate the possibilities of differentiating the uncertain variables to a lower level of detail, when it chooses to implement the results of this research. This prevents the probability of having multiple uncertain variables in one distribution, which makes it difficult to fit a theoretical probability distribution to the dataset. Differentiation is possible to a very small level of detail since there are a lot of factors influencing the processes and thereby a lot of variables to differentiate on. A disadvantage of differentiation, is that the complexity of the problem and the solution increases fast, giving a lot of different probability distributions for all different situations. This may threaten the robustness of the solution.

5 Robust capacity planning at the refuelling process

To determine the effect of a more robust capacity planning on the performance of the refuelling process, we need to incorporate uncertainty and unforeseen events into the capacity planning of KLM and then determine the resulting performance. The outcomes of this chapter give managers the opportunity to take well-founded decisions about how much operators to deploy when certain performance targets need to be achieved.

5.1 Overview of the different models

In this chapter, we use four different models to get from a predefined flight schedule, extended with variances in the processes caused by uncertainty and unforeseen events, to the corresponding performance of the refuelling process. We first create a capacity planning using model 1 to 3, and afterwards simulate a day of operation using model 4. We explain the function of, and the connection between, the different models in this section, following Figure 5.1. A detailed explanation per model and the outcomes follow in Section 5.2 till 5.5.

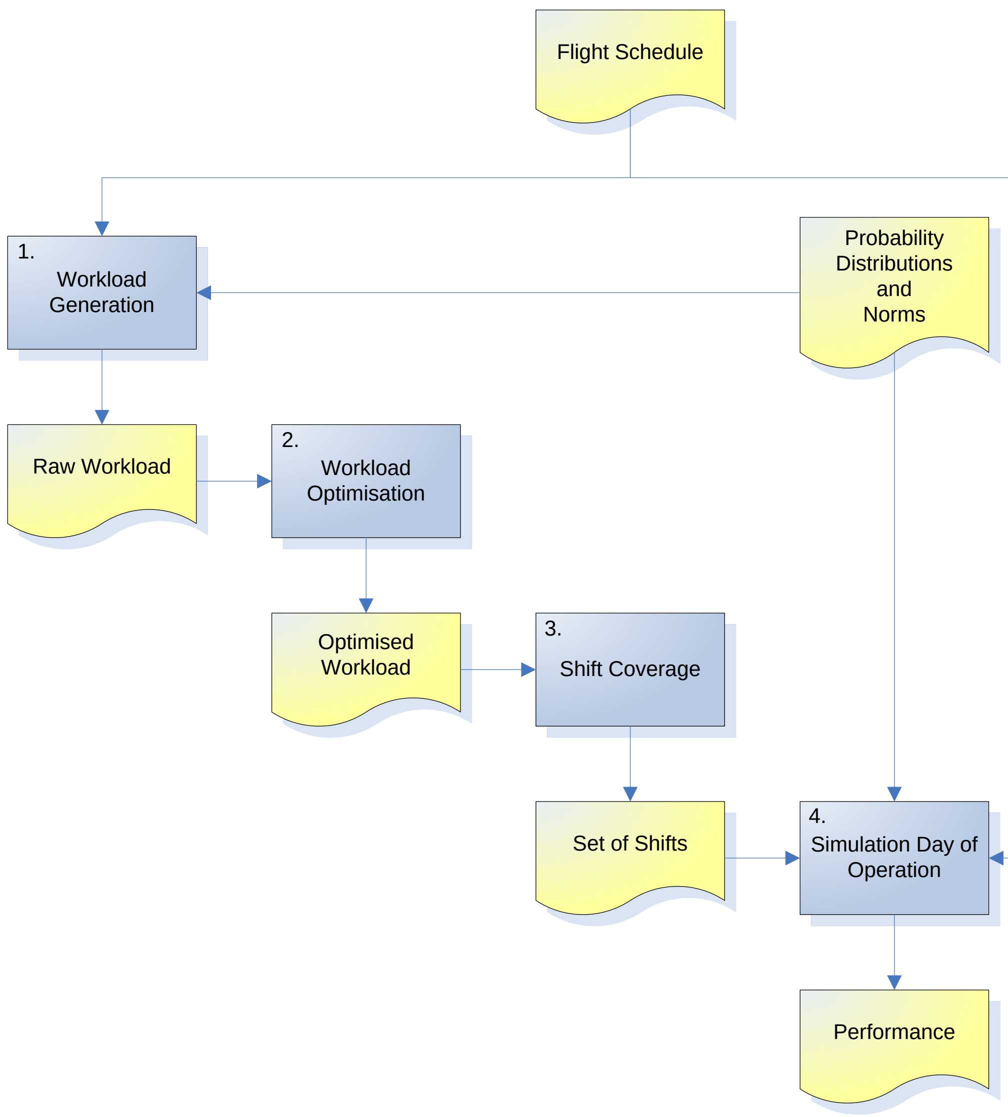

Figure 5.1: The correlation between the different models.

1. Workload Generation
The first step is to generate multiple workload lines using the probability distributions we described in Chapter 4 and the flight schedule. These workload lines indicate the number of employees needed to execute the planned tasks at a certain moment in time (see the definition in Chapter 1). The lines differ from the current situation due to variation in the arrival time of aircraft, in the process time of the refuelling process and due to possible disturbances during refuelling. We create multiple workload lines using a Monte Carlo Simulation, which we program in VBA using Microsoft Excel. The different steps of the Monte Carlo Simulation are explained in depth in Section 5.3.1

This model uses the flight schedule as input and transforms the defined arrival times by adding or subtracting an arrival punctuality defined by the probability distributions from Chapter 4. We then use the planning norms, as defined in Appendix C, for assigning earliest start and end times. We also generate the duration of refuelling tasks and possible disturbances in the refuelling process, according to the probability distributions as defined in Chapter 4. Combining this gives, as output, a set of task with an earliest start and a latest end time.

2. Workload Optimisation
The goal of KLM is to execute all tasks over the day with a minimum number of operators, so the second modelling step is to optimise the workload from our previous model in such a way that the peaks are minimised. Minimising the peaks is necessary since operators work in shifts which last eight hours. This causes them to be present for a longer time than just the peaks and thereby employing operators to cover the workload in the peaks leads to a lot of excess capacity in the off-peak moments (and thereby high costs).

Optimisation is in this case possible, since every refuelling task works with flexible permissible time windows. This means that the difference between the latest end time and the earliest start time is in general larger than the total refuelling process time, giving a window for optimisation.

This optimisation determines the optimal starting time for every task within its permissible time window of execution. We optimise the starting times of all tasks in such a way that there are as less as possible simultaneously executed tasks per time interval. This minimises the peaks in the workload. A detailed description of the optimisation can be found in Section 5.3.2.

After the optimisation, we aggregate the resulting workload lines to a percentile picture. This picture gives the bandwidth of the distribution of the workload lines. Generally explained, the x^{th} percentile is defined as the smallest value that is greater than or equal to x% of the values. The 50^{th} percentile is also called the median (Lane, 2010). The output of this model step is the workload profile of Figure 5.5 and the dataset used to create this.

3. Shift coverage
In an ideal situation, we would be able to change the number of operators every minute. In reality, this is not possible. KLM schedules her operators to work in shifts of eight hours. We need to fit a set of these shifts to the optimised workload to schedule personnel. The goal of our third model is to compose a set of shifts, covering the workload in such a way that the total number of shifts is minimised and thereby the personnel costs.

One could say that covering, for example, the 80th percentile of workload gives a performance of 80%. However, due to the excess capacity generated by the eight hour shifts used to cover this percentile, the achieved performance is much higher. A shift does not cover only the peaks since the shift lasts eight hours which causes operators to be present longer than just the peak times, leading to excess capacity. This makes that covering different percentiles of workload does not have a linear relationship with the achieved performance.

4. Simulation Day of Operation
Our fourth and last model determines the effects of a more robust capacity planning on the performance of the refuelling process. For this, we use a discrete event simulation executed in Siemens Technomatix Plant Simulation.

Using this model, we simulate a day of operation, where we use the planning created in modelling step 1 to 3. The optimised set of shifts is now fixed and predetermined by the model from step 3. The starting times of the refuel tasks are determined by the simulation and thereby completely independent from the starting times of step 2 (used to create a planning) and only bounded by the earliest start and latest end time of the task. The model determines when to start a task, based on a set of predefined rules. At the end of the simulation, we compare the performance of the capacity planning which incorporates uncertainty and unforeseen events with the old deterministic planning.

5.2 Assumptions and modelling choices

To keep the models manageable, we have to make some assumptions and choices which we use in all models. Some of these are predefined in the KLM planning process (Section 5.2.1) and some are our own modelling choices (Section 5.2.2). We evaluate the impact of assumptions at the end of this chapter.

5.2.1 KLM planning choices and assumptions

1. When an aircraft arrives late, KLM checks whether it still has the minimum ground time needed for a turnaround. If the available ground time is considered too short, the departure of the aircraft is delayed to increase the ground time to the minimum needed, without letting this influence the task performance of one of the ground processes. When an aircraft is delayed by an extended refuelling task, this does influence performance. In reality, the other ground processes, as defined in Chapter 2, can also delay the aircraft, but we assume them to be non-existent. For this reason we cannot measure the overall performance of AS. We measure purely the task performance of refuelling when using a predefined number of operators.
2. We follow the norms when assigning pre-fuel task to the different aircraft, although in reality these norms are not always followed. In reality a pre-fuel is only assigned by a controller when he has some spare capacity. When we follow the norms, a pre-fuel is assigned when an aircraft has a ground time longer than a predefined border (depending on the aircraft type).
3. We do not use exact driving times for the refuelling vehicles because this highly depends on the gate allocation of Schiphol. We previously explained that this is not under the influence of KLM and thereby very hard to plan. As a simplification, we add to every task the average driving time (5.5 minutes) as defined in the norms research report made by Visser (2010). Adding a fixed number of minutes to the norm duration of refuelling is the current way of dealing with driving times at AS, but this does not represent reality where

the driving time is subject to the traveling distances between gates plus a certain probability distribution for possible delays. We use here the same method as used by KLM, but this only approximates reality roughly.

4. We use the shifts as defined by the resource planners of AS to create a coverage of workload. These shifts all have a duration of eight hours, plus 30 minutes for lunch, minus 20 minutes for personal hygiene. We do not allow the operators to work in overtime. An overview of the shifts is given in Table 5.1.
5. We measure the performance of the refuelling process based on the task performance definition set in Section 1.2: "the percentage of tasks finished before their latest end times as defined in the norms." Due to KLM guidelines, we only consider final fuel tasks in this performance measure.

5.2.2 Modelling choices and assumptions

6. We assume that the refuelling process is completely independent from all other ground processes, which we do not model. However, when the refuelling process is delayed by other ground processes, this is incorporated in the disturbances.
7. We draw our values from the probability distributions, as defined in Chapter 4 (arrival punctuality, refuelling process times and disturbance times), independent from each other. In reality, these samples are correlated. For example: when one aircraft is delayed on a day with bad weather, it is very likely that the other aircraft are also delayed. Notice that these dependencies are assumed to be non-existent in our models. Resulting from this, our models are only valid for relatively normal situations.
8. We generalise the different aircraft types again to their general aircraft type (in the same way as we did in Chapter 4). We also generalise the norms to the norms of the 3 big airliners at Schiphol: KLM, Transavia and Delta, since these airliners are responsible for about 98% of all departing flights (see Table 4.3). We do make a distinction between the norms of these 3 airliners, when they occur.
9. We keep the flight links, as defined in the flight schedule. In reality, these links are sometimes changed when aircraft have big delays. KLM can in this case change the assignment of aircraft to the different flights, which influences the ground time (as indicated in Chapter 4). The linking of flights is in the current deterministic situation not an issue, because this planning does not incorporate delays in the arrival process. This means that in the current situation, the flight links are always optimal as defined in the flight schedule.
10. We define the earliest start and latest end times of a refuelling task always as defined in the planning norms (see Appendix C), split per general aircraft type, where the class "Other" always starts at the arrival time plus 10 minutes.
11. We model our workload without breaks. Every operator needs at least a 30 minute lunch/dinner break and, when time available, they also get a 15 minute coffee break. However, KLM cannot give us exact rules for when to schedule these breaks. For this reason we do not model breaks, but one should notice that this leads to an underestimation of the total workload.
12. Our models use just one type of equipment and do thereby not make the difference between fuel-dispensers and bowsers (refuel trucks). In some cases, operators are changed over the day from a fuel-dispenser to a bowser (or vice versa). This is also not incorporated in our models.
13. We described in Section 4.5 that the refuelling process is disturbed in 15.84% of all refuelling tasks. We thereby let a disturbance occur in 15.84% of all refuelling tasks, according to a uniform distribution. We define the duration of these disturbances according to the lognormal distribution fitted to the historical data in Figure 4.16.

5.3 Workload generation model

In this section we describe model 1 in Section 5.3.1 and model 2 in Section 5.3.2 (see Figure 5.1). We create these models using VBA in Microsoft Excel, which results in a workload profile which incorporates uncertainty and unforeseen events, aggregated per 10th percentile.

5.3.1 Model design

We take several steps to generate multiple workload lines by executing a Monte Carlo Simulation. In this section, we explain these steps following the flowchart of Figure 5.2.

When initialising this model, we can choose to run the model with or without variation. When the probability distributions are disabled, the model uses the scheduled arrival times, the norm duration for the refuelling process and there are no disturbances.

When starting the Monte Carlo Simulation, one has to choose the number of iterations. Every iteration is an extra workload line, so for a robust solution, the number of iterations needs to be high. We use 1000 iterations and create thereby 1000 different scenarios for the course of workload over the day. This is the equivalent of about 3 years of operation.

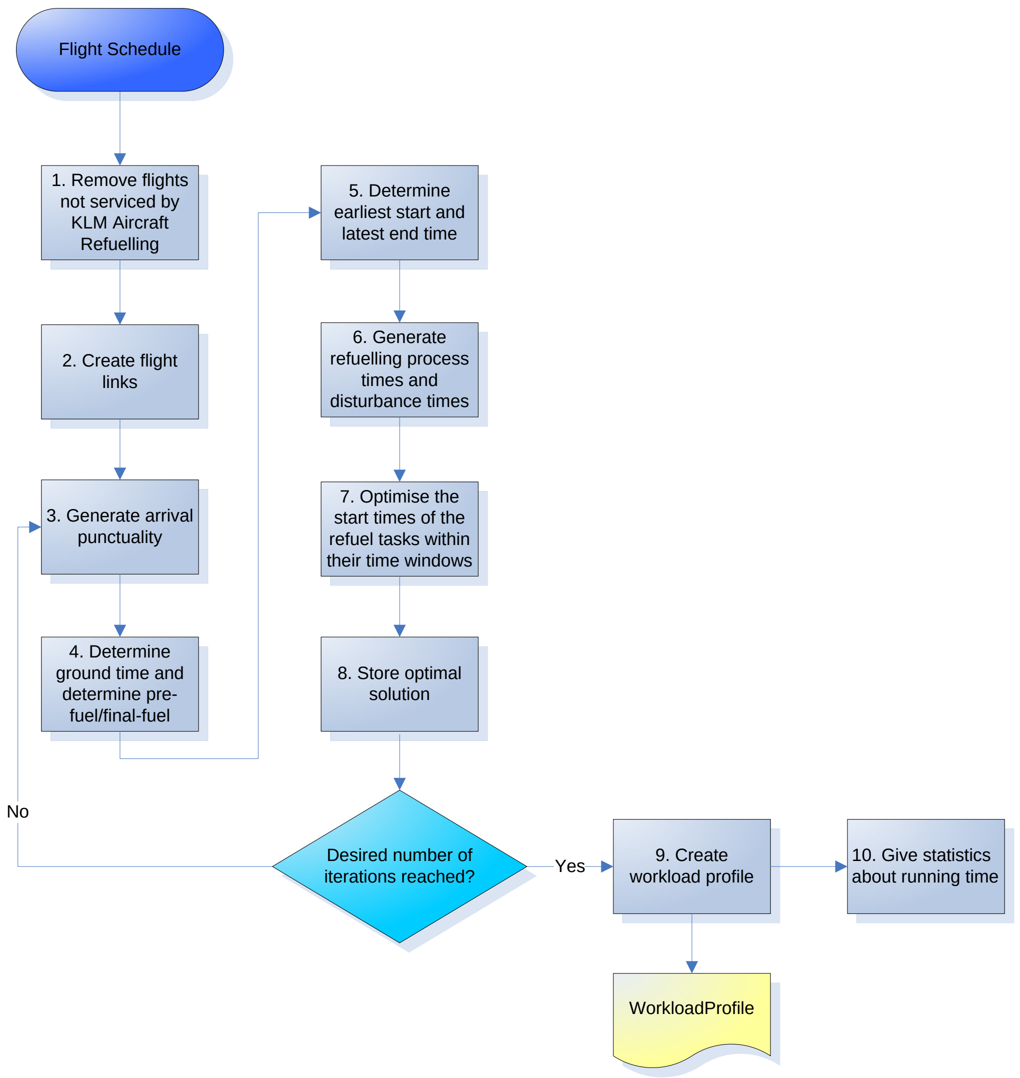

Figure 5.2: A flowchart of the workload generation model.

1. Remove flights not serviced by KLM Aircraft Refuelling
Since we receive a flight schedule which incorporates all flights arriving at Schiphol, we need to remove the flights which are not serviced by KLM Aircraft Refuelling. These are all flights from airlines which are not in Table 4.3.

2. Create flight links
The next step is to transform the flight schedule to an overview of the flight links: a couple of an arriving and a departing flight, executed by the same aircraft. These flight links all have a scheduled arrival and departure time.

3. Generate arrival punctuality
We assign to every flight link an arrival punctuality based on the 3-parameter lognormal distribution fitted to the historical data in Figure 4.15. After assigning the arrival punctuality, we compute the actual arrival time by adding the arrival punctuality to the scheduled arrival time.

4. Determine ground time and determine pre-fuel/final-fuel
By subtracting the actual arrival time from the departure time, we get the ground time of the aircraft. Using this, we can, according to the norms as attached in Appendix C, determine whether we need to schedule a pre-fuel for the flight. When a pre-fuel is needed, we add this as an extra refuelling task.

5. Determine earliest start and latest end time
For every refuel-task, we determine the earliest start and latest end time according to the planning norms as attached in Appendix C. We set, as an initial solution, the start time of every refuel-task to the earliest start time.

6. Generate refuelling process times and disturbance times
We generate a process time for every refuelling task based on the probability distributions defined in Table 4.11. We also generate disturbances in 15.84% of all refuelling tasks, according to a uniform distribution. For these disturbances we generate a disturbance time based on the lognormal distribution fitted to the historical data in Figure 4.16. The total duration of refuelling is the refuelling process time, plus a possible disturbance time, plus 5.5 minutes (which is the average driving time of a refuel vehicle between two aircraft).

7. Optimise the refuel tasks within their time windows
We explain this in Section 5.3.2.

8. Store optimal solution
Save the optimal solution. If one wants to execute more than one iteration, to create multiple workload lines, go to point 3.

9. Create workload profile
Create a percentile picture based on the saved optimal solutions from point 8. In this picture, we aggregate the different workload lines and display them per 10^{th} percentile. We display this in Figure 5.5.

10. Give statistics
We display running time statistics: start time, end time and duration, to be able to schedule model runs. We also display statistics about the number of changes executed in the optimisation, which is used for model validation.

5.3.2 Workload optimisation

The goal of KLM is to execute all tasks over the day with a minimum number of operators. To minimise the number of operators needed, the peaks in the workload have to be minimised. Peaks in the workload are very disruptive since KLM Aircraft Refuelling uses shifts of about 8 hours to schedule her operators. This causes operators to be present for a longer time than just the peak moments. Employing operators to cover the workload in the peaks leads therefor to a lot of excess capacity in the off-peak moments and high costs.

KLM Aircraft Refuelling works with flexible permissible time windows which gives the possibility to optimise the workload. The goal of the optimisation is to set the start moment of every refuelling task in such a way that the number of simultaneously executed tasks per time interval is minimised. This minimises the peaks in the workload. When the workload is optimised, KLM can make effective use of the available ground time of an aircraft, using a minimum amount of capacity in terms of equipment and operators. We lower the peaks in the workload by minimising the following objective function:

$$\min_{s \in S} Z = \sum_{t=1}^{T} (W_t)^3$$

W_t = the workload at time interval t
T = the number of time intervals over a single day
S = the set of all possible starting times for all refuelling tasks
s = a subset of S, containing a possible optimal solution, with for every task a starting time

This formulation dictates that we minimise the value of the objective function by searching the best set of possible starting times (s) out of the complete set with possible starting times (S) for all refuelling tasks over a single day. To determine the value of the objective function, we elevate the workload per time interval (W_t) to the third power and then sum over all time intervals (T). We elevate the workload to the third power to give more weight to high workload values and thereby force the function to lower the peaks in the workload. We select the third power empirically, by experimenting with the second to the fifth power. From the third power and higher, the outcome is good and stable over the different powers, so we choose the lowest giving good results: the third power.

We use the Simulated Annealing heuristic to determine the start time of every refuelling task within its time window. Simulated Annealing is a heuristic algorithm originated from physics, which obtains good, but not necessarily optimal, solutions to optimisation problems. The heuristic starts as a random search algorithm and ends as a local search algorithm. According to Eglese (1990), the algorithm avoids to get trapped in a local optimum, by in some cases accepting a neighbour solution which worsens the value of the objective function. When the algorithm gets in a later state, it only accepts improvements. We add to this a modification suggested by Glover and Greenberg (1989), which stores the best solution obtained during the optimisation. With this modification, there is less need for the algorithm to rely on a strong stabilizing effect at the end of the runtime to make sure the optimum reached is the best solution and not a local optimum.

The Simulated Annealing heuristic uses our previously defined objective function and the following steps to come to a good solution:

Initialise:

- Load the initial situation as set s (see step 5, Section 5.3.1)
- Determine the value of the objective value

Loop till end criterion:

- Start with set s
- Pick a random refuelling task
- Update set s by determining a random new starting time within the permissible time window. We only accept the new start time if the expected end time (the start time plus the expected total duration of refuelling) is less than the latest allowed end time, so we need a value from set S (the complete set of allowed starting times).
- Determine the value of the objective function based on set s
- If the value of the objective function is lower than the previous value, accept the new solution and save set s. Also save this set as the best solution so far. If the value of the objective function is higher or equal than the previous value, accept the new solution with a certain probability, as defined by the Simulated Annealing parameters (see Section 5.3.3), and save set s. In all other cases: undo the changes in set s.
- Loop

At the end of the heuristic:

- Compare the current solution with the best solution so far. Choose the solution with the lowest value of the objective function. The set s belonging to this value is the result of the heuristic. Figure 5.3 illustrates the result of the heuristic. The blue line is the workload before optimisation and the red line is this workload after using the Simulated Annealing heuristic.

