An Integrated Planning Approach: Maintenance Task Scheduling Optimization

MSc Thesis

D.A.J. Peschier

An Integrated Planning Approach: Maintenance Task Scheduling Optimization

MSc Thesis

by

D.A.J. Peschier

Student number:	4232739	
Field of Study:	MSc. Flight Performance and Propulsion	
Department:	Aerospace Engineering	
Project duration:	April 2018 - February 2019	
Supervision:	Dr. Ir. W.J.C. Verhagen	TU Delft
	Dr. Ir. B.F. Santos	TU Delft

Preface

The present study has been conducted as the author's Master thesis research in Aerospace Engineering at Delft University of Technology and as per assignment by an undisclosed aircraft maintenance MRO Service Provider. The study was conducted over the period from April 2018 to February 2019. The study was issued by the MRO officials as part of an ongoing effort to improve the maintenance planning process through optimization, automation and integration. The present study concerns the development of a task scheduling model which has been designed to provide greater insight into the relation between the actual task specific base maintenance workload and the actual available resource capacity of the maintenance checks and has been applied to the maintenance of the wide-body fleet of one of the MRO's customers. This insight is of particular interest to the modification tasks as there currently is not a sufficient means to provide task specific insight into the modification workload and available resource capacity on a tactical planning horizon. The task allocation model optimizes for maximum resource and interval efficiency. The present report explains the formulation of the model based on the requirements and needs of the MRO, presents the insights it provides, and aims to identify the potential benefits of these insights for the MRO's maintenance planning and airline operations in general.

With no prior knowledge of operational research optimization models, I have aimed to develop the model to the best of my knowledge of the needs of the MRO and of mathematical model formulations. As such I have used the best information available to me in the form of books, the web, experts and my own common sense. I proudly yet humbly present you with my thesis study. I hope that it may proof to be of value to you.

Please note that per the request of Delft University of Technology and the MRO this report has been rewritten to exclude any confidential information. As such, all sensitive data has been purposely replaced by letter combinations, blurred or removed from this report.

D.A.J. Peschier
Delft, January 2019

Summary

Maintenance is one of the major contributors to aircraft fleet unavailability. As such, optimization of maintenance programs is essential to the profitability and competitive ability of airline operators, especially legacy carriers. The MRO service provider for whom this research was conducted is responsible for all maintenance programs of a major European airline, including its current fleet of thirteen wide-body aircraft.
However, the current tactical maintenance planning provides insufficient insight into the relation between the required workload and the actual available capacity. This results in suboptimal maintenance task and resource scheduling, which potentially can have a negative impact on the airline's fleet availability. This problem appears to be particularly true for the modification workload for the airline's wide-body fleet, of which the actual workload continues to exceed the expected available resource capacity.
The present research aims to provide better insight into the required maintenance production capacity for the complete workload of the airline's wide-body fleet over a tactical planning horizon by developing a resource and task scheduling optimization model.

The task and resource scheduling optimization model developed in the present research has been formulated as a deterministic mixed integer linear programming model which optimizes for resource capacity efficiency and task interval efficiency. The model has been extended with a iterative solution technique resulting in multiple consecutive optimization runs, where the output of the previous run is used as an input for the subsequent run. A CPLEX solver has been used for the optimization.

The scope of the present research includes all routine base maintenance tasks for the routine A- and B-checks, as well as the active workload of all modification tasks applicable to the airline's wide-body fleet. Additionally, the model utilizes the MRO's Non Routine Prediction algorithm to include an estimate of the non routine workload for all routine tasks. Furthermore, all letter checks are included in the model (A-, B-, C-checks) as well as the optional extra SB-checks. The routine maintenance tasks of the C-checks have not been included in the present model. Furthermore, the model's scope is limited to base maintenance only, as such the model does not include any maintenance tasks and checks smaller than A-checks.
The model is formulated such that it aims to find the most efficient allocation of all routine out-of-phase and modification tasks to the various maintenance checks given the task due date constraints and the limited available resource capacity of the maintenance checks over a four year planning horizon.

The scope and formulation of the model set it apart from task scheduling models in existing literature. Contrary to conventional aircraft maintenance task scheduling literature, the present research includes *all* base maintenance task types: routine, non routine and modification tasks. The additional inclusion of modification tasks sets this research apart from the current state-of-the-art.
Furthermore, whereas conventional resource constrained task scheduling models minimize the maximum completion time of all tasks subject to due date and precedence constraints of tasks and resource capacity constraints, the present model optimizes for maximum resource efficiency. As such, precedence constraints between tasks are ignored. Moreover, the model includes soft resource capacity constraints. Penalties are incurred when the allocated capacity of the tasks either exceeds or runs short of the available resource capacity. In addition to the allocation of tasks to resources, the model also considers the allocation of resources to aircraft, combining resource and task allocation in one model. Thus, the model adds to the existing literature on applied task scheduling and resource allocation models.