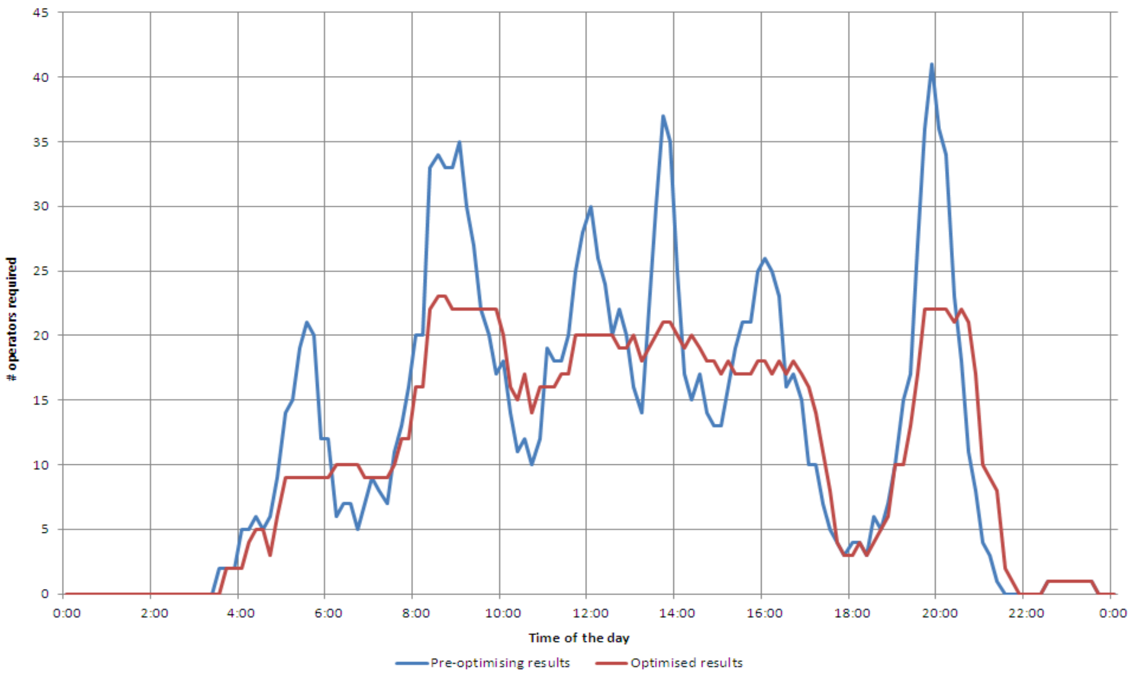

Figure 5.3: The effect of optimising the workload.

Using a heuristic is necessary, since our problem can be simplified to the Travelling Salesman Problem (TSP), which is known to be NP-hard and that its formulation is NP-complete (Wilf, 1994). Dantzig et al (1954) define the TSP as: *"Find the shortest route (tour) for a salesman starting from a given city, visiting each of a specified group of cities, and then returning to the original point of departure."* In our case, the cities can be seen as the different aircraft which need refuelling, where the refuelling process time can be seen as the travelling distance between aircraft. The salesman can be seen as a fuel-operator who has to service all the aircraft. When we use more than one fuel-operator, the TSP is extended to the multi-TSP, which works with multiple salesmen. Our problem is extra complex due to the time windows for servicing aircraft. This makes the problem comparable to a 'multi-TSP with time windows'.

We can also show that our problem is NP-hard by analysing the size of the problem. In our case, we have X tasks and Y start-moments per refuelling task. This gives a solution space of Y^X possibilities and thereby the problem size grows exponentially with the number of task to be executed. Solving this problem to optimality is almost impossible in polynomial time (Eglese, 1990).

5.3.3 Validation

Before we use the results of our model, we need to validate whether our model works well with the given input parameters. The Simulated Annealing heuristic works with a cooling schedule, with, according to Eglese (1990), the following control parameters: the start temperature (C_0), the end temperature (C_{stop}), the cooling factor (α) and the length of the Markov Chain (k). We work with the following setting: C_0 = 20, C_{stop} =0.001, α = 0.95, k = 100. This setting gives good results in a reasonable amount of time. We create Figure 5.3 to demonstrate the results. This figure gives the results of 4 runs of model 1 and 2 (see Figure 5.4) combined, where the probability distributions are disabled and thereby there is no variation modelled in the processes. We can see that the 4 runs are optimised to almost the same solution, so we conclude that the algorithm gives a robust solution using our predefined settings.

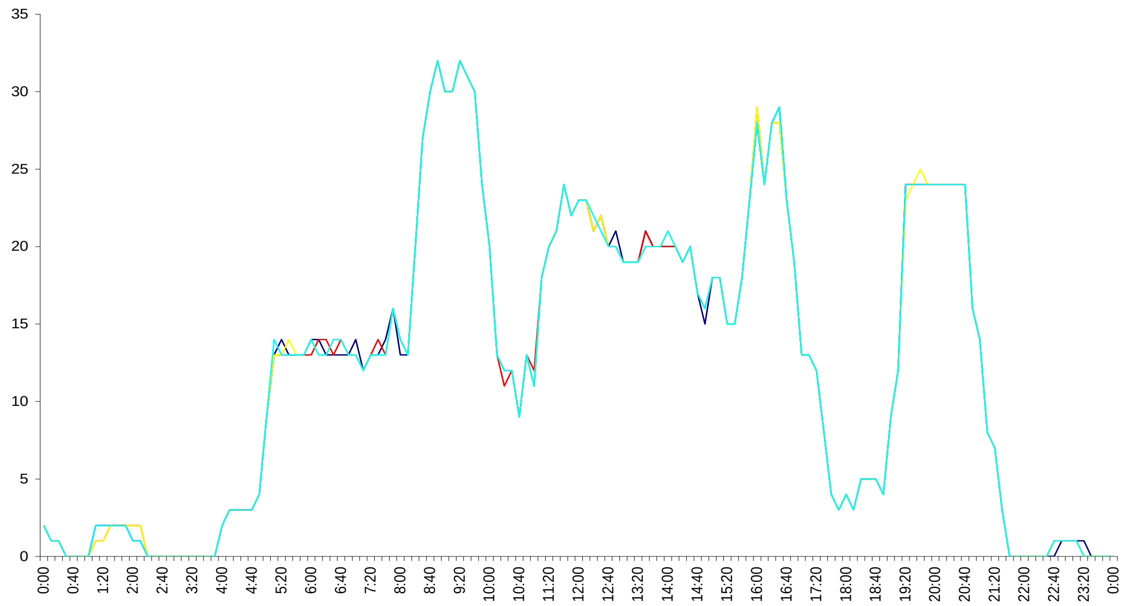

Figure 5.4: The result of 4 runs of the workload generation heuristic with no variance modelled in the processes.

For further validation of the output, we cannot simply execute a statistical comparison, because we built our model based on data from the day of execution and not the forecasts from which the workload profiles of KLM origin. Furthermore, we also made assumptions affecting our model, which are given in Section 5.2.

Since a quantitative validation is not possible, we need to rely on expert opinions. We discussed the outcomes of our model with management and compared the peaks in our workload profiles with the peaks in Figure 4.1 (which describes the number of arrival and departures). These two validation steps both indicate that our model is valid. The input distributions are a good representation of reality as shown in the analysis of Chapter 4, where we first obtained clean data and then determined and validated theoretical probability distributions for the different processes. We conclude from this that our model is valid.

5.3.4 Model results

The output of the Monte Carlo Simulation using our optimised workload generation model is the amount of operators required per ten minute time interval over a day. Here the start time of every task is optimised within its time window. We give a percentile picture of the results in Figure 5.5. This is the result of a Monte Carlo Simulation with 1000 runs, thereby creating 1000 different workloads which differ based on the probability distributions of the arrival punctuality, refuelling process times and possible disturbances in the refuelling process. We also add a line indicating the current deterministic situation (without uncertainty and unforeseen events), where all this variation is not taken into account.

The peaks in especially the morning and evening occur due to short turnarounds, giving narrow time windows for the completion of a refuelling task. Narrow time windows do not give many opportunities to spread the workload and thereby for optimisation of the starting time of the task. For this reason, a peak arises when a lot of short turnarounds are needed at a certain moment in time.

We see that the largest bandwidth is situated at and around the peaks. The highest bandwidth is 15 employees, varying between 6 and 21. The bandwidth is that large, because peaks can move a bit over the x-axis due to changes in the arrival time of aircraft caused by the arrival punctuality. Since the difference between the 0th and the 100th percentile is relatively large, we conclude that uncertainty and unforeseen events have a big influence on the planned capacity.

The deterministic line lies generally at high percentile values. This means that using the current planning, the available capacity is sufficient for most situations. However, we have to take the assumptions we made in Section 5.2 into account. The outcome of the workload generation model (see Figure 5.5) is only valid for normal days (not days which are extremely disturbed). We removed extreme values from the data when we defined the probability distributions, since one should not plan for extremely disturbed days. An extremely disturbed day due to bad weather leads to all fuel operators working at a slower pace and many delayed flights, or for example the formation of an ash-cloud due to a volcano eruption leads to almost all flights being delayed. This can lead to workloads higher than our 100th percentile. When a day is highly disturbed, the data gets positively correlated, so the assumption that we draw independent samples from the probability distributions is not valid in this case, and thereby the same accounts for our model.

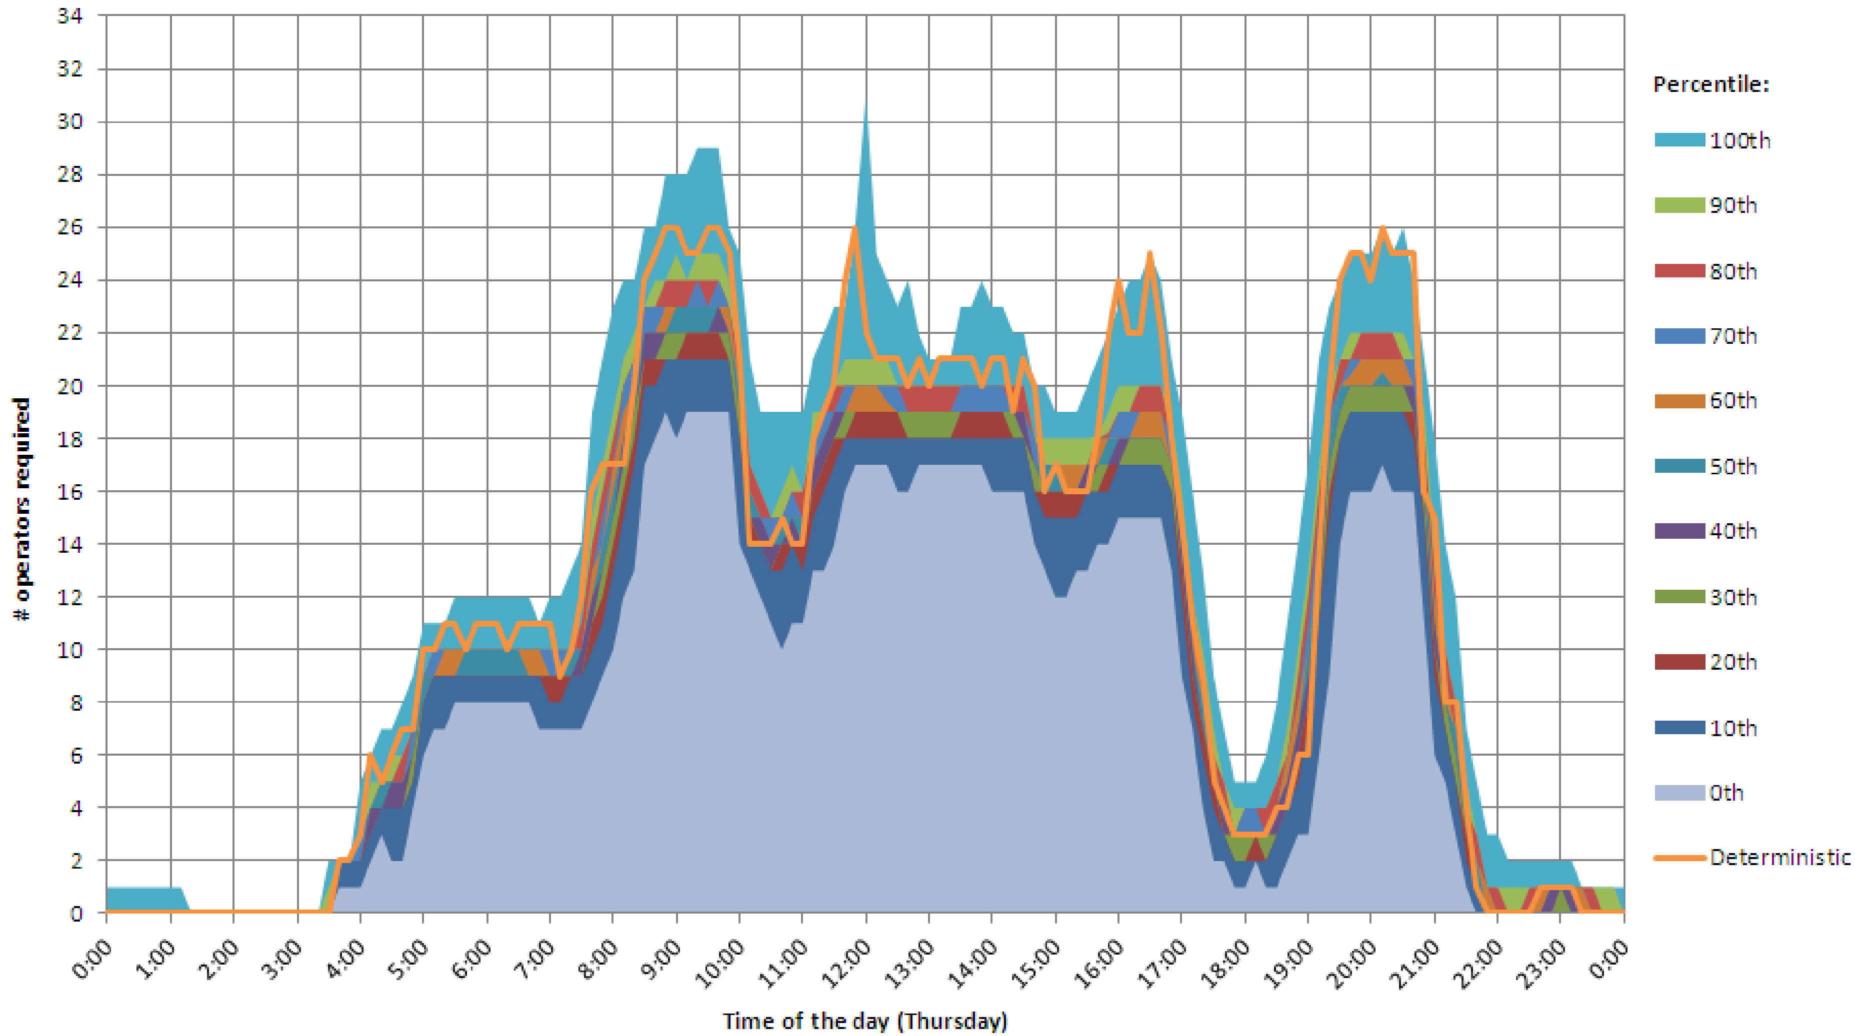

Figure 5.5: Percentile picture of the daily workload, made based on 1000 single workload lines resulting from the workload generation model. The orange line indicates the current deterministic situation (the workload without incorporating uncertainty and unforeseen events).

5.4 Shift optimisation

We need to cover the optimised workload from our previous model with a set of shifts to schedule personnel. In this section, we compose sets of shifts in such a way that the total number of shifts is minimised and thereby also the personnel costs.

The personnel of KLM Aircraft Refuelling works in eight hour shifts (plus 30 minutes for lunch, minus 20 minutes for personal hygiene). We give an overview of the shifts used by KLM in Table 5.1.

shift	from	till
N	22:30	6:40
SV	5:00	13:10
V	6:00	14:10
D	7:00	15:10
A1	13:00	21:10
A2	13:30	21:40
A3	15:00	23:10

Table 5.1: The shifts used at KLM Aircraft Refuelling.

We need to use a multiple of the shifts from Table 5.1 to cover the daily workload to schedule personnel. To create an optimal shift coverage of the workload, we optimise the deployment of the shifts to be as close as possible to the planned workload, on the levels of the in Figure 5.5 defined percentiles. The goal of the optimisation is to create an optimal covering for every percentile, while using a minimal amount of shifts. We execute this optimisation by solving the following LP-formulation, using the Simplex method as defined by Dantzig et al (1955), in Excel:

$$\min Z = \sum_{s=1}^{S} X_{st}$$

s.t.

$$\sum_{s=1}^{S} X_{st} \geq a_t \qquad \forall t$$

(The total number of operators must be larger than the amount of capacity needed at time t)

$$X_{st} = int$$

(All decision variables are integers)

X_{st} = *the number of shifts of type s planned at time t*
a_t = *the amount of capacity needed at time t, following from Figure 5.4*
t = *the ten minute time intervals (t = 1,…,144)*
s = *the shifts from Table 5.1 (s = 1,…,7)*

We solve this LP-formulation for every percentile of workload, as we defined in Figure 5.5 (the output of the workload generation model after optimisation). Table 5.2 gives an overview of the results, which is the optimal number of shifts required per type, needed to cover the planned workload over the day on a predefined level (percentile). We also add the amount of shifts needed over the day to cover the deterministic line from Figure 5.5, indicating the shifts needed when we do not incorporate uncertainty and unforeseen events and execute all processes according to their static norms. We use the outcomes of Table 5.2 as input to simulate a day of operation in Section 5.5, where we use these sets of shifts as input.

Percentile	N	SV	V	D	A1	A2	A3
0th	4	7	1	11	5	1	11
10th	5	8	1	12	5	1	13
20th	5	8	1	13	4	1	14
30th	6	8	1	13	5	1	14
40th	6	8	1	14	4	1	15
50th	6	8	1	14	5	1	15
60th	6	9	1	13	5	1	15
70th	7	9	1	14	4	1	16
80th	7	9	1	14	5	1	16
90th	7	9	1	15	4	1	17
100th	9	11	1	19	3	1	22
Deterministic situation	7	10	1	15	5	1	20

Table 5.2: The number of shifts used per type and per percentile of workload to be covered, including the number of shifts needed to cover the deterministic situation.

Figure 5.6 gives an overview of the total number of shifts needed over the day to cover the workload at a certain percentile. We conclude that we need a lot of extra shifts for covering the highest percentiles of workload. The green area gives the amount of shifts needed to cover the current deterministic situation. We conclude from this figure that the current situation lies at a relatively high percentile, since our graph crosses this line at about the 94th percentile.

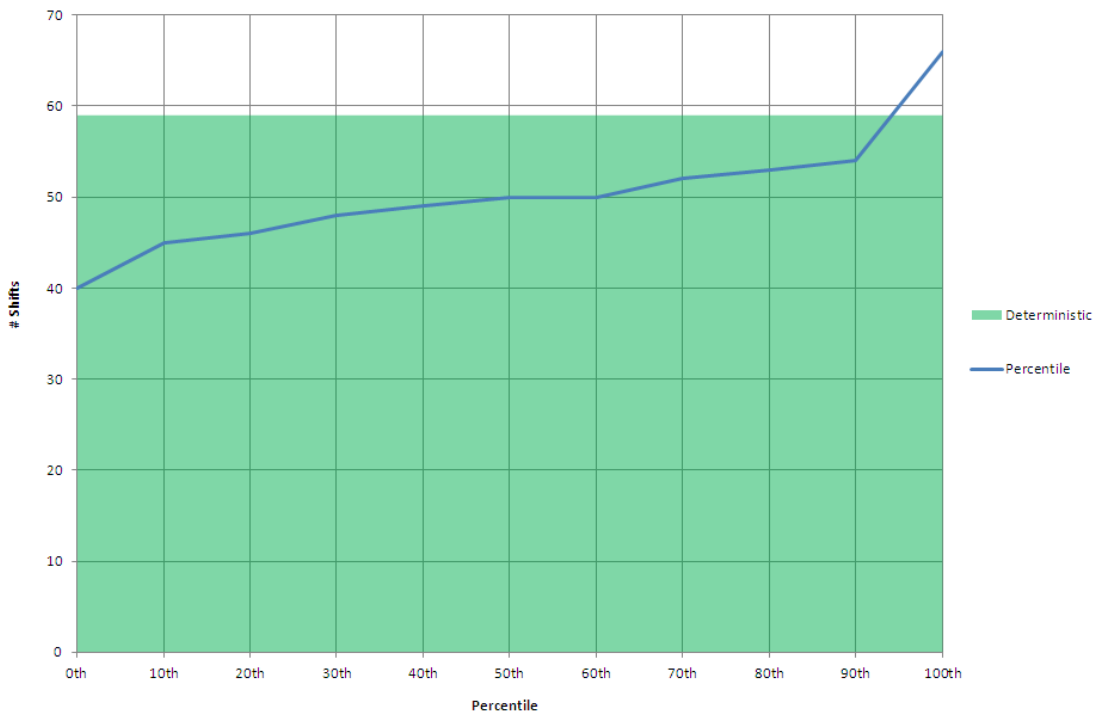

Figure 5.6: The total number of shifts required per day to cover the workload of the xth percentile, including the current deterministic situation.

Figure 5.6 shows that we can use less shifts if we choose to cover a lower percentile of workload. One operator costs KLM about €50,000 a year on wages. This operator works 210 days a year, so the cost of one shift is about €238 a day per operator. This is the amount of money we can save per day by reducing the number of shifts used by one shift. Notice, however that this is an indication, since in reality not every shift has the same costs. For example: a night shift is more expensive for KLM than a day shift.

5.5 Simulation Day of Operation

In this section, we translate the coverage of workload on different levels, to the performance of the refuelling process. We use the optimal number of operators per shift, as defined in Table 5.2, as input for a discrete event simulation model to obtain insights into the effects of a more robust capacity planning on the performance of the refuelling process. Using this model, we compare the performance of the capacity planning which incorporates uncertainty and unforeseen events with the current planning, the situation based on static norms, which is completely deterministic.

5.5.1 Model design

We create a discrete event simulation model to determine the effects of using different coverages of workload on the performance of the refuelling process. We describe this model in the flowchart of Figure 5.7. We also add Figure 5.8 and 5.9, which are screenshots of the model, to give an illustration of the model layout. In this section, we describe the model step by step, following the flowchart.

1. Load and update flight schedule

First, we load the flight schedule and update this with an arrival punctuality according to the probability distribution defined in Chapter 4. When this is finished, the simulation starts and aircraft start to arrive.

2. Determine the task type(s) and process time(s)

When an aircraft arrives, the model first determines the available ground time. Using this, we can, according to the norms as attached in Appendix C, determine whether we need to schedule a pre-fuel for the flight.

Afterwards we determine, according to the probability distributions as defined in Chapter 4, the process time for the final fuel, for a possible pre-fuel and for possible disturbances. The aircraft then moves to the “Store”: a buffer where the aircraft waits till it is allowed to be processed.

3. Pick an aircraft for refuelling

Due to KLM guidelines, we always prioritize final fuel tasks over pre-fuel tasks. When an aircraft is available in the Store and refuel capacity is available, we check whether we are allowed to start a final fuel. When multiple aircraft are available, we pick the aircraft with the earliest due date: the latest end time which comes first. When no final fuel is allowed to start, we check whether we are allowed to start a pre-fuel.

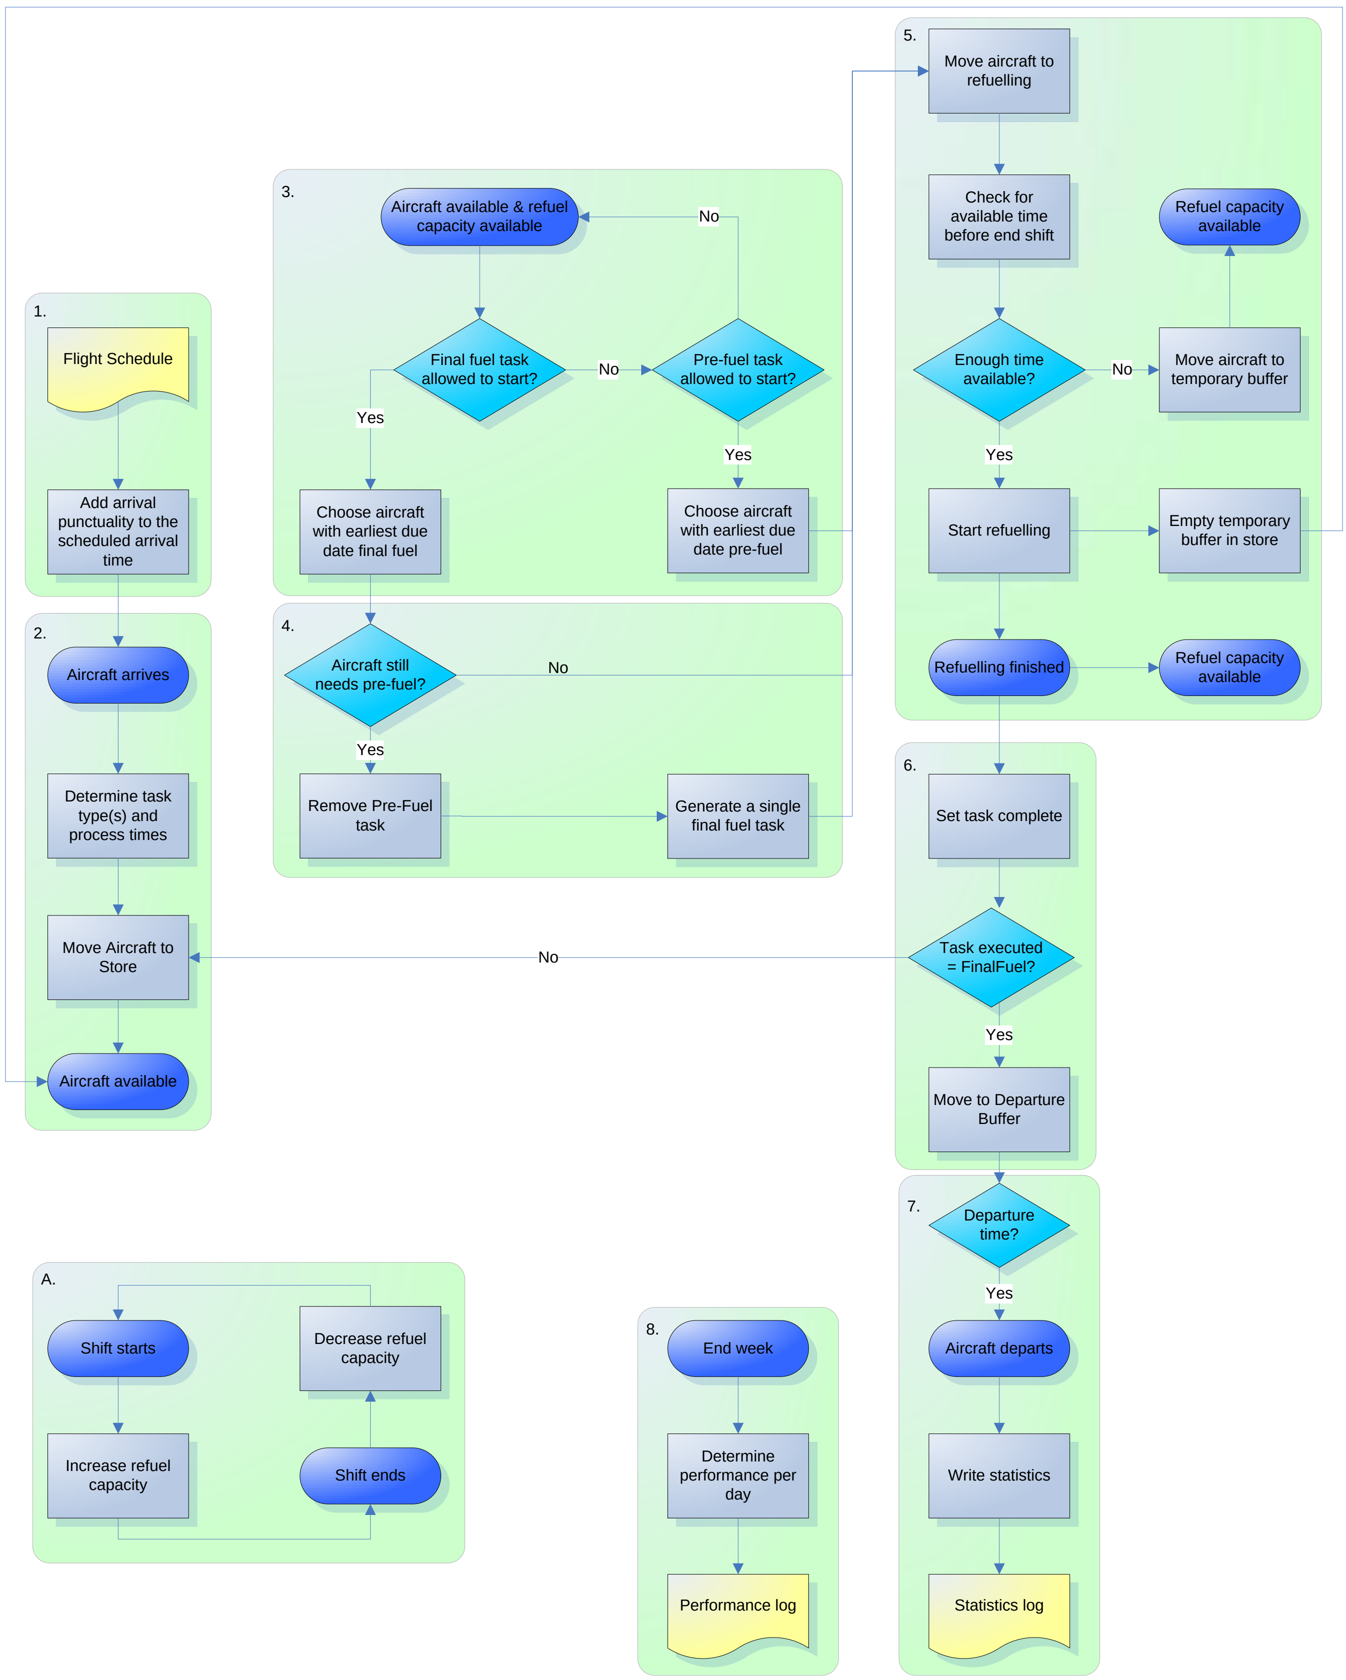

Figure 5.7: Flowchart of the discrete event simulation model.

4. Remove pre-fuel task when task is not executed
When we execute a final fuel, we check whether there was a pre-fuel scheduled for the aircraft. When this pre-fuel has not been executed, we cancel the pre-fuel task and use a single final fuel task, with a duration based on the probability distribution for single final fuel tasks (see Chapter 4).

5. Check availability within a shift
We move the chosen aircraft to refuelling and check whether there is time available before the end of one of the current shifts. If there is, we start refuelling. If not, we move the aircraft to a temporary flow buffer (which is only used to make sure we do not pick the same flight again) and check whether we are allowed to handle the aircraft with the next earliest due date. When we have found an aircraft which can be refuelled, we move all aircraft from the temporary buffer back to the Store. We also empty the temporary buffer every minute to prevent delays or a locked system due to flights trapped in this buffer.

6. When the refuelling task is finished, choose where to move the aircraft
When the refuelling task is finished, we set the task state to complete. If the task executed is a pre-fuel, we move the aircraft back to the Store. If the task executed is a final fuel, we move the aircraft to the Departure Buffer.

7. Allow the aircraft to depart
When the current time is equal to the departure time, we allow the aircraft to depart and write statistics about the start and end time of the refuelling task. These statistics are used later, to determine performance.

8. We determine the daily performance at the end of the week
We determine the daily performance of the refuelling process, based on the task performance as defined in Section 1.2: “the percentage of tasks finished before their latest end times as defined in the norms.” Due to KLM guidelines; we only consider final fuel tasks in this performance measure, since this is where the refuelling department is judged on.

A. Determine the available capacity
Table 5.2 defines the number of operators available over the day. We load this table as input to our simulation. At the start and end of a shift, we adapt our available capacity based on this table. Refuelling stations are only open when the shift operating the specific station is working. The refuelling stations can increase or decrease their capacity when the number of operators per shift changes.

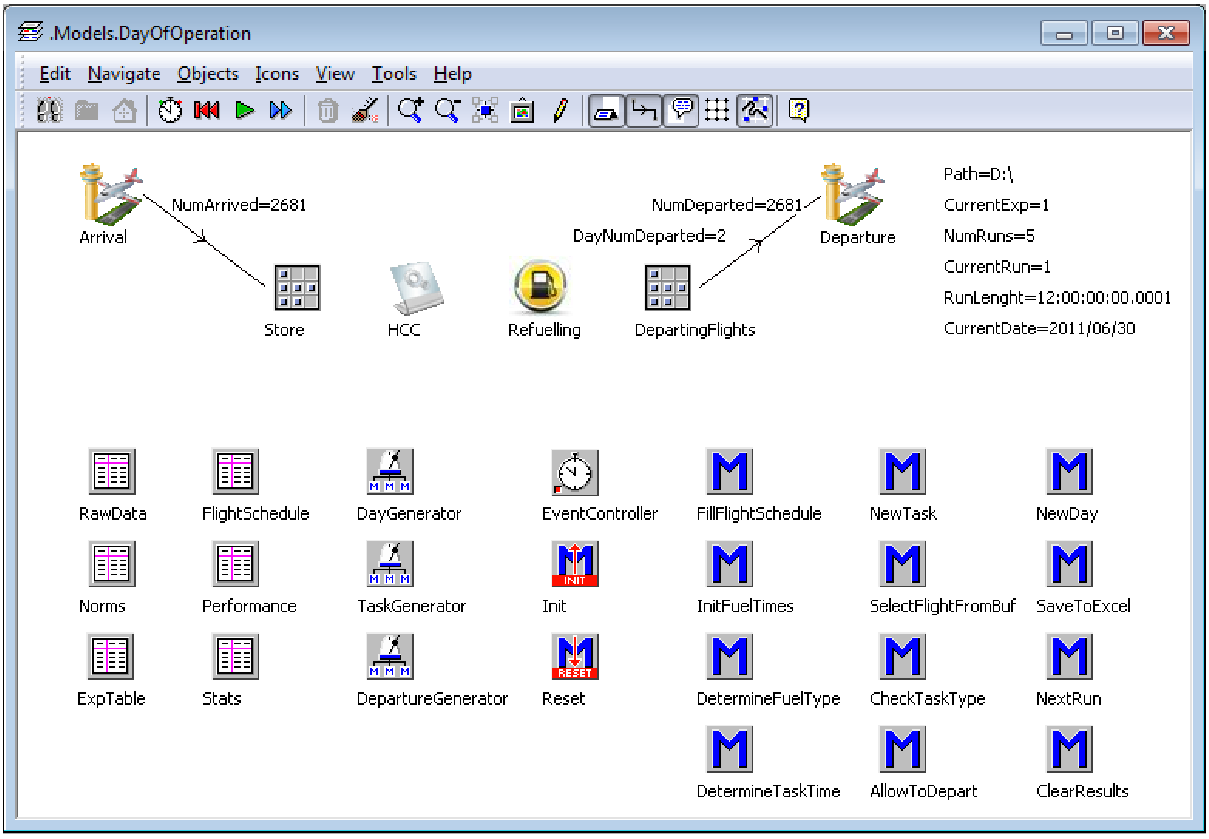

Figure 5.8: Main window Simulation Day of Operation.

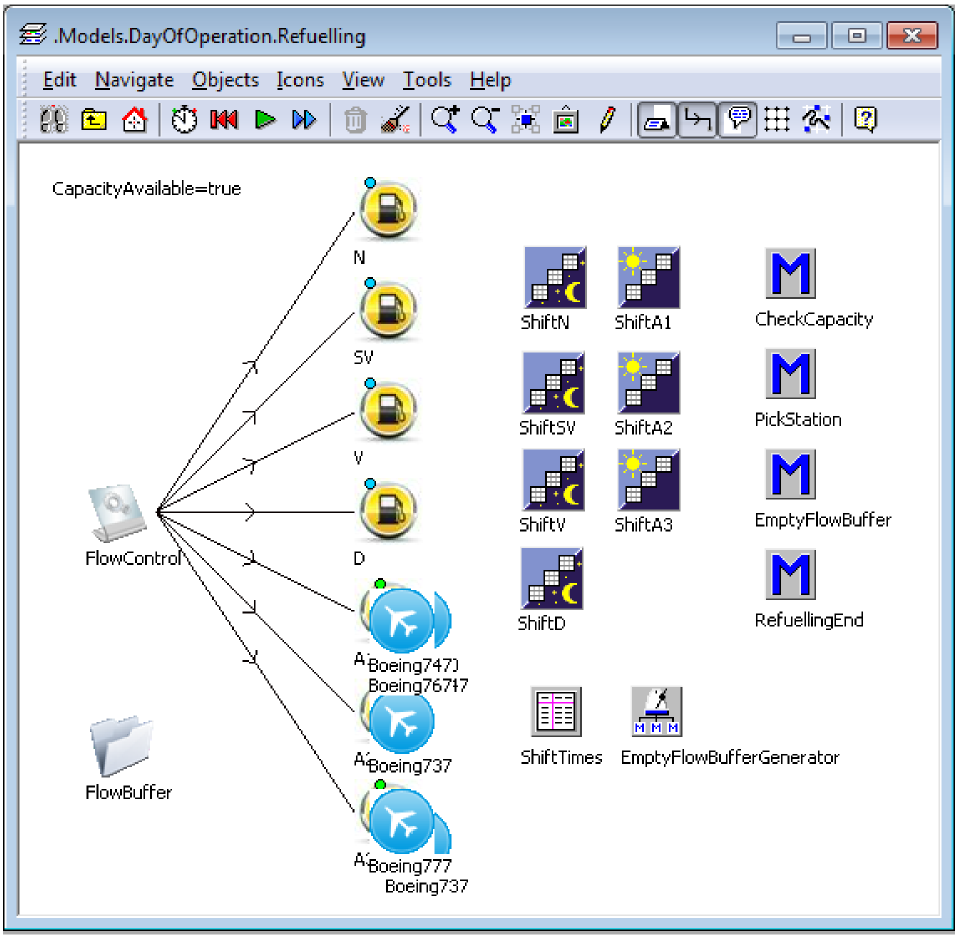

Figure 5.9: Refuelling window Simulation Day of Operation.

5.5.2 Model configuration

To conduct experiments with this model, we first need to set the correct configuration. We do this by defining the experiments we want to execute, the required number of replications and the required warm-up period.

Experiments

We execute several experiments to determine the impact of different sets of shifts. An overview of the composition of shifts is given in Table 5.2. We also execute an experiment using an unrestrained number of operators per shift. We use this for model validation, because this has to result in a performance close to 100%.

Number of replications

We execute multiple replications of all the experiments to obtain statistically significant results within the relative error margin (γ=0.05) and measure performance for every replication. The performance per experiment is the average over the different replications.

The performance statistics per replication are independent and identically distributed random variables. For this reason, we can, according to Law (2007), construct confidence intervals for the performance measure and use the sequential procedure to solve these for the number of replications. We determine the number of replications required using the following method, depending on the average and the standard deviation of the obtained performance:

$$\boldsymbol{n = min\left\{ n\text{: } \frac{t_{n-1,1-\alpha/2} * \sqrt{S_n^2/n}}{|\bar{X}_n|} \leq \frac{\gamma}{1+\gamma} \right\}}$$

This gives an outcome of about 2 replications per experiment. However, since Law (2007) dictates *“to be better safe than sorry”*, we use 5 replications per experiment.

We make use of common random numbers (CRN) in the used probability distributions to make sure that we use the same prediction of the future for every experiment. This keeps the experiments mutually comparable. However, within experiments, we use a different prediction of the future for every run, to create a confidence interval of the expected performance (instead of 5 identical results). We also want to keep the drawings from the different probability distributions independent of one another, so we never use the same random number for different processes.

Warm-up period

A warm-up period of a simulation is according to Law (2007) necessary when the initial conditions of a system are not representative for the steady state system behaviour (i.e. when we start with an empty system). A warm-up period can be defined by for example one set-up run. Welch method is an example of a good graphical method to determine the length of the warm-up period (Law, 2007). Other methods are evaluated by Mahajan and Ingalls (2004), who give an overview and evaluation of the performance of different methods.

In our case, the warm-up period is defined easier. Before our system reaches a steady state, it first needs to be filled with aircraft. The flight schedule supplied by Network incorporates this. This flight schedule consists of more flights than just the flights from one week, because aircraft stay in some cases overnight (or for a few days) at the airport. We need to incorporate the arrival of these aircraft too; otherwise we would miss some flights in our analysis. In our case, we call the period where the flight schedule is filled with flights the warm-up period of our system. Figure 5.10 illustrates the flight schedule, defined for 20 till 26 June 2011. The dates outside this range are only to fill and empty the system and are thereby not considered when analysing the system.

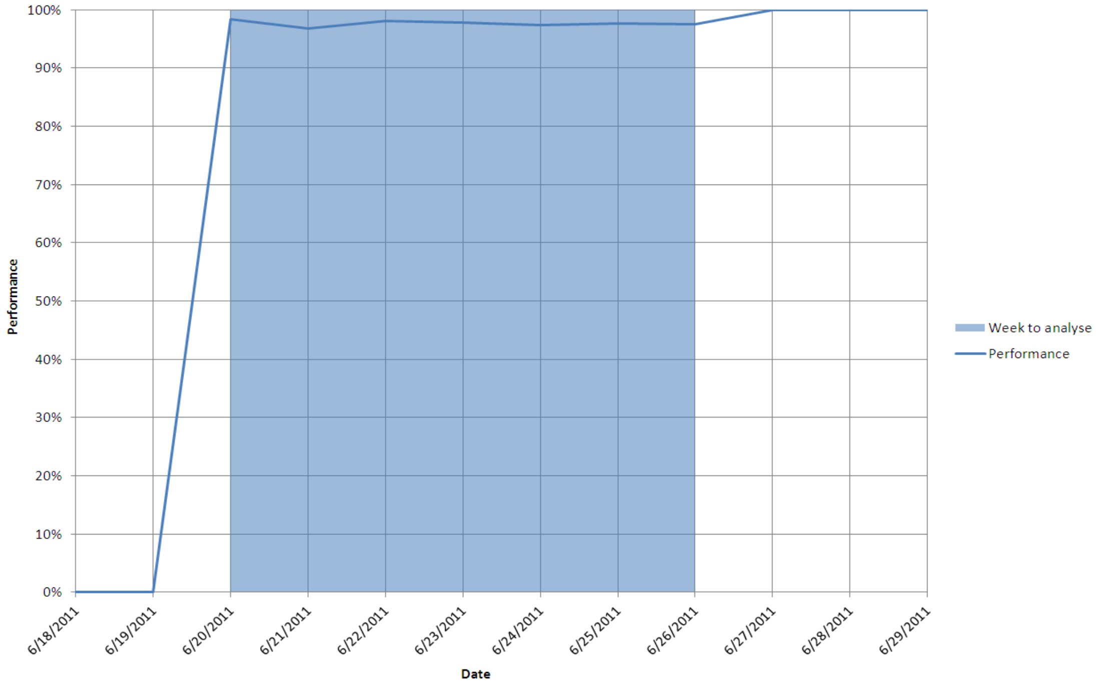

Figure 5.10: The simulated performance, for only the experiment at the 10th percentile.

5.5.3 Validation

For validation of our model, we check whether it complies with the conservation of material law. This law states: Created + Initial = Destroyed + Remaining (Standridge, 2004). By using this law, we make sure we do not loose flights during the process (which should be physically impossible).

To check the routing between the different stations of the model, we perform a visual check. In this visual check we check whether the routing of flights goes according to the predefined paper model (see Figure 5.7). We also check the sequence in which the flights are being refuelled, to make sure our selection algorithm works correctly. We execute these checks with different numbers of flights in the system and at different moments in time.

To check whether the process time is a good drawing of the corresponding theoretical probability distribution, we check the single arrival and process times for a number of different flight types. We also check the start and end times of processes and whether they comply with the earliest start and latest end times per aircraft type, as defined in the planning norms.

We check how the system works with extreme values. If we enter very high values for the capacity of the system, the system has to run smoothly and the performance should be high. When we enter low values for the system, the system should be congested and the performance should be low.

We keep a log file of all the non-performance and check whether we can explain the causes of this non-performance. If our model is valid, we should be able to identify the source of every non-performing task.

The last check we execute is checking the results with management. The results should not deviate much from what can be expected from reality. Management can judge this based on their experience with the refuelling process.

5.5.4 Results

We display the results of the simulation in Figure 5.11. This figure gives an overview of the simulated performance per 10th percentile of workload coverage, when we cover these percentiles of workload entirely with the shifts from Table 5.2.