The model successfully produces an optimized, integral, task specific planning of the base maintenance tasks and checks over a four year time period for the airline's wide-body fleet. Up until now an equivalent task specific tactical maintenance planning did not exist within the MRO. As such, the present model's schedule provides valuable insights which can benefit the MRO and the operator.
The schedule for the coming four years highlights the need for the MRO to take action to avoid maintenance disruptions, as the currently known workload exceeds the available resource capacity. The model provides suggestions on when and where to schedule additional resource capacity for each aircraft registration in the

form of extra maintenance checks to facilitate the excessive workload. Additionally, the model may also be used to assess the effectivity of proposed measures by the MRO planning policy makers, e.g. an increase of the norm for the letter checks. As such, the additional insight provided by the model will help to reduce the potentially negative impact of poor planning on the airline's fleet availability.
Furthermore, the insight provided by the model's schedule is considered particularly valuable for scheduling modification tasks. It is anticipated that the model's schedule will help the MRO to implement modifications more quickly, yielding operational and financial benefits for the airline. Moreover, the schedule will help the MRO to provide a more firm scheduling commitment for modifications and it will assist the MRO in negotiations with the operator on accepting future modification workloads.

Although the present model already provides valuable insights to the MRO and the airline, analysis have also shown that the model is still limited by its current design and scope.

Firstly, the model's iterative solution technique restricts the optimality of the model's solution. The iterative solution technique was introduced to manage the recurring routine out-of-phase tasks. However, the resulting successive optimization runs yield an overall suboptimal task allocation. It is therefore recommended that the model's formulation be adjusted to enable a single optimization run, which will further optimize the task schedule.
Furthermore, the current formulation of the excess hour penalty is such that in case of very large excess hours, the model converges towards a suboptimal solution. As such, it is recommended to reformulate the post-slope of the piecewise linear functions which control the excess hour penalty. This will ensure that the model converges to an optimal solution even in case of very large excess hours.

In addition to these recommended design changes, the present research has also shown the need to expand the current scope of the model. Most importantly, the present model is restricted to a resource capacity optimization only. As such it does not include any precedence relations between tasks or any form of ground time optimization. This restriction currently results in a proposed schedule that is not always practically feasible. Thus, the proposed model extension to include ground time optimization and additional scheduling constraints is essential to improve the accuracy and feasibility of the model's schedule.
Moreover, the MRO's non routine prediction algorithm does not provide an accurate enough prediction of the non routine workload. This in turn has a negative effect on the reliability of the model's schedule. As such, it is recommended that further investments are made to improve the accuracy of the non routine predictor.
Other recommended extensions to the model include the addition of C-check type routine maintenance tasks, such that the model provides a truly complete overview of the base maintenance workload. Moreover, it is recommended to replace the current approximation method for maintenance check dates with an optimization of the maintenance check allocation. The MRO already uses an optimization algorithm for its A-checks. However, this model is limited to A-checks and a three month planning horizon only and as such not usable by the present model. Thus, it is recommended to expand the scope of the existing maintenance check optimization model to include all letter checks and to also cover a four year planning horizon. This would allow for the present model to use the outcome of such a maintenance check optimization model as an input for the task allocation. This addition will improve the operational consistency of the present model and hence its usability as an actual tool to generate an optimized maintenance planning.

In conclusion, the present model already provides valuable insights that will potentially contribute to cost reductions for both the MRO and the airline. Moreover, the model performs satisfactory within the limited scope of the present research. However, limitations in the model design and limitations in the model scope restrict the added value of the present model. This limits the outcome of the model to an initial planning suggestion, rather than a maintenance schedule which can fully be depended upon. Nevertheless, the present research highlights the potential of an integrated, automated and fully optimized task schedule. It is therefore strongly recommended that the necessary model design improvements and capability extensions be considered and implemented, such that the full potential of the task scheduling optimization model may be used to benefit the efficiency of the MRO's maintenance planning and subsequently improve the availability and profitability of the airline's wide-body fleet.