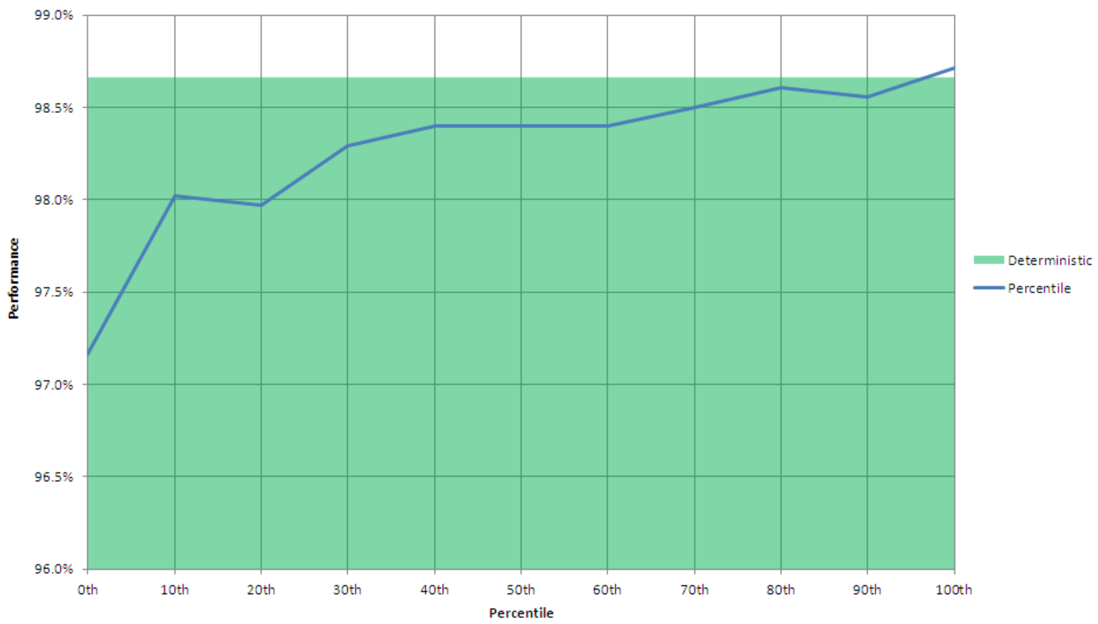

Figure 5.11: The simulated performance per fully covered percentile of workload, including the current deterministic situation.

Figure 5.11 indicates that the simulation gives a very high performance. We comment in this section on these results and in Section 5.5.5 we comment on our assumptions and modelling choices.

When we run the experiment with the maximum capacity from Table 5.2, we obtain the same performance as in the 100th percentile experiment. This means that the set of shifts fitted to the 100th percentile causes no delays due to capacity restrictions and thereby the 100th percentile is validated as the maximum workload which can be obtained. The corresponding performance is 98.72%. We do not reach a 100% performance due to some tight norms, set by KLM, for earliest start times. In these cases the process time is longer than the time between earliest start and latest end time.

Figure 5.11 does not give a smooth line. This is due to the optimisation of shifts in Section 5.4. We only minimise the number of shifts to reach minimal personnel costs, but we do not optimise the placing of these shifts; we accept the first feasible solution. The moment in time where we place extra shifts can also be optimised, but this is a topic for further research.

On-time performance and the impact of shifts

We conclude from Figure 5.11 that our model gives very good results for the on-time performance of the refuelling process. Even when we cover the lowest possible workload from 1000 Monte Carlo runs (the 0th percentile), we obtain a performance of 97.17%. This is partly caused by the shifts we use. On first instance, one should say that covering for example; the 80th percentile of workload would give a performance of 80%. However, due to the eight hour shifts, the achieved performance is much higher. The impact of these shifts is that a lot of operators are needed to cover the peaks, but due to the eight hour shifts, these operators are present for a longer time than just the peak times. This makes the schedule instantaneously more robust, because peaks can now move within a shift without influencing performance. When we choose to cover a lower percentile of workload with our shifts, we only cover less of the peak. The large amount of workload next to the peak is still being covered completely, due to the remaining time in the shift. This makes that covering different percentiles of workload does not have a linear relationship with the achieved performance.

We give an example: when we plot the set of shifts optimised to the 10th percentile of workload (red line), and the 90th percentile workload (blue area), in one figure, we obtain Figure 5.12. This figure gives a good overview of how a set of shifts optimised to a low percentile of workload performs in our model. The red line in the figure still covers 95.35% of the blue workload, even though it is a set of shifts optimised to a much lower percentile of workload. We conclude that a shift optimised to fit on a low percentile of workload, still performs well on high amounts of workload. This is caused by the combination of the shifts covering all the peaks in the workload and the requirement to use eight hour shifts.

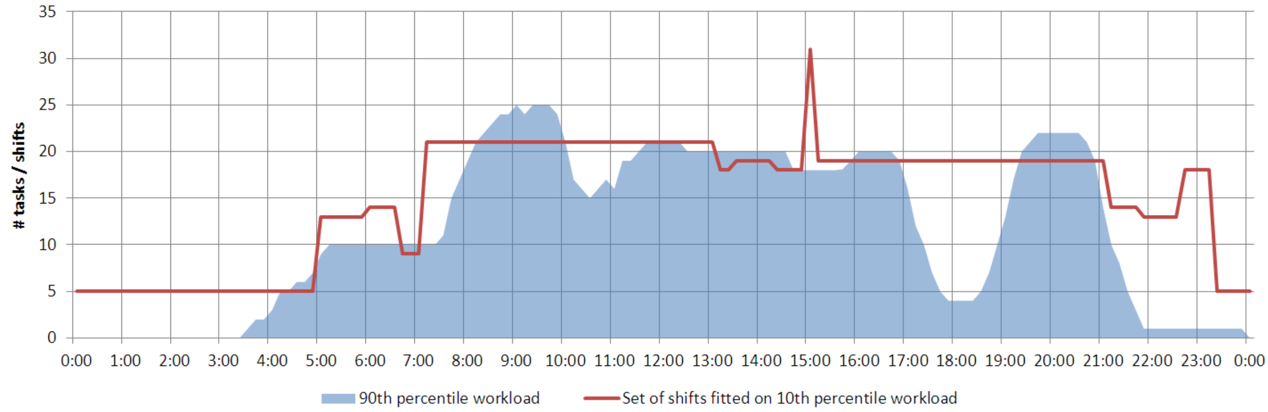

Figure 5.12: The set of shifts fitted on the 10th percentile of workload plotted at the 90th percentile workload line.

When we use a lower amount of shifts then defined by our lowest percentile, we cut through a bigger part of the workload. We check how this influences the performance by executing three extra experiments (named -1, -2 and -3), as defined in Table 5.3. For every experiment we reduce the number of shifts with five shifts, starting from the set of shifts optimised to our lowest percentile (the 0th).

ExpName	N	SV	V	D	A1	A2	A3
0th percent.	4	7	1	11	5	1	11
-1	3	6	1	10	4	1	10
-2	2	5	1	9	3	1	9
-3	1	4	1	8	2	1	8

Table 5.3: Additional experiments.

We simulate a day of operation with the shifts as defined in Table 5.3. Extending the results from Figure 5.11 with these results leads to the performance we display in Figure 5.13. This figure shows clearly that when we cut through a bigger part of the workload, by using sets with less shifts than the sets optimised to our lowest defined percentile (from Figure 5.5), the performance drops significantly.

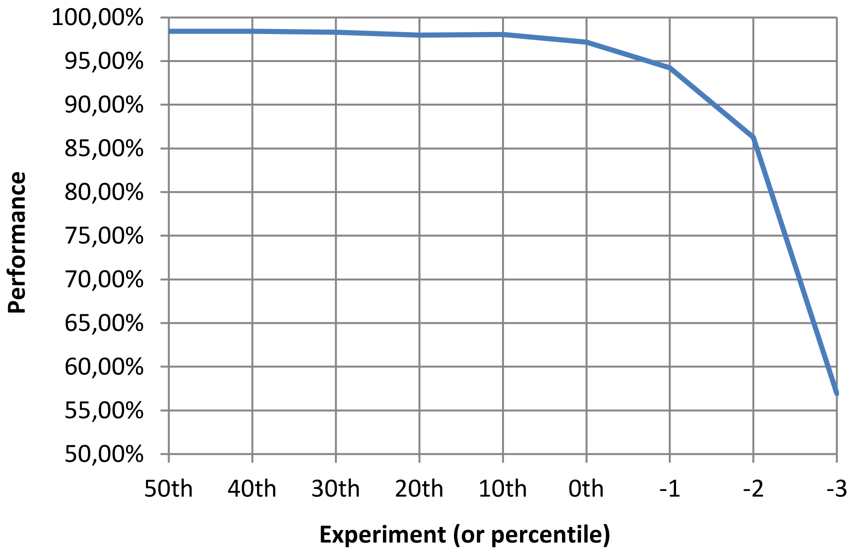

Figure 5.13: The performance resulting from the additional experiments.

Resulting from Figure 5.12 en 5.13, we conclude that choosing just one percentile of workload to cover over the day is not the best choice, since this is too much influenced by the peaks. When KLM wants to implement this research, they need someone who can make smart choices which part of the planned workload to cover and at what level to cut through the peaks. This person is able to make even better choices when KLM steps down from the restriction to use only 8 hour shifts, because then the expected workload line can be followed more closely. Our recommendation is to take a relatively low percentile in the peaks, since this decrease the amount of shifts needed (and thereby the costs) significantly and the performance does not decrease much. Figure 5.13 shows that the performance decreases fast when we let our shifts cut through a bigger part of the workload. For this reason, we recommend covering the off-peaks periods completely.

Fluctuation between runs

The standard deviation in the performance between different runs of one experiment varies between 0.001 (high percentiles) and 0.005 (low percentiles), which means that the performance does not vary much between different runs of the simulation and grows when we reduce the capacity. We therefor conclude that our model gives a robust solution over the different runs, independent of the covered percentile of workload.

Costs

In Section 5.4, we determined the cost of one shift to be €238 a day. Using the number of shifts needed to cover the workload per percentile (see Figure 5.6) and the simulated performance per percentile (see Figure 5.11), we can determine the costs corresponding to a certain performance. We give an overview of the yearly costs corresponding to the simulated performance in Figure 5.14.

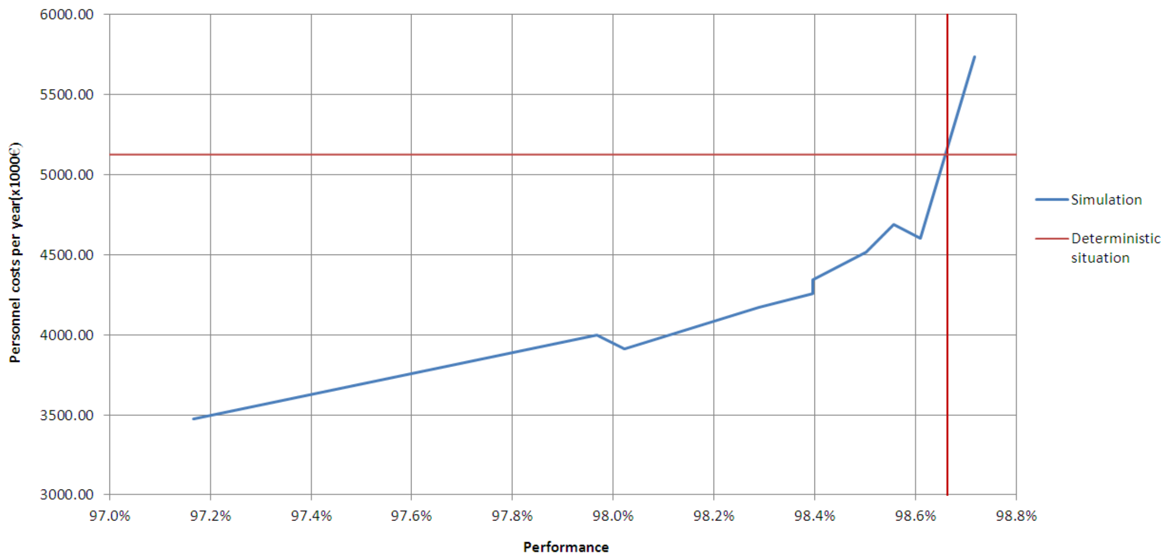

Figure 5.14: Performance versus yearly personnel costs, compared to the deterministic situation.

We conclude from Figure 5.14 that our model gives a difference in costs, between the current deterministic situation and the lowest obtained performance (corresponding to the 0^{th} percentile), of €1,650,530 a year. Another conclusion is that it is very expensive to obtain the performances corresponding to the highest percentiles, since the cost curve indicates an almost exponential growth of costs.

5.5.5 Impact of assumptions and modelling choices

Due to the model simplifications and assumptions made in Section 5.2, the performance and costs obtained from our model cannot be used to make decisions in the real world. This section gives an overview of how our model differs from reality. We treat only the assumptions which have a big impact and explain how they influence our model results.

Driving times

KLM underestimates the driving times of the fuel-operators between different aircraft in her planning. As described in assumption 3, KLM uses in the planning an estimation of an average driving time of 5.5 minutes per task. This average is created under the assumption that the next task of an operator is always an aircraft close by (which should be the case in an optimal situation). However, in reality an operator has to travel much further when the planning is shifted due to disturbances or due to delays in the arrival of aircraft. This is extra uncertainty, leading in reality to much more driving time and thereby a reduction in operator capacity and lower performances. We do not incorporate this extra traveling distance in our model due to the mentioned uncertainty in gate allocation, and thereby we obtain a higher performance than in reality.

Shifts

Following from assumption 4, we do not allow operators to work in overtime. An operator always finishes his task, and thereby it can happen that in reality, he works in overtime due to a disturbance. In our model, we do not allow overtime and thereby use the simplification to plan all tasks in such a way that everything is finished before the end of a shift. We know the exact duration of a task in advance and thereby we make a more optimal decision when to start a task. This makes our model give higher performances. However, planning the tasks is in this case also more complicated, which can decrease the simulated performance.

Extreme days
Following from assumption 7, the outcome of our models is only valid for normal days. As also mentioned at the results of our workload generation model in Section 5.3.4, we removed extreme values from the data when we defined the probability distributions, since KLM should not create their general planning for extremely disturbed days. When a day is highly disturbed, the data gets positively correlated, which leads to mutual dependent data. We built our models using independent samples from the probability distributions, so our simulation model does not generate the dependencies in delays needed to simulate an extremely disturbed day.

Breaks
Every operator has right to a 30-minutes lunch/dinner break and when time is available also a 15-minutes coffee break. We chose in assumption 11 not to model these breaks. We mentioned earlier that since a shift lasts eight hours, we have a lot of excess capacity next to the covered peaks. The impact of not modelling the breaks is that we overestimate the excess capacity in a shift next to the peaks. Our model uses this excess capacity as buffer capacity to finish tasks which could not be scheduled at their planned time due to for example a delayed arrival. In reality, this buffer capacity is less available, since the excess capacity is used to schedule breaks. We conclude from this that not modelling breaks gives an overestimation of both the excess capacity under a shift and the scheduling flexibility. This results in an overestimated simulated performance.

Changes in equipment
At Schiphol, not all gates are equipped with a hydrant fuelling system (as explained in Section 2.2.1). This forces KLM to use two types of refuelling equipment with different properties and deployment options. When we created our models, we used assumption 12 and thereby follow the KLM planning methods to not split the workload to different types of equipment. However, on the day of execution, KLM controllers do make this distinction and also sometimes change the allocation of operators (which takes about 40 minutes) to different types of equipment over the day. We assume all hydrant systems, and use thereby only one type of equipment in our simulation of a day of operation. This assumption makes that we use our capacity more efficient then in reality. For example: when an operator allocated to a dispenser, which services the larger aircraft, has in reality only 30 minutes left in his shift, he is not able to service another large aircraft anymore so he leaves early. In our simulation, this operator can still service a small aircraft, which is in reality serviced by a bowser (fuel truck), and thereby make more efficient use of his time. Concluding from this, we use the available capacity more efficient in our simulation and thereby obtain higher performances.

5.6 Conclusion

In this chapter, we created a link between the flight schedule extended with variances in the processes caused by uncertainty and unforeseen events, and the performance of KLM Aircraft Refuelling. We created this link by using four models to first generate workload lines (based on the predefined probability distributions and the flight schedule) and then optimise the starting times of the tasks to obtain an optimised workload profile. Third, we created an optimal coverage of the workload on different levels, using shifts to schedule personnel. In the last model, we simulated a day of operation using the obtained sets of shifts to determine the corresponding performance.

We conclude that when we make a workload planning which incorporates uncertainty and unforeseen events, the difference between the minimum and maximum predicted workload is relatively large. From this, we conclude that uncertainty and unforeseen events have a big influence on the planned capacity giving a wide range of values the workload can take on every moment in time. Covering this workload on just one percentile value over the entire day is not the best choice, since then the coverage is too much influenced by the peaks. We also conclude that the eight hour shifts used by KLM have a large impact on the robustness of this planning, since they cover more than just the peaks in the workload due to their duration. This allows peaks to move within a shift without consequences for performance. The shifts thereby neutralize a large part of the discrepancy between planning and reality, caused by the arrival punctuality of flights.

When we simulate a day of operation, using the optimal set of shifts based on the optimised planning, we obtain performances around 97-98%. These performances do not differ much over different executions of the simulation. The high performances and the low fluctuations between different runs indicate a robust solution. Here we see that the lowest covered percentile gives still a good performance, with €1,650,530 less cost a year, compared to the current deterministic situation. Due to the assumptions we made to be able to model the refueling process with the limited scheduling guidelines KLM can supply, our simulation model overestimates performance. We cannot extent the conclusions from the simulated day of operation directly to make decisions in the real world. However, we can use the conclusions from the other models, and the lessons learnt from the simulation model in the upcoming chapters.

6 Robust capacity planning at other services

In this chapter, we make an assessment, based on the process characteristics as defined in Chapter 2, for which processes, and to what extent, we can apply the insights of Chapter 5. The main characteristics we consider are:

- The type of process (arrival/departure *oriented* or *strictly* arrival/departure processes, see Section 2.2) and thereby the presence of flexible permissible time windows and the dependency on other processes.
- The duration of the process.
- The amount of variance in the process.
- The shape of the workload profile.

An overview of the different services and their characteristics is given in Table 2.1 and the Figures 2.7, 2.8 and 2.9.

In the following sections, we treat the different services individually to check whether we can extent the insights about incorporating uncertainty and unforeseen events to the workload planning of these services.

We also determine the buffer types we are allowed to use per type of process. As defined in Chapter 1, uncertainty and unforeseen events cause variability in the planning. We introduced the Law of Buffering (Hopp and Spearman, 2008), saying that variability in a production system will always be buffered by some combination of inventory (or cycle time), capacity and time (longer lead times). An example of buffering with cycle time is widening the norms, allowing more time to complete a task. Buffering with capacity is done by hiring extra operators or allocating extra equipment. Buffering with time moves the starting time of a task. Mostly tasks are then delayed to a less busy moment in time.

6.1 Water and Toilet service

The most important characteristic of the water and toilet service is that these tasks have a relatively short duration and large flexible permissible time windows. This gives a high degree of scheduling flexibility and thereby the workload is relatively flat without large peaks. This is indicated by Figure 6.1 and 6.2 which give an example of the workload profiles for these services. These workload profiles illustrate the current (deterministic) situation.

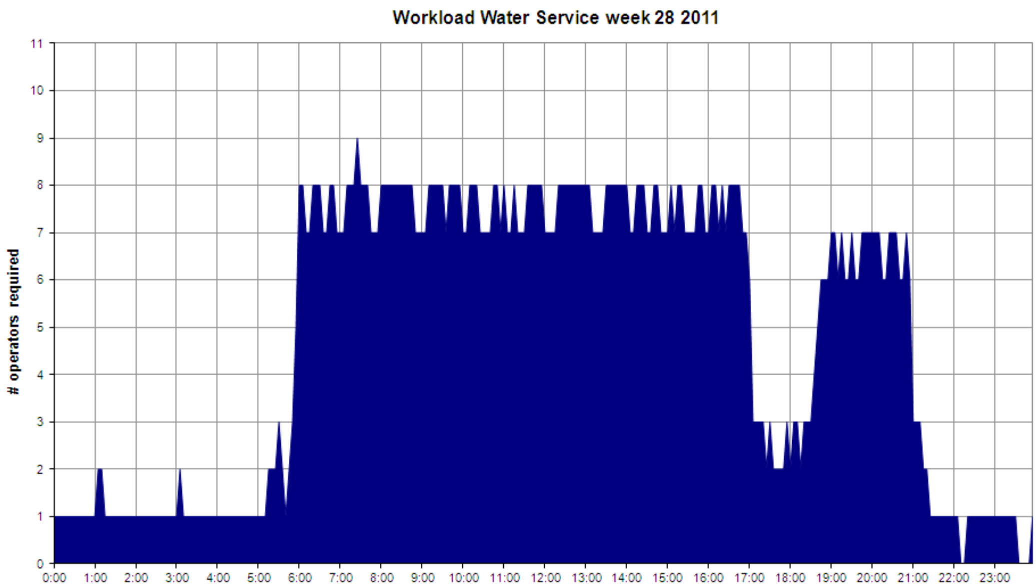

Figure 6.1: An example workload profile (deterministic) from the Water Service.

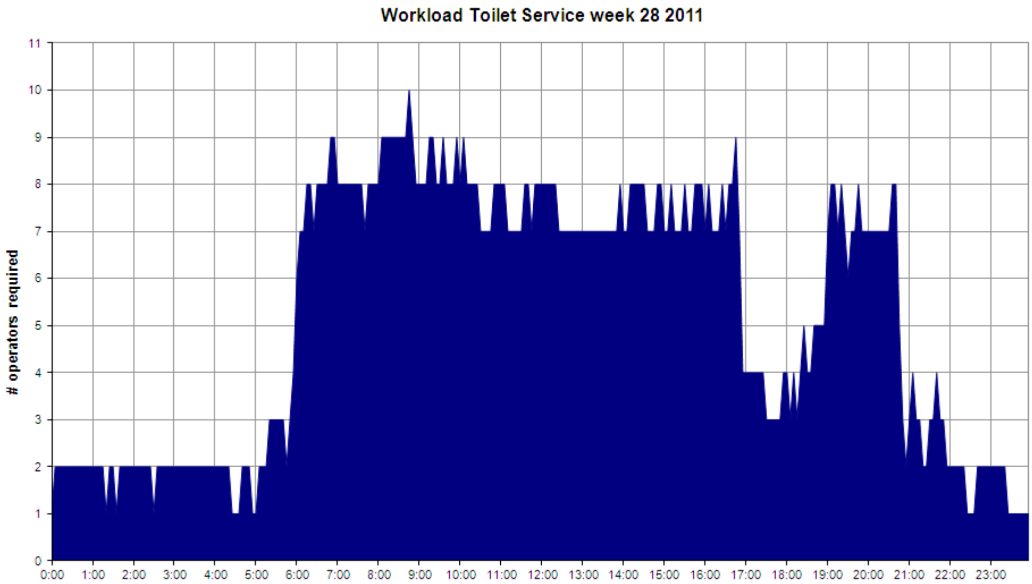

Figure 6.2: An example workload profile (deterministic) from the Toilet Service.

Figure 2.8 and 2.9 indicate that there exists significant variation in the duration of these tasks (especially at the water service). When we incorporate uncertainty and unforeseen events in these workload profiles, similar to as we did for refuelling in Chapter 5, we expect to obtain a relatively large difference between the lowest and highest values of workload. Due to the absence of many high peaks, the workload profiles are relatively flat. This gives the opportunity to work with a close fitting set of shifts covering the workload of these services. The breaks now just add up to the workload and do not affect our scheduling flexibility. A tight shift fit gives a very close covering of workload, and thereby good potential for a strong link between the cost of covering a certain workload and the expected performance. When using the method of Chapter 5, it is possible to create a robust capacity planning, where the model results can directly be translated to reality.

Since the water and toilet service are both arrival/departure *oriented* processes, these processes are allowed to buffer with all three option of buffering. The planning norms for the process can be widened (giving more time to complete the task), extra capacity can be scheduled and the process can be delayed in time due to the presence of flexible permissible time windows.

6.2 Airside Handling Support

Airside handling support (AHS) connects the aviobridge to the aircraft and facilitates the crew briefings. They also transport crew to the aircraft when an aircraft departs from a buffer. Connecting the aviobridge makes AHS a *strictly* arrival/departure process, because when they are late the passengers are unable to de-board, or the aircraft is not able to depart. Since this service is heavily reliant on arrival and departure, it does not have flexible permissible time windows. The tasks of AHS are relatively short (see the characteristics defined in Chapter 2) and thereby the peaks in the workload follow exactly the peaks in the flights schedule (see Chapter 4), which gives a very variable workload. We give an example of an AHS workload profile in Figure 6.3.

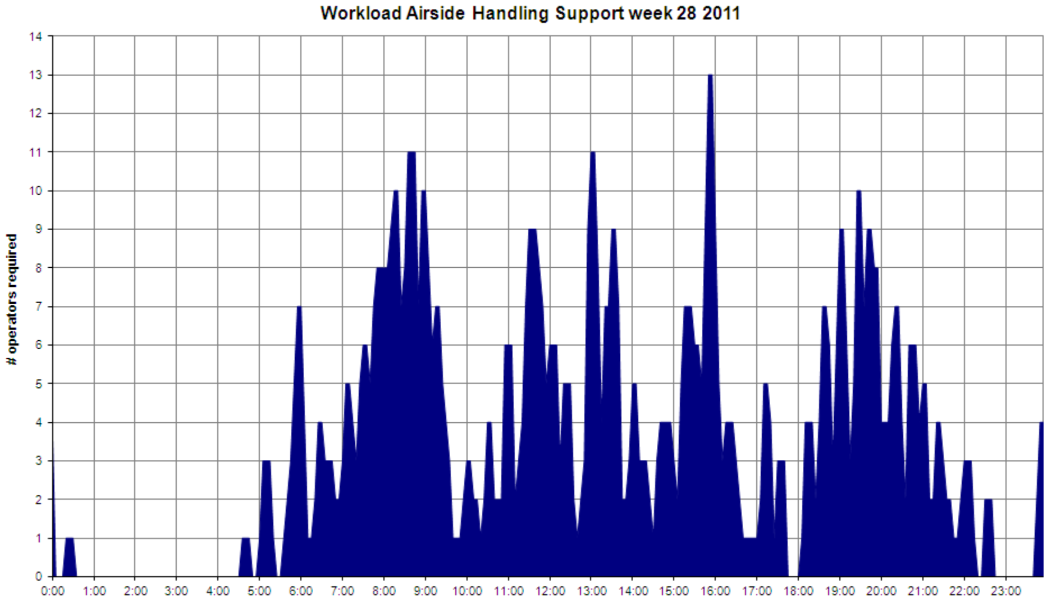

Figure 6.3: An example workload profile (deterministic) from Airside Handling Support.

According to Figure 2.8 and 2.9, the duration of the processes of AHS is relatively certain and the process does not take much time (see Figure 2.7). This means that using probability distributions to determine the process time of connecting the aviobridge is not very relevant. Disturbances in this process can be of relevance, so KLM needs to investigate whether there are many disturbances and whether they have long durations. When this is the case, they should be described according to a probability distribution. For creating a robust planning at AHS, we have to deal especially with the arrival punctuality. We indicated in Chapter 5 that uncertainty and unforeseen events in the arrival punctuality can very well be described using a probability distribution and that this has impact on the location of the peaks. Chapter 5 also tells us that the arrival punctuality at the refuelling process is for a large part covered by the shifts. KLM needs to investigate whether this is also valid at AHS.

Since AHS is a *strictly* arrival/departure process, it needs an operator to be present to attach an aviobridge at the moment an aircraft arrives. This means that buffering with time is not desirable, i.e. by moving the start time of a task (delaying the task). We have to buffer the variability in the process with cycle time by setting broader norms (more time available per task) or by scheduling extra capacity (Hopp & Spearman, 2008).

6.3 Towing and Pushback service

The required capacity for towing and pushback is planned simultaneous, similar to the refuelling service with hydrant and non-hydrant fuelling. A controller determines on the day of operation which operator executes which task.

As indicated in Chapter 2, KLM only tows an aircraft from a gate to a buffer when they are obliged to, due to the gate allocation determined by the Schiphol Airport Authority. KLM also tows aircraft to Schiphol East for maintenance, which happens according to the flight schedule or due to an aircraft breakdown. There exist a dependency between towing and the other ground services. When an operator wants to start towing, the other services have to stop and continue work on the new location after towing. Almost all ground services depend on the towing service which has to deliver the aircraft on the right time at the right location.

The pushback service is a strict departure process and thereby there are no flexible scheduling possibilities for these tasks. The start time of pushback depends, in case of a short turnaround critical in time, on the arrival punctuality of aircraft and always on the variability in other ground services executed before the aircraft is allowed to depart. All of these services can cause a delay in the departure of the aircraft. This makes the process vulnerable to uncertainty and unforeseen events.

We give an example workload profile of the pushback and towing service in Figure 6.4. This workload profile indicates that there are a lot of peaks in the workload of these services. From Chapter 5, we learned that covering a workload profile with this amount of peaks with eight hour shifts, gives a lot of excess capacity under the shifts.

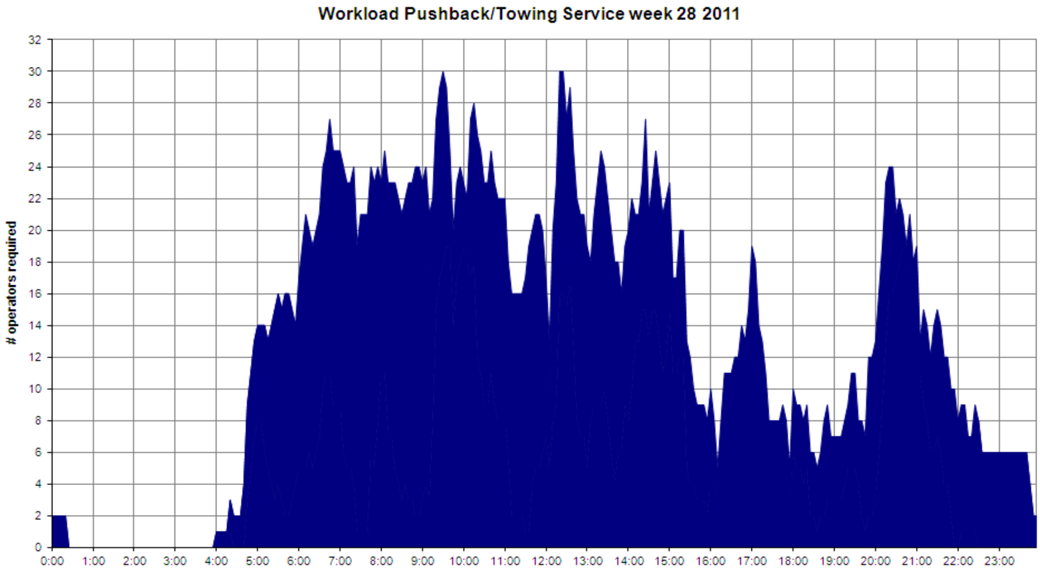

Figure 6.4: An example workload profile (deterministic) from the pushback and towing service.

Pushback is a *strictly* arrival/departure process and does thereby not have flexible permissible time windows. In Chapter 2, we did not describe the pushback and towing service in terms of average, standard deviation and coefficient of variation, because pushback occurs in the building block flight and towing is not done at a short turnaround. Those characteristics make this a special service. We give small summary statistics of the pushback and towing service is Table 6.1.

	Average	**StDev**	**CV**
Towing	25.704	12.790	0.498
Pushback	19.176	11.696	0.610

Table 6.1: Description of the duration of the pushback and towing service.

We conclude from the high coefficients of variation (CV) in Table 6.1 that the pushback and towing service is highly variable relative to the average duration. This gives good possibilities for making the workload planning more robust by incorporating uncertainty and unforeseen events in the planning. Just as at AHS, KLM needs to investigate whether it makes sense to describe disturbance by means of a probability distribution. Describing the arrival punctuality is of less relevance for pushback, since this is a strict departure process. We recommend KLM to investigate whether it makes sense to describe the departure punctuality stochastically. However, they have to keep in mind that there has to be an operator available at departure. For towing the arrival punctuality is more relevant, but also this impact needs further research.

Since pushback is a strict departure process and since a lot of other ground processes are dependent on towing, we do not want to buffer against variation by delaying the tasks (buffering with time). This means that we need to buffer either with cycle time (more time available per task) or capacity (Hopp and Spearman, 2008).

6.4 Other services

There are also some services which we do treat extensively. Extending robust capacity planning to these services will not lead to improvements or is not applicable. We describe these services in this section.

We do not consider services which are outsourced, because their capacity planning is not under the influence of KLM. These services are the cleaning and catering (and board supply) services.

We do not consider the security check since their capacity planning is easy. They just schedule two cabin inspectors in the morning and two in the afternoon and they check aircraft at random.

We do not consider the flex-tasks since the scheduling of these tasks is easy. There are two or three operators scheduled in the morning and in the afternoon. These people execute the flex task and if they are finished they execute other temporary tasks. Due to weather circumstances (mostly high or low temperatures), sometimes more employees are necessary. This is covered by temporary workers.

The last service we do not consider is de-icing. This service is only relevant in the winter period, due to the weather circumstances. The amount of capacity needed depends on the severity of the weather and is determined one day in advance. The operators are volunteers (for example office personnel) and temporary workers.

6.5 Conclusion

Incorporating uncertainty and unforeseen events in the workload planning is possible for the Refuelling service, Water service, Toilet service, Aircraft Handling Support and the Pushback and Towing service. The relevance per service depends on the duration of the task and the amount of variation in the different aspects making a task perform on time (arrival punctuality, task process time, possible disturbance, etc.). For *strictly* arrival/departure processes the variation in the arrival/departure punctuality of aircraft is the most relevant and for arrival/departure *oriented* processes the variation in the task process time is the most relevant.

For the other services of Aircraft Services, it is not relevant to incorporate uncertainty and unforeseen events in the workload planning, since either these services are not under direct control of KLM or their scheduling is very easy or they dependent so much on specific weather circumstances that they need an entirely different kind of planning.

When there are a lot of peaks in a workload profile, covering this workload with eight hour shifts leads to a lot of excess capacity. For this reason it is hard to link the cost of covering workload directly to performance. For services with very flat workload profiles, a relatively strong link can be created between the cost of covering a certain value of workload and the required performance. These are mostly arrival/departure oriented processes with short durations, because they have large flexible permissible time windows to be able to optimise in such a way that peaks in the workload are almost non-existent.

KLM has to introduce buffers to deal with the effects of uncertainty and unforeseen events in the workload planning. The Law of Buffering (see Section 1.2) defines which type of buffers can be used. To determine the type of buffers we can use, we need to make a distinction between *strictly* arrival/departure processes and arrival/departure *oriented* processes as defined in Section 2.2. *Strictly* arrival/departure processes are only allowed to buffer with cycle time (more time available per task) or extra capacity. Buffering with time (delaying the process) is not desirable, since this immediately delays the de-boarding of passengers (and thereby their possibility to make their connection to other flights) or the departure of the aircraft (while the passengers are already on-board). Arrival/departure *oriented* processes are allowed to buffer with all three types of buffers. This type of process does mostly have flexible permissible time windows, allowing them to schedule the starting moment of their tasks in a certain window. However, when they use too much time, the aircraft can still be delayed. Introducing buffers in the planning always increases the robustness; however this goes at the expense of optimality.

7 Implementation

In this chapter, we give recommendations for how KLM Aircraft Services should deal with uncertainty and unforeseen events in her planning to make this planning more robust. Dealing with uncertainty can be done by a combination of handling or managing this uncertainty, reduce it and buffer against the part which cannot be reduced. We give an overview of this process in Figure 7.1, which is also the main structure of this chapter.

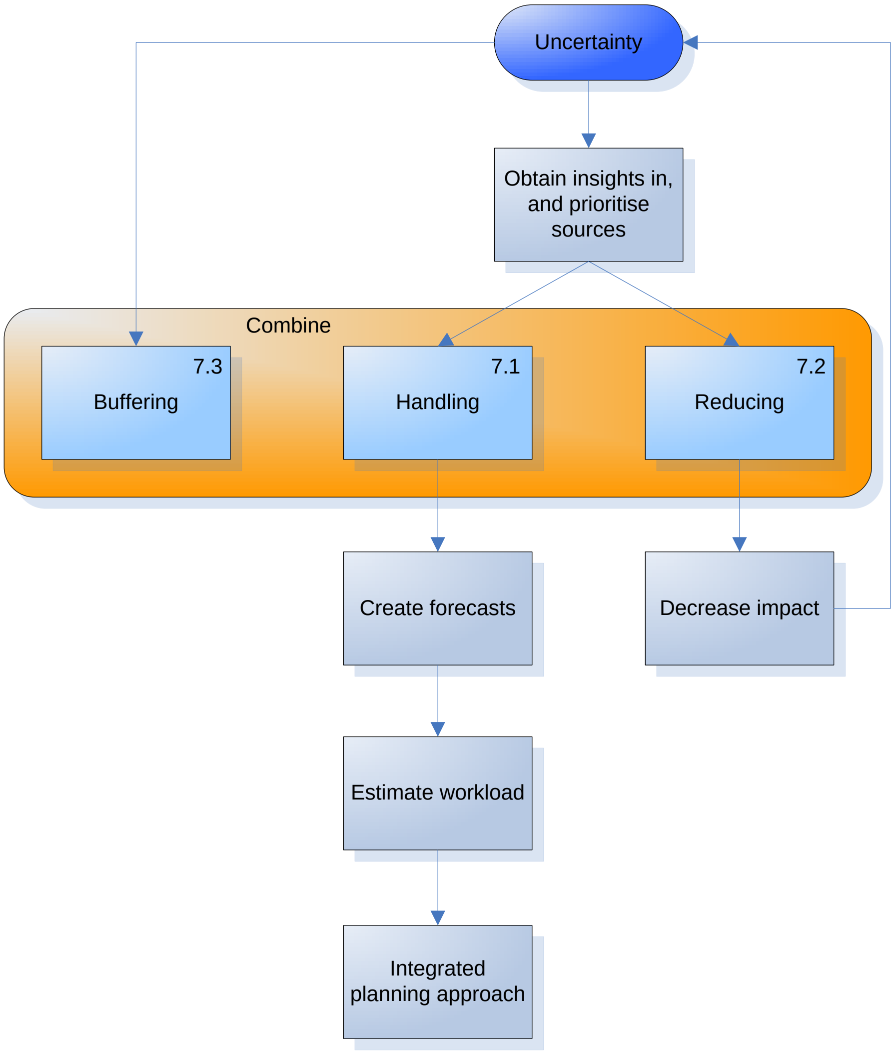

Figure 7.1: Steps to take to deal with uncertainty.

When KLM Aircraft Services wants to create a more robust capacity planning, they first have to find a good way to handle the current amount of variability caused by uncertainty and unforeseen events. This uncertainty needs to be reduced when they want to improve the robustness and optimality of the planning. This is an iterative process, since less uncertainty may lead to other handling methods. The part of uncertainty which cannot be reduced has to be covered by some way of buffering, following the Law of Buffering as described in Section 1.2.

Implementing a robust planning is a process which can take a long time and thereby it is essential to follow the steps of Figure 7.1. The sections where we handle the different solutions are marked in this figure.

7.1 Covering of workload

In our analyses from the previous chapters, we worked with historical data from which we determined theoretical probability distributions to generate new values, used to create a workload profile. We recommend substituting the planning norms for task durations by the probability distributions as defined in previous chapters and to use probabilistic arrival times instead of the static arrival times from the flight schedule. When KLM implements this method, they cannot simply use the same parameters for the probability distributions for every period, since this would give an outdated representation of reality. What KLM needs, is a good forecast for the new period, based on the historical data of multiple periods, and determine the parameters based on this forecast. Over the years, this forecast can deviate a lot due to differences in the flight schedule. This causes the need to analyse the parameters every time a change is made in either the flight schedule or in the working methods of the different departments. There are multiple methods for creating forecasts, based on the characteristics of the data set (like certain trends or differences between seasons). Hopp and Spearman (2008) give a short overview of the most commonly used forecasting methods (moving average, exponential smoothing, etc.). For more sophisticated models, we refer the reader to the book of Box and Jenkins (1970).

We learned from Chapter 5 that covering just one percentile of workload is not the best decision, however defining percentiles does give good insight on how the workload is distributed. To create a robust planning, we need to decide how to deal with the range of optimised workload lines resulting from incorporating uncertainty and unforeseen events in the workload (see Figure 5.5). This range gives an overview of all possible values the workload can take at a certain moment in time and provides thereby good insights for KLM to make well founded decisions about which part of the workload to cover.

In the current situation, the resource planners plan everything step-by-step (see Chapter 3), but better would be to integrate all steps: from choosing the workload to cover (which is now done by the tactical planning department (ST) of Ground Services), till determining a set of shifts. When we create a set of shifts covering the chosen workload values, we often obtain excess capacity due to the covering of peaks with eight hours shifts. We concluded that breaks can very well be used to fill this excess capacity. For this reason, we recommend an integrated planning approach. Since the steps to take require a lot of feeling for the processes, KLM is not able to develop clear scheduling rules, making an optimisation using an intelligent computer program not feasible. We recommend using an expert who executes multiple cycles of design, testing and adapting as illustrated in Figure 7.2. This expert needs a lot of knowledge of all the processes to take well-founded decisions about which value from the workload range to cover, for every moment in time. He has to make decisions to get a robust and relatively optimal covering of the workload. We illustrate this in Figure 7.2 and give a detailed explanation afterwards.

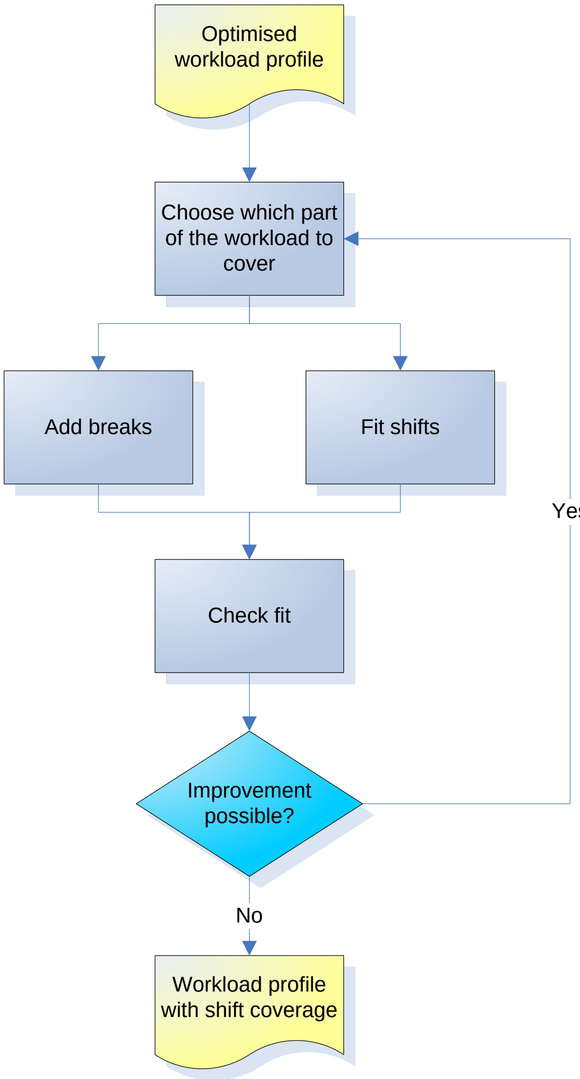

Figure 7.2: The recommended integrated planning approach.

Using an integrated planning approach, the expert uses a planning loop consisting of choosing the amount of optimised workload to cover (from the pre-defined range), covering this with shifts, planning breaks and checking feasibility. He can now schedule breaks in such a way, that they fall exactly into the excess capacity of the shifts. When there is not enough excess capacity available to schedule for example the breaks, an extra shift needs to be scheduled. The expert can then use sensitivity analysis to determine the best moment for scheduling an extra shift and start the planning loop again. Our recommendation is to make the schedule more robust by scheduling these shifts at the moment where the range between the minimum and maximum amount of workload is the biggest (on the horizontal and vertical axis of the workload profile). These shifts functions as a buffer against uncertainty. When we use this approach, we are able to use the capacity created by the shifts optimal, but still have a robust planning since we make well-founded decisions of which amount of workload to cover from the predefined range which displays the effect of all kinds of uncertainty and unforeseen events.

7.2 Decreasing uncertainty and workload balancing

We concluded in previous chapters that the most disturbing aspect in workload profiles is the high amount of peaks. When these peaks can be reduced, we are able to create a capacity planning which follows the optimised workload closely and thereby performs well against lower costs. For this reason, the primary goal of KLM for optimising workload should be to minimise the peaks. They should stop immediately using other optimisation goals (see Chapter 3), since, as indicated in Chapter 5, peaks in the workload translate directly to high personnel costs.

A recent development at KLM is the introduction of extra short turnarounds for the European flights. This concept decreases the flexible permissible time windows of AS and creates thereby more peaks due to the lack of scheduling flexibility. Extra short turnarounds work counterproductive in our goal of peak reduction. For KLM as a whole however, the revenue generated from executing more flights with less aircraft has the potential to exceed these extra costs. We recommend a thorough consideration of the positive and negative aspects of this development.

To decrease the variance in the task duration, we recommend splitting the task duration to set-up time and processing time. This because the set-up time is relatively certain and can thereby be treated as deterministic. The processing time is still relatively uncertain at the moment of planning and has thereby a lot more variation, i.e. at the refuelling service the processing times depends on the amount of fuel to deliver, which is uncertain till the last moment since this depends on the load of the aircraft. We also recommend using the driving time matrix (indicating the driving times between different locations on the airport), instead of the relatively unreliable determined average driving time used in the current situation, to incorporate more reliable driving times in the planning.

For the refuelling service, we recommend to break down the planning norms to a lower level than per aircraft type. When KLM creates different norms per aircraft type and then per destination, they can make a better estimation of the required fuel amount and thereby of the task duration. KLM should also investigate whether breaking down the norms to a lower level is relevant for other services.

To reduce the variability in especially the arrival punctuality, but also in the task durations, we recommend KLM to develop different scenarios for the workload planning, based on different circumstances with similar characteristics. These scenarios differ for example on the weather circumstances or on the month of the year. Variables to investigate are the different sources of uncertainty as defined in Chapter 4. Between the scenarios, we can change the parameters of the used probability distributions. This makes the workload predictions more precise and thereby it improves the fit of the workload profile to reality.

An interesting point to investigate in further research is how much KLM should move along with already known factors. For example: at this moment, we can give a reasonable prediction of the weather for the coming days, or early next week. Do we need a dynamic planning which adapts to every change in circumstances, or is one fixed planning for a longer period better?

7.3 Buffering

As defined in Chapter 1, uncertainty and unforeseen events cause variability in the planning. According to the Law of Buffering, variability in a production system will always be buffered by some combination of inventory (or cycle time), capacity and time (Hopp and Spearman, 2008). We explain the application of buffering at AS in this section.

Time buffers

We concluded in Chapter 6 that arrival/departure *oriented* processes can buffer with time (longer lead times). For *strictly* arrival/departure processes (Pushback, AHS) it is not desirable to buffer with time.

Buffering with time as used in Chapter 6, means that we are able to move the start time of a task. For example: delay the moment of execution of a task within a certain window. Another way of buffering with time is designing for non-performance. This means that we plan for a performance of for example 90% and accept that we do not achieve the last 10%. We are now able to delay a task till a less busy moment of the day, outside the flexible permissible time window. This almost certainly delays the aircraft and leads to costs of delay. However, these costs of delay can be less than the costs which had to be made to perform on-time and thereby it can be profitable. Design for non-performance can lead to cost saving, but it also leads to dissatisfied customers who might choose another airline the next time they fly. For this reason, KLM has to be very careful when choosing this option.

When designing for non-performance, it is important to not always delay the same flights. This is unacceptable in terms of customer satisfaction and due to agreements with Schiphol. The Schiphol Airport Authority and Air Traffic Control give every flight a time slot for departure. When KLM cannot make this time slot, they are scheduled for another, but when they are consequently unable to make it to this time slot, they might lose it to another airline.

Capacity and inventory buffers

We concluded in Chapter 6 that all services of AS can buffer with cycle time or extra capacity to cover the variability in their tasks. Buffering with cycle time (more time available per task) means that we extend our norms to have more time to finish a task. Buffering with extra capacity means that we hire more operators or deploy more equipment. Using these buffers we can buffer for the arrival punctuality and for task disturbances or delays. KLM can determine the required size of these buffers by using the knowledge of the earlier in this chapter introduced expert. This expert can use the predefined range of workload values, to make a good estimation of the required buffer size.

7.4 Conclusion

Uncertainty needs to be dealt with, by first being able to manage the current amount of uncertainty, then start initiatives to reduce it and finally buffer against the amount of uncertainty which cannot be reduced. We recommend the approach illustrated in Figure 7.1 to reach solid results and obtain a permanent improvement.

Managing uncertainty in the workload planning of KLM needs to be done by first being able to make a good forecast for the parameters of the defined probability distributions. These probability distributions have to be used, instead of the current planning norms for tasks, in the process of creating workload profiles. This allows KLM to create an estimated range of values the workload can possibly take. Then KLM needs an integrated planning approach (as described in Figure 7.2) to make well founded decisions about which parts of the workload to cover, which shifts to use and how to schedule other activities (like breaks). This is a process with multiple cycles of design, testing and adapting, which have to be executed by an expert with a lot of knowledge of all the processes.

KLM has to start reducing uncertainty with executing further research to define the impact of all sources of uncertainty, implicating the ground processes. When this is clear, steps need to be taken to decrease the most significant sources. We recommend the following steps:

- Decrease the peaks in the workload profiles as much as possible and use this as the only goal for optimising the planned workload.
- Investigate the concept of extra short turnarounds, since this concept works counterproductive against the goal of peak reduction.
- Split the task duration to set-up time and processing time. Since the set-up time is relatively certain, this can be treated deterministic. The processing time can now be broken down further than then just the aircraft type. We decrease the variation when we first break down to the type of aircraft and afterwards to the destination of the flight, since we are then able to make a better estimation of the required amount of fuel and thereby the task duration.
- Use the driving time matrix, instead of the relatively unreliable determined average driving times used in the current situation.
- Develop different scenarios for the workload planning, based on different circumstances with similar characteristics (for example weather circumstances or different periods of the year) to make workload predictions more precise and better fitted to reality.

Buffering against uncertainty is according to the Law of Buffering automatically being done by some combination of inventory (or cycle time), capacity and time. KLM Aircraft Services needs to make well founded choices which buffer to use in what situation. Arrival/departure *oriented* processes (Refuelling, Water, Toilet) can buffer with all three options. For *strictly* arrival/departure processes (Pushback, AHS) buffering with time is not desirable.

KLM can also design for non-performance. This leads to costs of delay, but these can be less than the costs of performing exactly on-time. KLM has to be very careful when choosing this option, because next to cost savings it also may lead to dissatisfied customers.

8 Conclusions & Recommendations

In this chapter, we draw our conclusions by answering the research questions we defined in Chapter 1. By doing so, we refer back to the goal of this research: "To provide insight into the effect of uncertainty and unforeseen events on the dynamics of personnel capacity planning, related to the performance of KLM Aircraft Services."

We recall the research questions from Chapter 1:

1. How can the different processes of KLM Aircraft Services be characterised?
2. What is the current way of planning at KLM Aircraft Services?
3. What are the main elements affected by uncertainty and unforeseen events in the planning of KLM and KLM Aircraft Services?
4. What performance can be expected when a robust capacity planning is used at the aircraft refuelling process?
5. How and in what extent can robust capacity planning be extended to the other ground processes of KLM Aircraft Services?
6. What is the best way to implement this robust way of capacity planning at KLM Aircraft Services?

After the conclusions, we give recommendations for further research which KLM Aircraft Services should execute to improve the planning and performance of the ground processes, next to the recommendations for implementation from Chapter 7.

8.1 Conclusions

To achieve the goal of this project, we only analysed one of the processes of KLM Aircraft Services in depth and then extended the insights of this analysis to the other processes. We concluded that the refuelling process is the best process to analyse in depth, based on that this process is arrival/departure *oriented*, relatively flexible to schedule and has a relatively long duration and high variation in this duration. To choose between the different processes, we first characterised them as an answer to the first research question. The most important characteristics in the scope of this research were:

- The type of process (*strict* arrival/departure or arrival/departure *oriented*)
- The flexibility of scheduling the process (dependency on other processes and whether the process has flexible permissible time windows).
- The duration of the process
- The variation in the duration of the process

We concluded from research question 2 that the capacity planning currently used, is a completely deterministic coverage of the total workload. The planning is primary based on the workload directly derived from the flight schedule and static norms, where Aircraft Services does not account for uncertainty and unforeseen events. This means that Aircraft Services uses just one single option for the workload and creates her entire capacity planning based on this single option. The management of Aircraft Services does thereby not get any insights into the variability affecting the processes, which we proved to be present and significant.

Answering research question 3, we concluded that there a lot of different sources of uncertainty and unforeseen events affect the planning of KLM, from which some are truly uncertain and some are relatively predictable elements. KLM should react to the relatively predictable elements and buffer against the truly uncertain elements. The most important elements used for planning at KLM Aircraft Services are the arrival punctuality of aircraft, the duration of the ground processes and possible disturbances. The arrival punctuality of aircraft varies over the time of the day and over different periods of the year and can be reasonably predicted (on a short horizon) when analysed in depth. The process time of the refuelling service differs a lot over the different aircraft types and per type of refuelling task. This is for a large part due to the amount of fuel an aircraft has to take in. The refuelling process is disturbed in 15.84% of all tasks, due to unforeseen events. We concluded that theoretical probability distributions can be used to describe the behaviour of these uncertain elements.

To answer research question 4, we needed to create a workload planning which incorporates uncertainty and unforeseen events, using the defined probability distributions. This leaded to a large difference between the minimum and maximum predicted workload. From this, we concluded that uncertainty and unforeseen events have a big influence on the planned capacity, giving a wide range of values the workload can take over the day. We concluded that covering the workload on just one percentile value over the day is not the best choice, since then the coverage is too much influenced by the peaks. We also concluded that the eight hour shifts used for scheduling have a large impact on the robustness of the planning, since they cover, due to their duration, more than just the peaks in the workload. This allows peaks to move within a shift without large consequences for performance. The shifts thereby neutralize a large part of the discrepancy between planning and reality, caused by the arrival punctuality of flights.

When we simulated a day of operation, using the set of shifts based on the workload planning which incorporates uncertainty and unforeseen events, we obtained steady performances around 97-98%. Here we indicated that the lowest covered percentile still gives a good performance, with €1,650,530 less cost a year, compared to the old situation. However, due to the assumptions we made, to be able to model the refueling process with the limited scheduling guidelines KLM could supply, our simulation model overestimates performance. This means that we could not draw conclusion from our simulation model which we can directly extend to reality, but we could extent the obtained insights to the other processes of KLM Aircraft Services to answer research question 5.

We concluded that incorporating uncertainty and unforeseen events in the workload planning is possible for the Refuelling service, Water service, Toilet service, Aircraft Handling Support and the Pushback and Towing service. For *strict* arrival/departure processes the variation in the arrival/departure punctuality of aircraft is the most relevant and for arrival/departure *oriented* processes the variation in the task process time is the most relevant. If a workload profile has a lot of peaks, covering the workload with eight hour shifts leads to a lot of excess capacity. For this reason it is hard to link the cost of covering workload directly to performance. A relatively strong link can be created between the cost of covering a certain value of workload and the required performance, for services with very flat workload profiles. These are mostly arrival/departure *oriented* processes with short durations, because they have large flexible permissible time windows to be able to optimise in such a way that peaks in the workload are almost non-existent.

As the best way to implement a more robust capacity planning at KLM Aircraft services (research question 6), we recommend implementing the approach we described in Chapter 7 and illustrated in Figure 7.1. This chapter contains all steps KLM needs to take to benefit the most from the results of this research.

To deal with uncertainty, we recommend implementing an approach where KLM first manages the current amount of uncertainty and simultaneously start initiatives to reduce this uncertainty. Finally KLM should buffer against the amount of uncertainty which cannot be reduced. To manage uncertainty, KLM needs to make a good forecast for the parameters of the defined probability distributions. These probability distributions have to be used in the process of creating workload profiles (instead of the current static planning norms) to create an estimate of the range of values the workload can possibly take. Then KLM needs an integrated planning approach (as described in Figure 7.2) to make well founded decisions about which parts of the workload to cover, which shifts to use and how to schedule other activities (like breaks). This is a process with multiple cycles of design, testing and adapting, which has to be executed by an expert with a lot of knowledge of all the processes.

Reducing uncertainty starts with executing further research to define the impact of all sources of uncertainty and afterwards taking steps to decrease the most significant sources. We recommend reducing uncertainty in an iterative way, alongside managing and buffering against uncertainty.

KLM has to introduce buffers to deal with the effects of uncertainty and unforeseen events on the workload planning. Buffering against uncertainty is according to the Law of Buffering automatically being done by some combination of inventory (or cycle time), capacity and time. We concluded that arrival/departure *oriented* processes (Refuelling, Water, Toilet) can buffer with all three options, but for *strictly* arrival/departure processes (Pushback, AHS) buffering with time is not desirable. KLM must take in mind that introducing buffers in the planning always increases the robustness; however this goes at the expense of optimality.

8.2 Recommendations for further research

During this research, we found several points that KLM should investigate further. These are points we could not investigate as thorough as we wanted due to time limitations and/or points we want to mention, but are outside the scope of our research.

- KLM Aircraft Services has to execute an extensive research to the sources of uncertainty and unforeseen events in her processes. We set the first steps in Chapter 4 of this report, but when KLM wants to fully understand the magnitude of every single source of uncertainty, more research in this direction is necessary.
- We described in Chapter 4 that operator behaviour has a large influence on the quality of the data KLM registers. Since operators do not clock the starting moment of their task at the same moment in time, more variability in the data is being created. This also means that, for example at the refuelling service, we cannot make a distinction between set-up times and actual refuelling time. We recommend KLM to take immediate action to align the clocking moment of their operators.
- We recommend KLM to investigate the possibilities of workload optimisation in depth. KLM should optimise always in such a way that the peaks in the workload are minimised, but this can be combined with optimising the placing of shifts in such a way that the personnel costs are minimised and buffer capacity is placed at the right place. In Chapter 5, we concluded that KLM does not optimise much in the placing of shifts. We recommend investigating the optimal moment in time to place extra shifts. Furthermore, we concluded that eight hour shifts are very disturbing for the optimality of the planning, so we recommend KLM to investigate the possibilities of using shorter shifts.

- We indicated in Section 7.2 that it could be interesting to investigate how KLM should deal with already known factors. We can, for example, give a reasonable prediction of the weather for the coming days, or early next week. Do we need a dynamic planning which adapts to every change in circumstances, or is one fixed planning for a longer period better? Both options can be good, but they require different planning methods.
- The performance per task is in the current situation only defined as finished on-time or too late. However, this causes KLM to not make any distinction between a little too late and much too late. To give more insight into the obtained performance, we recommend KLM to also monitor the amount of minutes a task finishes late.
- KLM should strive to use integral norms for all airliners. During our modelling, we noticed a lot of (sometimes small) differences between norms for the same aircraft type, owned by different airlines. These norms are also very poorly documented and can differ per documentation system. We recommend KLM to generalize all norms per aircraft type and generalise all documenting systems to one up-to-date system. The GOMS system has the potential to function well, but it has to be updated and expanded in such a way that it incorporates all aircraft for all airlines.
- We recommend to KLM to use theoretical probability distributions to describe the duration of their ground processes in the way we illustrated in this report. However, when KLM chooses not to switch and to keep using static norms for these durations, than they should conduct a thorough investigation to the way the current norms are designed. In the current situation, this is still being done by using a stopwatch to take measurements on the platforms. This method is not fail proof and often leads to norms based on an insufficient number of measurements. KLM keeps an extensive database of all historical data, so when they do not want to use theoretical probability distributions, we recommend creating the norms based on historical data retrieved from this database.

8.3 Additional insights

Our opinion is that KLM has to implement this research as we described in Chapter 7 of this report. The most important reason for this is that it gives KLM much more insights in their planning process. By using our approach, an explicit insight is given in how uncertainty and unforeseen events affect the daily workload. Using our approach, a business manager receives a workload profile which illustrates the entire range of values the workload can take on every moment of the day, created using the best forecast for that moment. The expert mentioned in Chapter 7 will give an advice on how to cover this workload to the business manager. This gives the business manager the opportunity to take well-founded decisions, because he has a clear overview of all uncertainty influencing his process.

There exists a high probability that implementing our recommended way of planning leads to resistance within KLM, since we expect that it will be seen as a far more complex way of planning. The current capacity planning works reasonably well, but is not transparent due to all the adjustments made to it from the moment of creation till the day of execution. This leads to hidden assumptions and extra buffering in different stages of the planning process. Our recommended way of planning leads to a transparent and more robust planning, where all uncertainties are clearly made visible to the responsible managers.

References

Anthony, R.N. (1965). *Planning and Control Systems: A Framework for Analysis.* Harvard University Graduate School of Business, Boston, Massachusetts.

Bedeian, A.G. and Mossholder, K.W. (2000). *On the use of the coefficient of variation as a measure of diversity.* Organisation Research Methods, vol 3, no 3, pp 285-297.

Box, G.E.P. and Jenkins, G.M. (1970). *Time Series Analysis, Forecasting and Control.* San Francisco: Holden-Day.

Cacchiani, V. and Toth, P. (2012). *Nominal and robust train timetabling problems.* European Journal of Operational Research. 219. pp 727-737.

Dantzig, G., Fulkerson, R. Johnson, S. (1954). *Solution of a Large-Scale Traveling-Salesman.* Journal of the Operations Research Society of America, Vol. 2, No. 4, pp.393-410.

Dantzig, G., Orden, A., Wolfe, P. (1955). *The generalised Simplex method for minimizing a linear form under linear inequality restraints.* RAND Corporation, Princeton University.

Dekkers, K. (2010). *A Design for KLM Aircraft Services' planning – roostering – indeling chain*. Stan Ackermans Institute, Eindhoven.

Eglese, R.W. (1990). *Simulated Annealing: A tool for Operational Research.* Department of Operational Research and Operations Management, The Management School, Lancaster University, Lancaster, UK.

Galbraith, J. (1973). *Designing Complex Organizations.* Addison-Wesley, Reading, Massachusetts.

Gao, C., Johnson, E. and Smith, B. (2009). *Integrated Airline Fleet and Crew Robust Planning.* Georgia Institute of Technology.

Glover, F., and Greenberg, H.J. (1989). *New approaches for heuristic search: A bilateral linkage with artificial intelligence*. European Journal of Operational Research 39, 119-130.

Halton, J.H. (1970). *A retrospective and prospective survey of the Monte Carlo method.* SIAM Review, vol 12, no 1, January 1970.

Hans, E.W. (2001). *Resource Loading by Branch-and-Price Techniques.* PhD thesis. Twente University, Enschede.

Hopp, W.J., Spearman M.L. (2008). *Factory Physics.* Third edition, McGraw Hill, New York.

Lane, D. (2010). *Percentiles*. Retrieved from http://cnx.org/content/m10805/2.11/ at 16-08-2012.

Law, A.M. (2007). *Simulation modelling and analysis.* Fourth edition. McGraw-Hill, New York.

Mahajan, P.S. and Ingalls, R.G. (2004). *Evaluation of methods used to detect warm-up period in steady state simulation.* Proceedings of the 2004 Winter Simulation Conference. Oklahoma State University, Oklahoma.

Mulvey, J.M. and Vanderbei, R.J. (1994). *Robust Optimization of Large-Scale Systems.* Operations Research, Vol. 43, No.2, March-April 1995.

Schiebaan, H. (2002). *Integral Fleet Service Control: Design Project at Schiphol Fleet Services*. Stan Ackermans Institute, Eindhoven.

Standridge, C.R. (2004). *How factory physics helps simulation.* Proceedings of the 2004 Winter Simulation Conference. School of Engineering, Grand Valley State University, Grand Rapids, Michigan.

Visser, A (2010). *Normentraject Tankdienst: De weg naar een norm.* KLM Royal Dutch Airlines, Amsterdam Airport Schiphol.

Walfish, S. (2006). *A review of statistical outlier methods.* Statistical Outsourcing Services, Rockville, Maryland.

Wilf, H.S. (1994). *Algorithms and Complexity.* First edition, University of Pennsylvania, Philadelphia.

Appendix A. Glossary

AS: KLM Aircraft Services, the department were this research takes place.

AWF: In Dutch the aanwezigheidsfactor: the presence factor.

BAX: Baggage and/or cargo.

Bowser: a bowser is a generic name for a tanker of various kinds. At AS this term is used to refer to a fuel tanker.

Buffer: An aircraft parking place, not on a gate.

Flexible permissible time windows: Some slack in the norms of the processes which can be used by control to advance or delay a process (in Dutch: schuifnorm).

GOMS: Ground Operations Manual Schiphol.

GS: KLM Ground Services.

Hub: A place that forms the effective centre of an activity, region, or network.

Hub and spoke system: A system where passengers are delivered to a hub airport with small continental aircraft and then transfer to the bigger intercontinental aircraft to reach their destinations over sea.

KCS: KLM Catering Services, which is a daughter company of KLM.

OPC: Operational Plan Check, see Chapter 3 for an explanation.

PAX: A passenger or multiple passengers.

ST: The tactical planning department of KLM Ground Services.

Tug: An aircraft towing vehicle.

Workload: the number of employees needed to perform the amount of work which needs to be done at a certain moment in time.

Appendix B. Organisational structure

Air France KLM is a holding containing Air France S.A. and KLM N.V. were KLM is owned for the biggest part (93.41%) by the holding. This because Air France KLM is owner of the two investment foundations (SAK I and SAK II) and owns next to this all priority shares and a part of the common shares. All preference shares are owned by SAK I and II and by the State of the Netherlands, which has got 5.92% of the economic and voting rights. A small part is owned by other parties (0.67%).

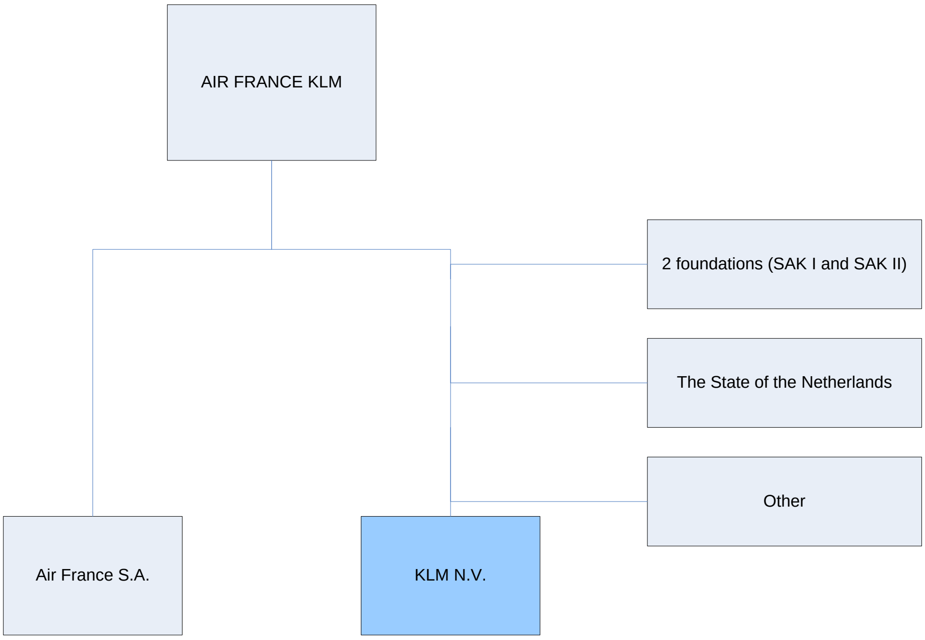

Figure B.1: Organisational structure AIR FRANCE KLM

Within KLM, there are seven divisions, supported by some supporting departments:

- Cargo
- Commercial
- Marketing, Revenue Management and Network
- Operations
- Flight Operations
- Inflight
- Engineering & Maintenance

The division Operations is the most important for this project and is split in six departments:

- Hub Operations Schiphol (KLM Ground Services)
- The Operational Control Centre (OCC)
- Mainport Strategy
- Alliances
- Security Services
- Outstation Management

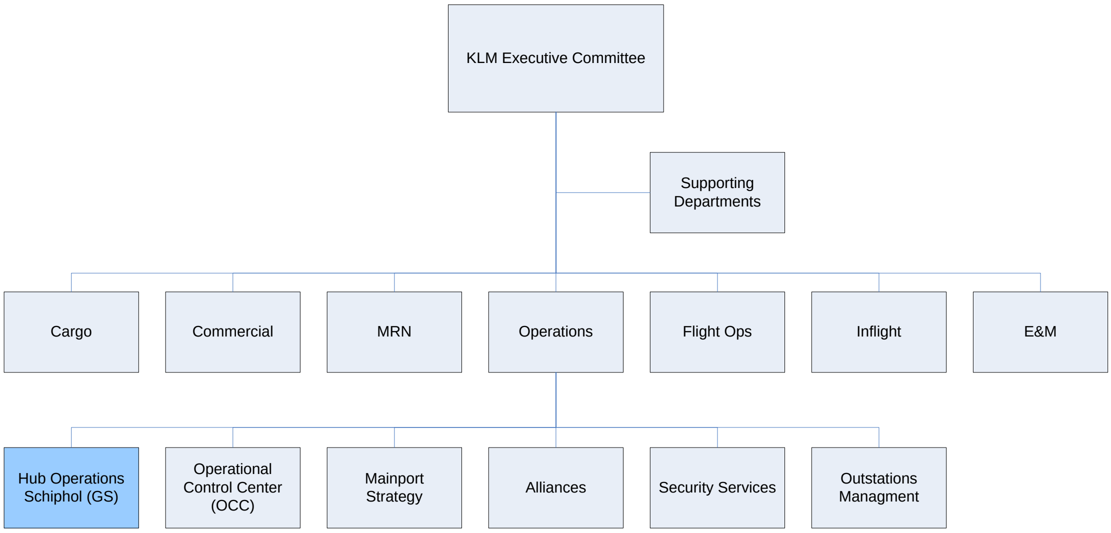

Figure B.2: Organisational structure KLM

The division Hub Operations takes care of all operations on KLM's hub: Schiphol Amsterdam Airport. Here it is called KLM Ground Services (GS). The services within Ground Services are:

- Passenger Services (PS)
- Baggage Turnaround Services (BTS)
- Aircraft Services (AS)
- Customer Ground Handling
- Hub Control Centre (HCC)
- Tactical Planning

Here Passenger Services is responsible for all processes around the boarding and deboarding of passengers (PAX). Baggage turnaround services takes care of all process around the loading and unloading of baggage and cargo (BAX). Aircraft Services is the place where this project is executed and is responsible for preparing the aircraft for the next flight and moving it around the airport. Customer Ground Handling is responsible for PAX and BAX flows of third party airliners. The Hub Control Centre (HCC) coordinates all ground processes of KLM and her partners at Schiphol Amsterdam Airport. The Tactical Planning department (ST) takes care of the planning of all processes executed by KLM Ground Services.

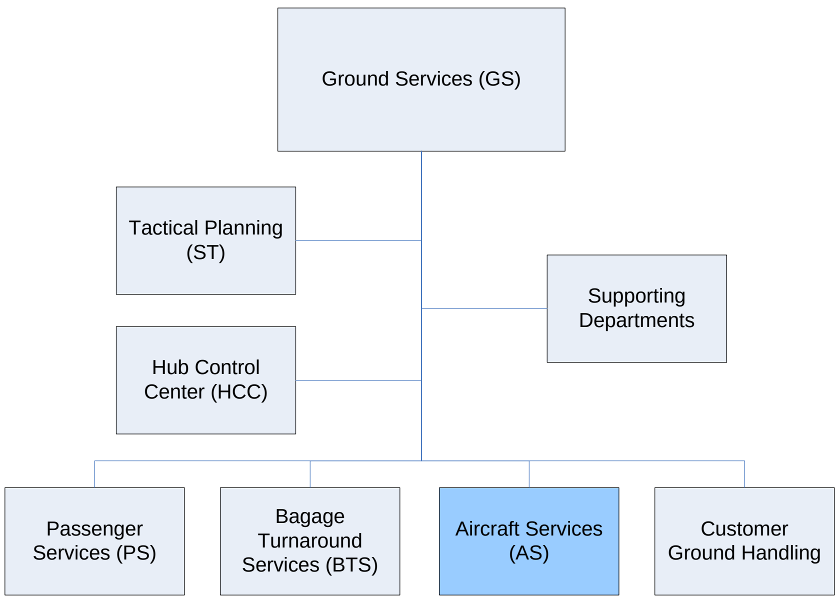

Figure B.3: Organisational structure Ground Services

Aircraft Services is organised in five departments:

- Movement & Aircraft Support
- Readiness
- Operational Support
- Ground Support Equipment (GSE)
- Jetcenter

Movement & Aircraft Support is responsible for the Airside Handling Support, De-icing and Pushback & Towing services. Readiness (in Dutch: Gereedstelling) is responsible for the Refuelling, Cabin Quality, Catering & Board Supply and Water & Toilet services. Operational Support is responsible for analysing the incoming data and supplying improvement measures for planning and management support. Control is responsible for controlling all the processes of AS and the steering of operators. Ground Support Equipment is responsible for servicing and managing the available equipment. The Jetcenter is a completely detached unit situated on Schiphol Amsterdam Airport and Rotterdam The Hague Airport, providing ground services to jets, but this is out of scope of this research.

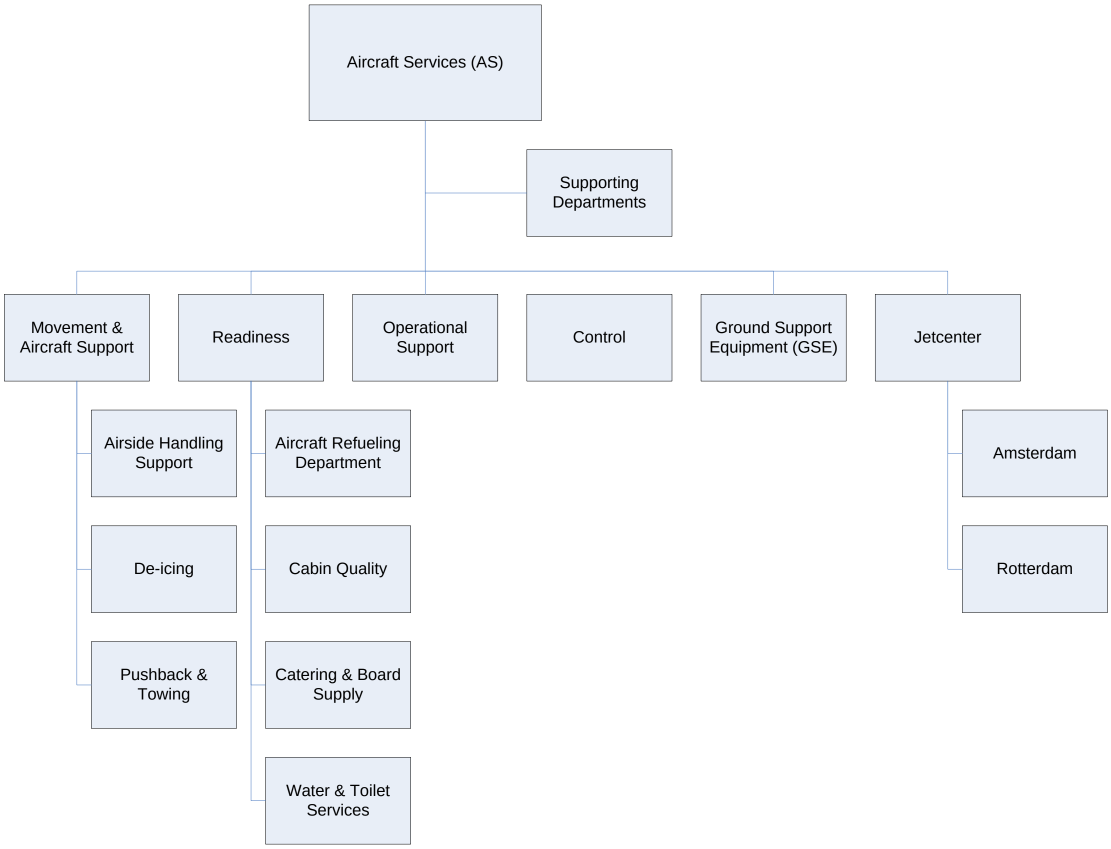

Figure B.4: Organisational structure Aircraft Services

Appendix C: Refuelling norms

Pre-fuel

carrier	AC-type	ground time [hh:mm]	duration [mm]	start [mm]	finish [mm]
KL	744	≥ 4:30	75	A+0	A+75
KL	744 Asia	≥ 4:30	75	A+0	A+75
KL	74E	≥ 4:30	75	A+0	A+75
KL	74E Asia	≥ 4:30	75	A+0	A+75
KL	777	≥ 4:30	65	A+0	A+75
KL	M11	≥ 4:30	65	A+0	A+75
KQ	763	≥ 4:30	40	A+10	V-105
KQ	772	≥ 4:30	40	A+10	V-125
SQ	772	≥ 3:00	65	A+0	A+75

Final Fuel

carrier	AC-type	ground time [hh:mm]	duration [mm]	start [mm]	finish [mm]
KL	744	≥ 4:30	45	V-90	V-10
KL	744 Asia	≥ 4:30	55	V-90	V-10
KL	74E	≥ 4:30	45	V-90	V-10
KL	74E Asia	≥ 4:30	55	V-90	V-10
KL	777	≥ 4:30	45	V-90	V-10
KL	M11	≥ 4:30	45	V-90	V-10
KQ	763	> 4:30	30	V-95	V-60
KQ	772	> 4:30	30	V-95	V-60
SQ	772	≥ 3:00	45	V-85	V-10

Table C.1: The refuelling norms for the summer period 2011, when executing a pre-fuel.

Final Fuel

carrier	AC-type	groundtime [hh:mm]	duration [mm]	start [mm]	finish [mm]
KL	744	< 4:30	90	V-110	V-10
KL	744 Asia	< 4:30	100	V-120	V-10
KL	74E	< 4:30	90	V-110	V-10
KL	74E Asia	< 4:30	100	V-120	V-10
KL	777	< 4:30	80	V-110	V-10
KL	M11	< 4:30	80	V-110	V-10
KL	332	-	60	V-110	V-10
KL	737/738/739	≥ 2:00	35	V-120	V-10
KL	737/738/739	< 2:00	35	A+0	V-10
KL	73W	-	45	V-70	V-10
KL	733/734	≥ 2:00	30	V-120	V-10
KL	733/734	< 2:00	30	A+0	V-10
KL	100	≥ 1:30	30	V-90	V-10
KL	100	< 1:30	30	A+0	V-10
KL	E90	≥ 1:30	30	V-90	V-10
KL	E90	< 1:30	30	A+0	V-10
KL	F70	≥ 1:30	25	V-90	V-10
KL	F70	< 1:30	25	A+0	V-10
A9	733/734	-	30	A+0	V-10
A9	735	-	35	A+0	V-10

AAN	320	-	35	A+0	V-10
CZ	332		60	V-110	V-10
CZ	772		80	V-110	V-10
CZ	332		60	V-110	V-10
DL	744	-	90	V-110	V-10
DL	330	< 2:15	65	A+5	V-10
DL	330	≥ 2:15 en < 4:30	65	V-130	V-10
DL	330	≥ 4:30	65	V-85	V-10
DL	757	< 2:15	45	A+5	V-10
DL	757	≥ 2:15 en < 4:30	45	V-130	V-10
DL	757	≥ 4:30	45	V-85	V-10
DL	767	< 2:15	60	A+5	V-10
DL	767	≥ 2:15 en < 4:30	60	V-130	V-10
DL	767	≥ 4:30	60	V-85	V-10
HV	73H	≥ 1:30	35	V-90	V-10
HV	73H	< 1:30	35	A+0	V-10
HV	73H	≥ 1:30 tussen 07:01 - 22:59	35	V-90	V-5
HV	73H	< 1:30 tussen 07:01 - 22:59	35	A+5	V-5
HV	73W	≥ 1:30	30	V-90	V-10
HV	73W	< 1:30	30	A+0	V-10
HV	73W	≥ 1:30 tussen 07:01 - 22:59	30	V-90	V-5
HV	73W	< 1:30 tussen 07:01 - 22:59	30	A+5	V-5
KQ	763	< 4:30	45	V-115	D-60
KQ	772	< 4:30	50	V-125	D-60
MP	763	< 2:00	60	A+5	V-10
MP	763	≥ 2:00	60	V-80	V-10
PY	340	-	90	V-110	V-10
SQ	772	< 3:00	80	V-110	V-10
SU	319	-	30	A+0	V-10
SU	320	-	35	A+0	V-10
SU	321	-	40	A+0	V-10
TWI	734		30	A+0	V-10
VQ	320		35	A+0	V-10

Table C.2: The refuelling norms for the summer period 2011, when only a final fuel is executed.

Appendix D: Flight schedule

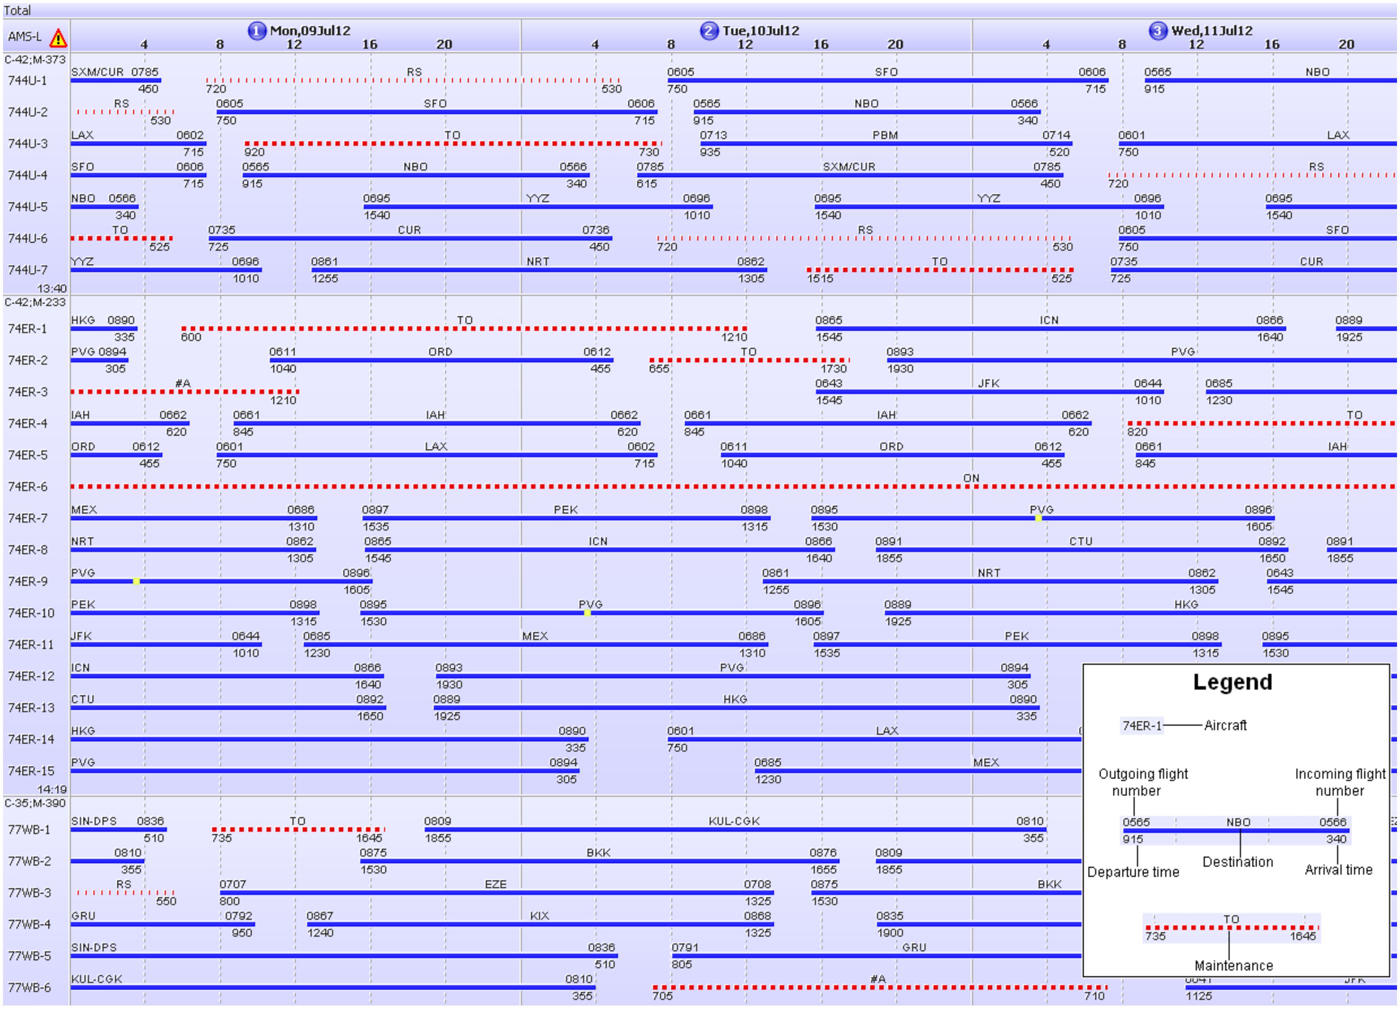

Figure D.1: An example of a flight schedule. This flight schedule has been used for the OPC of summer 2012.

Appendix E: Workload profile

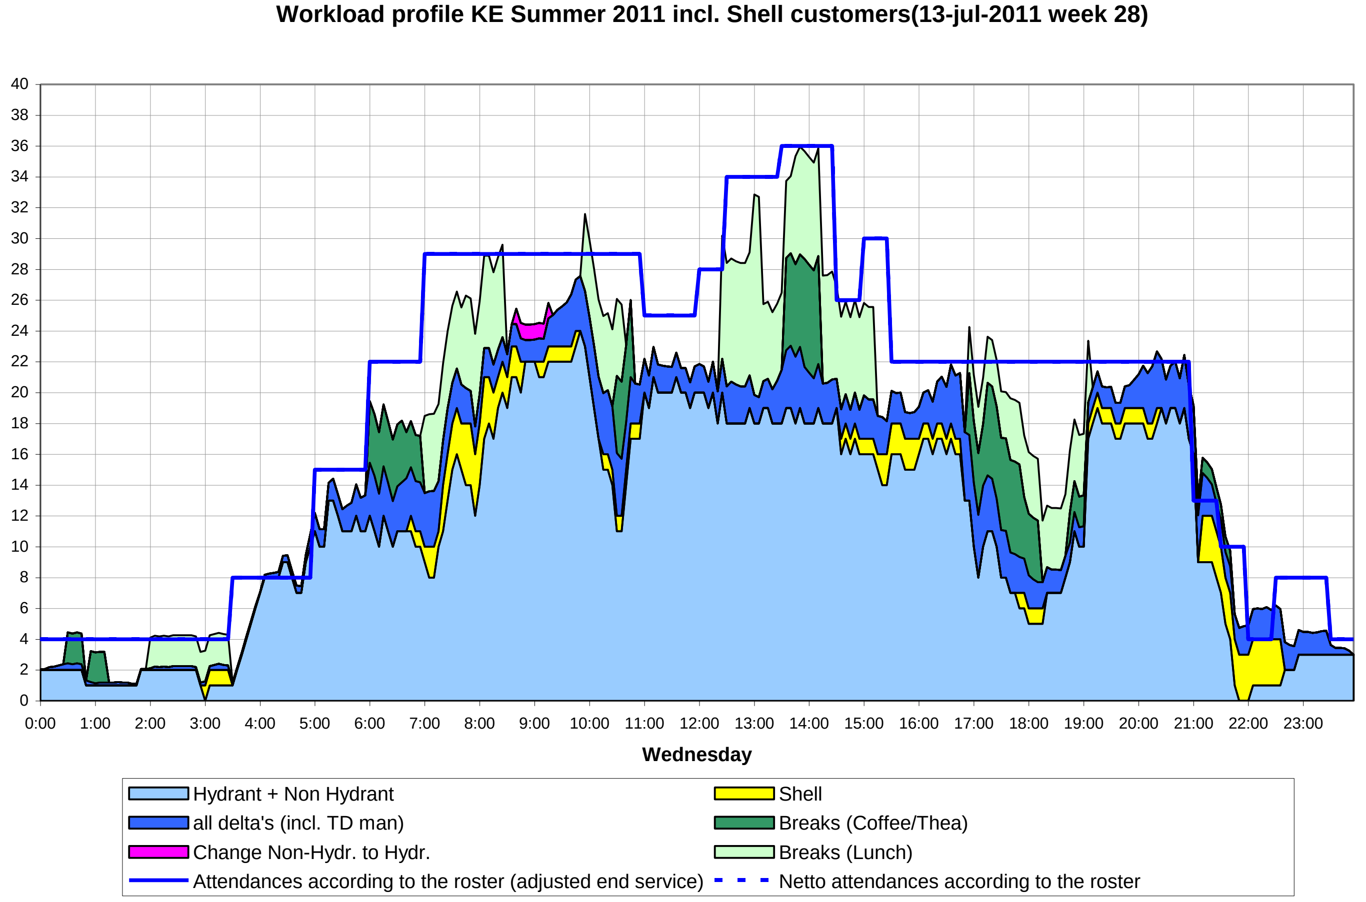

Figure E.1: The workload profile (including planned attendance) used for Wednesday of the OPC week for the summer of 2011. The surfaces are additive.

Appendix F: Histograms with fitted probability distribution

Arrival Punctuality

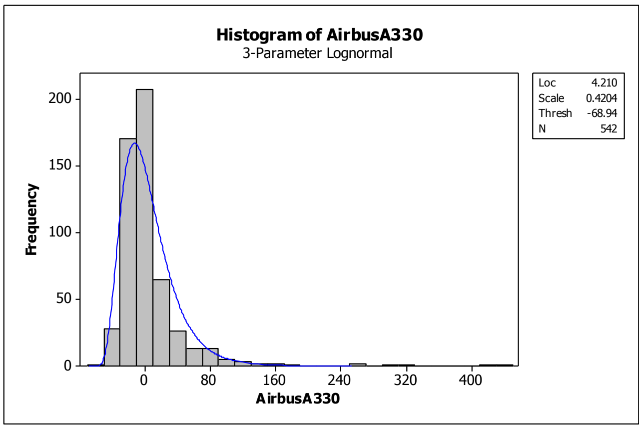

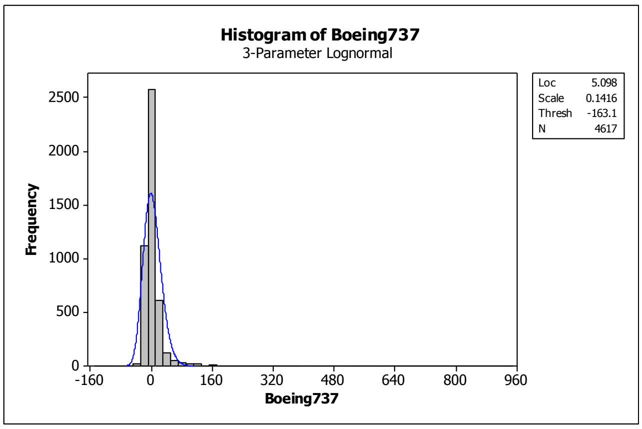

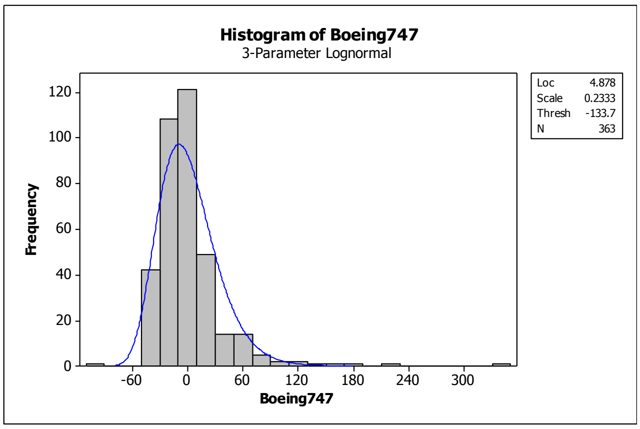

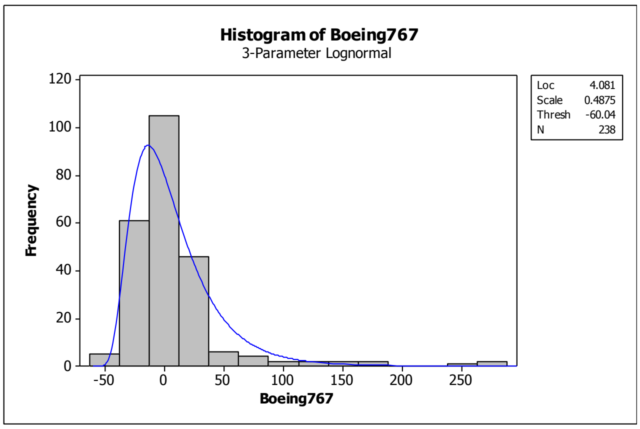

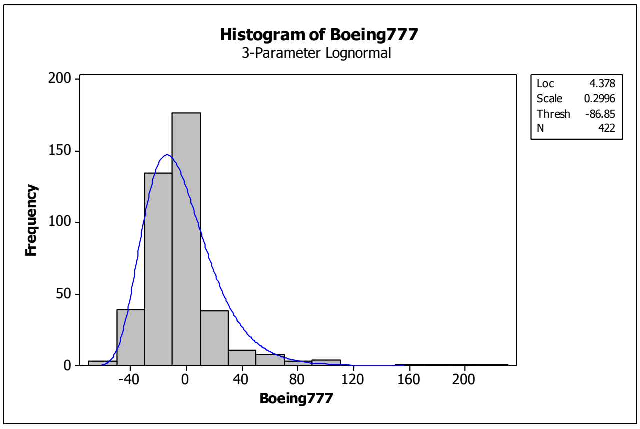

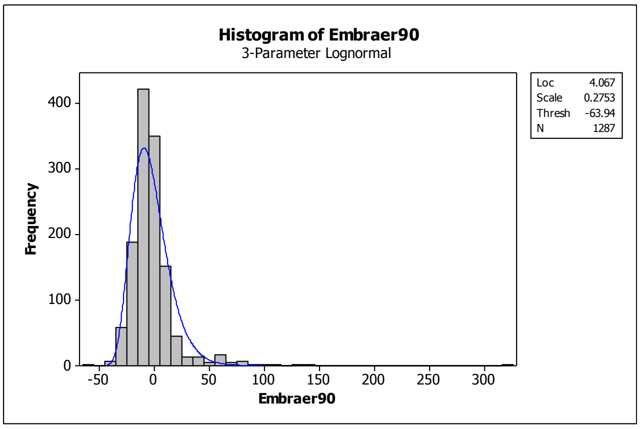

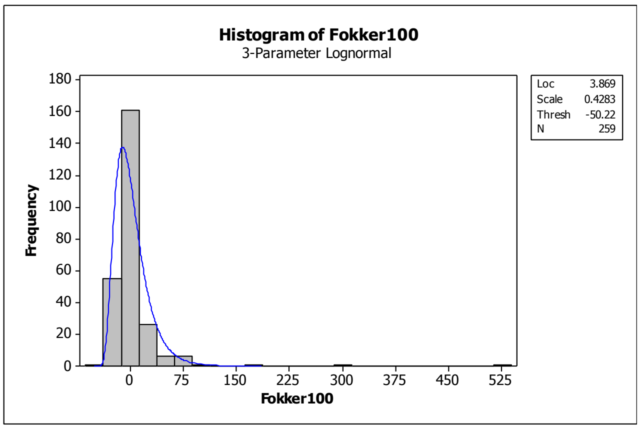

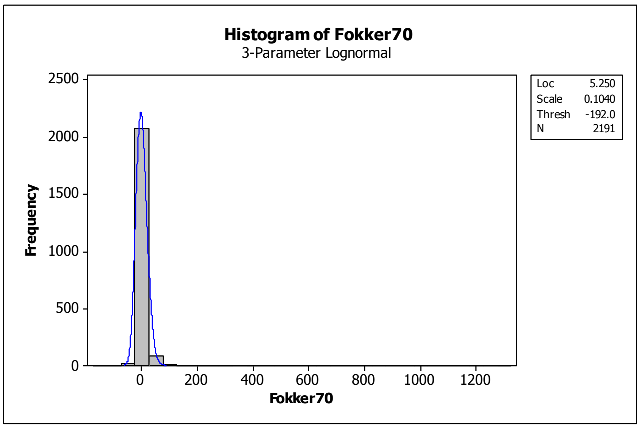

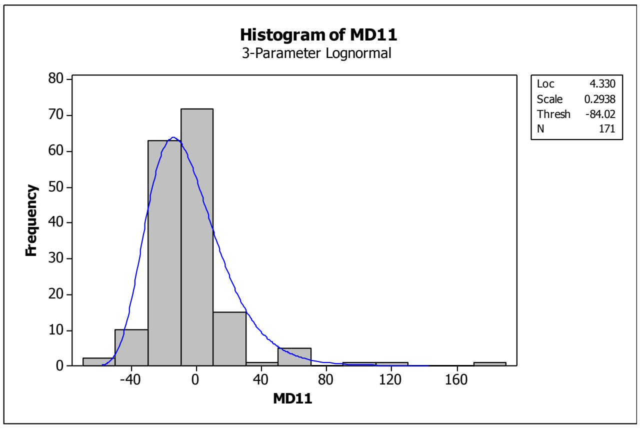

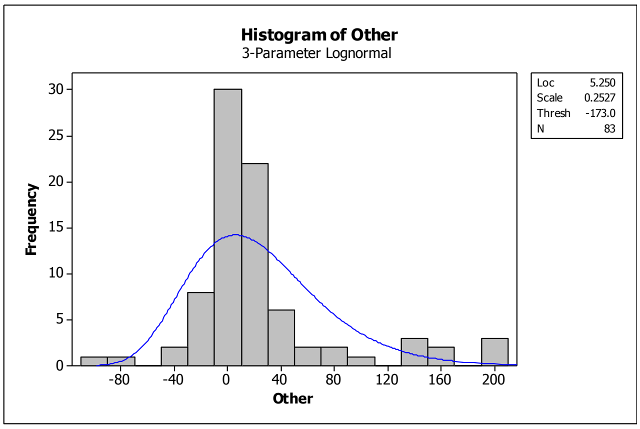

Refuelling Process Times

Pre-fuel

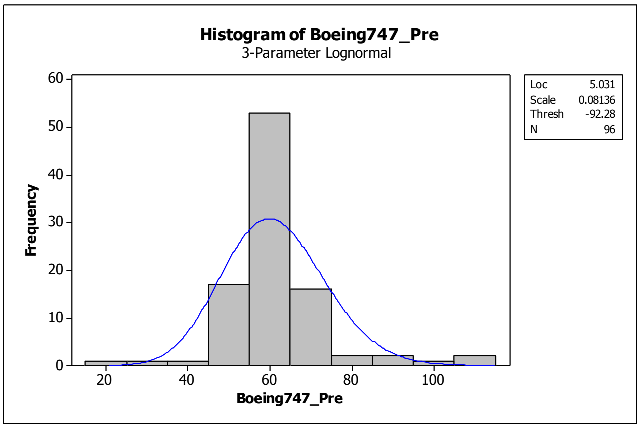

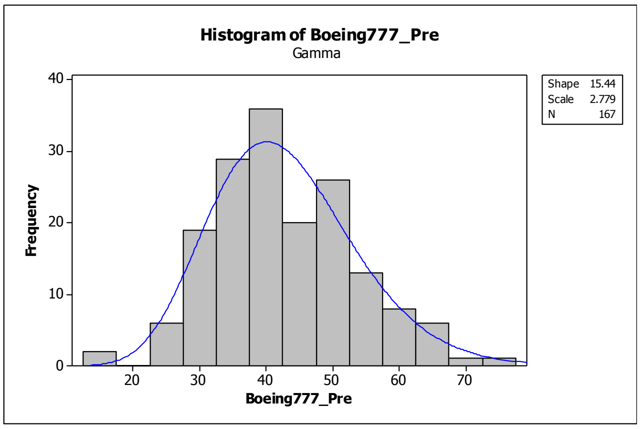

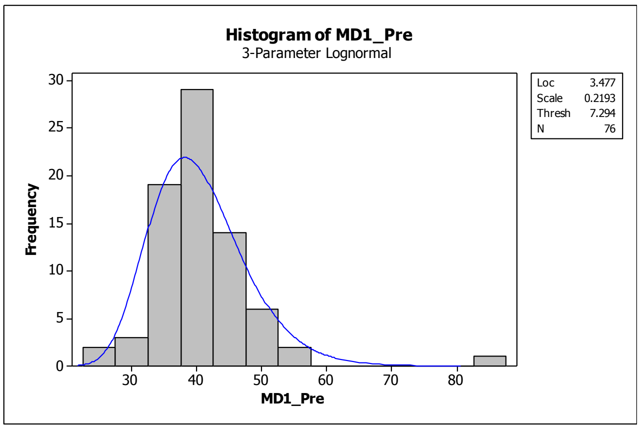

Final fuel after pre-fuel

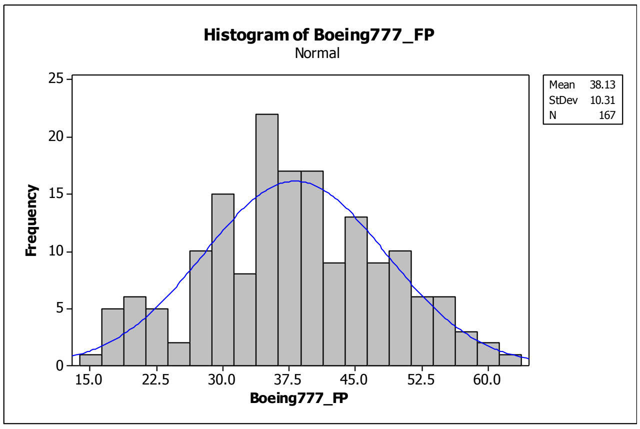

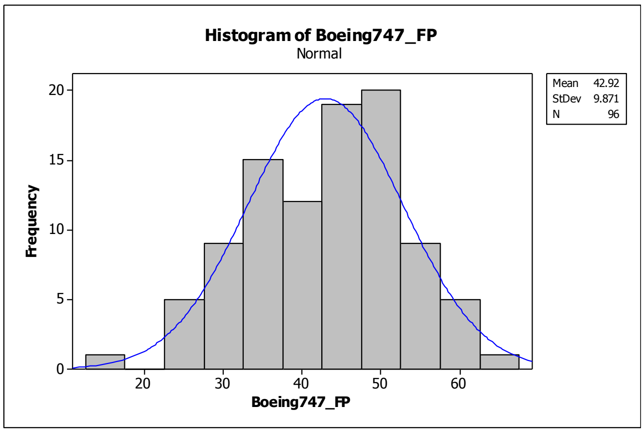

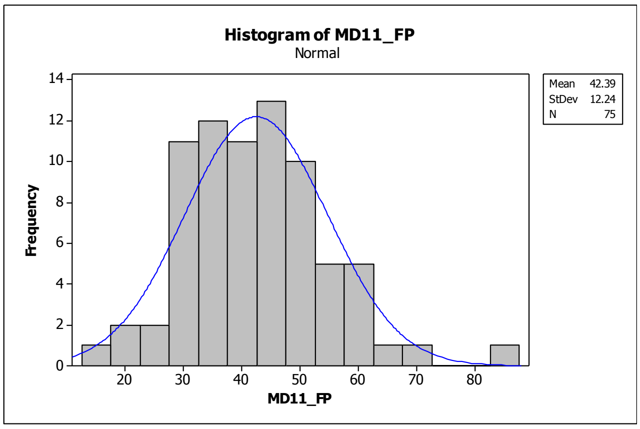

Final fuel without pre-fuel

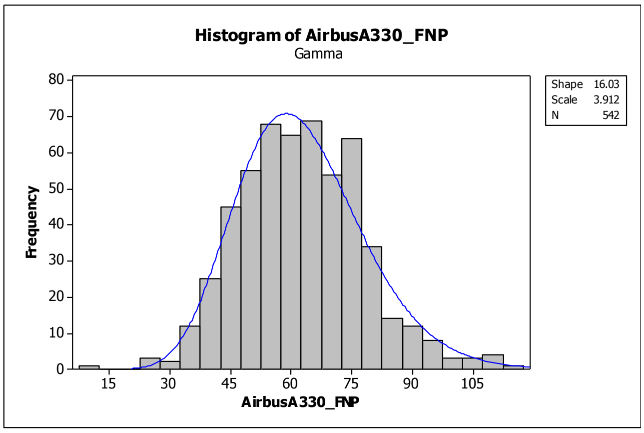

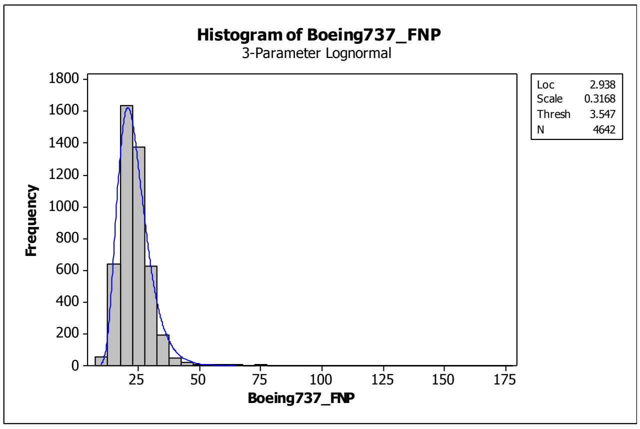

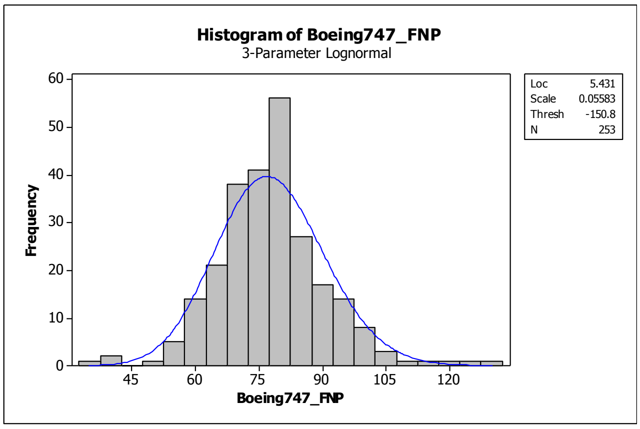

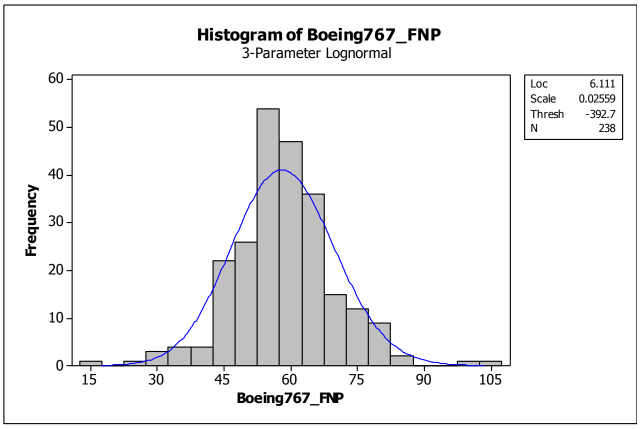

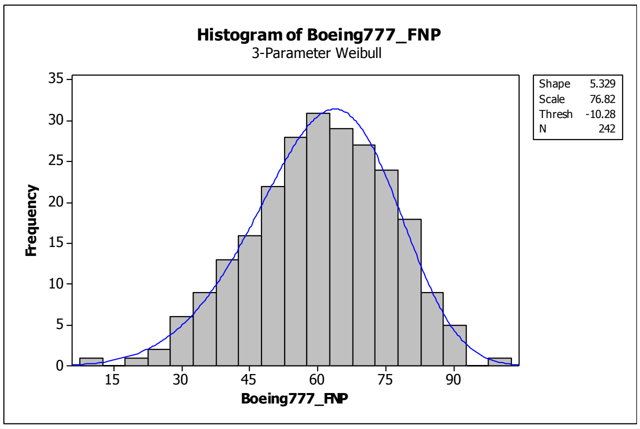

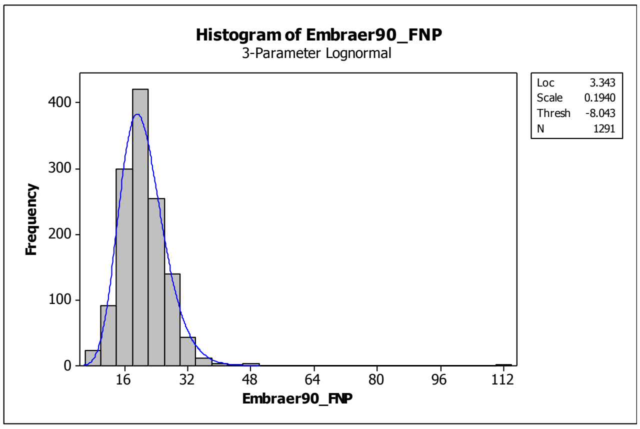

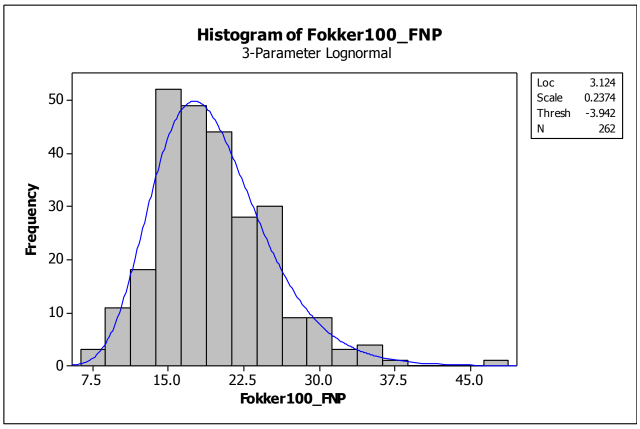

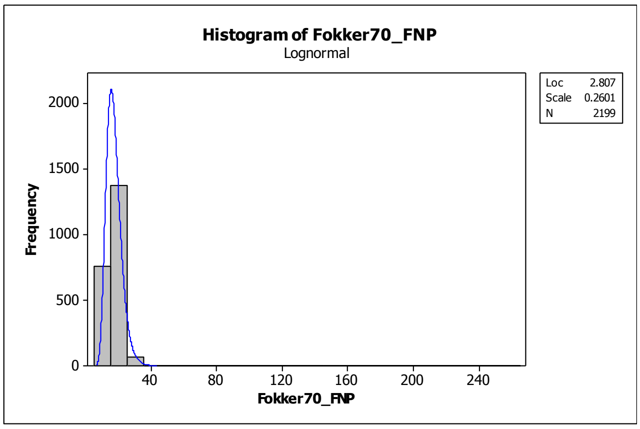

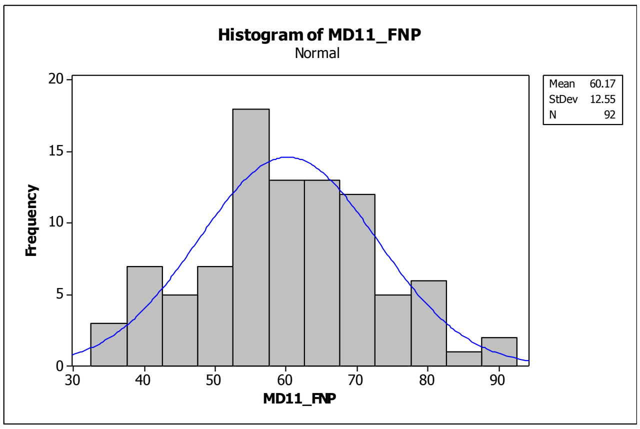
Histogram of MD11_FNP
Normal
Mean 60.17
StDev 12.55
N 92
Frequency
MD11_FNP

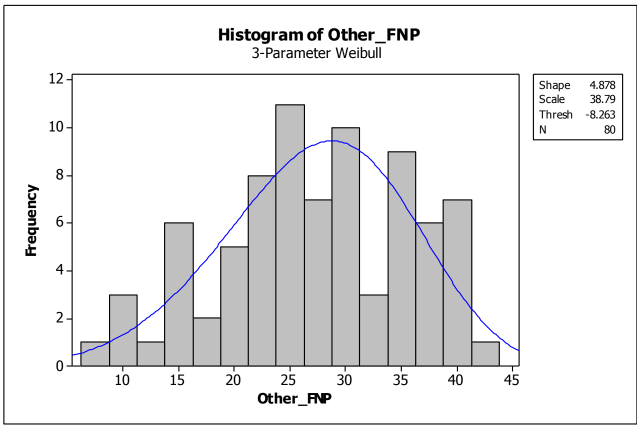
Histogram of Other_FNP
3-Parameter Weibull
Shape 4.878
Scale 38.79
Thresh -8.263
N 80
Frequency
Other_FNP