Contents

List of Figures

List of Tables

List of Acronyms

AD	Airworthiness Directive
AOG	Aircraft On Ground
AMM	Aircraft Maintenance Manual
AMP	Aircraft Maintenance Program
ASK	Available Seat Kilometers
ATA	Air Transport Association of America
AWL	Airworthiness Limitations
CASK	Cost per Available Seat Kilometer
CASM	Cost per Available Seat Mile
CMP	Configuration, Maintenance and Procedures
CMR	Certification Maintenance Requirement
DD	Deferred Defect
DMC	Direct Cost of Maintenance
EASA	European Aviation Safety Agency
EO	Engineering Order
ETOPS	Extended-range Twin-engine Operational Performance Standards
FAA	Federal Aviation Administration
FC	Flight Cycles
FH	Flight Hours
FJSSP	Flexible Job Shop Scheduling Problem
GA	Genetic Algorithm
IATA	International Air Transport Association
IMC	Indirect Cost of Maintenance
JIC	Job Instruction Card
JSSP	Job Shop Scheduling Problem
KPI	Key Performance Indicator
LRU	Line Replacement Units
MCC	Maintenance Control Center
MIS	Maintenance Information System
MPD	Maintenance Planning Document
MRBR	Maintenance Review Board Report
MRI	Maintenance Requirement Item
MRO	Maintenance Repair and Overhaul
MSG	Maintenance Steering Guide
NR	Non Routine
NRP	Non Routine Predictor
OAM	Original Aircraft Manufacturer
OEM	Original Equipment Manufacturer
OMP	Operator’s Maintenance Program
OOP	Out-Of-Phase
ORP	Operator Requirement Proposal
PASK	Profit per Available Seat Kilometer
PSFC	Planning Scheduling and Fleet Control
PWC	Planning and Workload Control
PWLF	Piecewise Linear Function
RAP	Resource Allocation Problem
RASK	Revenue per Available Seat Kilometer
RASM	Revenue per Available Seat Mile
RCPSP	Resource Constrained Project Scheduling Problem
SB	Service Bulletin
TLW	Time Limited Work

1
Introduction

The following Chapter introduces the research study which has been conducted into base maintenance task scheduling optimization at a major MRO Service Provider (referred to as 'the MRO') which provides MRO services to a large European carrier (referred to as 'the airline'). Firstly, Section 1.1 provides some context on the economic need for airlines to optimize their maintenance. Section 1.2 elaborates on the limitations of the current maintenance planning policy at the MRO and thus highlights the practical need for this research. Moreover, this Section also briefly explains the academic relevancy of the present research. Sections 1.3 and 1.4 present the project objective and research scope of the study respectively. The research questions and research approach are discussed in Sections 1.5 and 1.6 respectively. Lastly, Section 1.7, gives a brief overview of the further setup of the report.

1.1. Industry Background

In the current airline market, where low-cost carriers dominate the increasing demand for cheap air travel; where personnel strikes force an airline to its knees; and where fluctuating oil prices eat away at the remaining profit margin, it is predominantly the legacy carriers which are forced to look for new ways to reduce operational costs in order to successfully compete and survive in tomorrows market. According to IATA, airlines around the world spent an average of 9.5% of the total operational cost on MRO services over the year 2016 [4]. These costs do not yet include so-called indirect maintenance cost, such as the cost of aircraft unavailability, or downtime, inherent to aircraft maintenance. Aircraft downtime equates to a further loss of revenue for airlines [10][11]. It is with these significant costs in mind that aircraft maintenance has evolved from a necessity to maintain the highest standards of safety, into a strategic interest for an airline's competitiveness [12][13].

Regardless of strategic concerns however, safety is to be considered first and foremost. As such, all MRO service providers are required to adhere to the stringent maintenance requirements outlined in the so-called MPD, or maintenance planning document. It contains the specifications of maintenance tasks as well as their intervals as identified by the original aircraft and equipment manufacturers (OAMs and OEMs) and as approved and endorsed by regulatory authorities. However, although strict adherence is required, it is up to the MRO service providers to determine the most efficient planning and scheduling of these tasks within the predefined limits specified in the MPD. As such, the MRO's Planning Scheduling and Fleet Control (PSFC) department is charged with the optimization of maintenance planning for a large European legacy airline's wide-body fleet.

Aircraft maintenance is a complicated, multi-disciplinary, optimization problem. The present research focuses on the optimization of maintenance task planning from a resource capacity perspective. As the available labor forces of any MRO service are finite, an optimal allocation of maintenance tasks over time has to be determined such that the available resources are used most efficiently whilst ensuring that all maintenance task are completed within their set intervals. The following Section will describe some of the limitations of the current maintenance planning at the MRO, which the present research aims to improve upon.

1.2. Problem Statement

The MRO's current maintenance task scheduling is limited by various factors which result in a suboptimal planning. First and foremost, the present MRO planning policy is mostly a manual, non-automated process, which is time consuming, subject to errors, and inherently suboptimal. It is therefore considered that an automated task planning optimization planning model will proof to be profitable by yielding a more efficient planning, as well as saving valuable time in the planning process itself.

Secondly, the tactical maintenance planning is based on estimations rather than the task specific workload. This can result in 'unexpected' discrepancies between the actual required maintenance workload and the anticipated workload, and therefore also the available workforce. This subsequently results in a reactive maintenance plan, where additional ground time may be required to cope with the extra workload. In turn, such suboptimal planning performance can also have a negative effect on the fleet availability of the airline.
Furthermore, the accuracy of the maintenance planning could be improved through a better estimate of the non routine workloads. Unscheduled maintenance is a fact of all maintenance operations. Hence, the ability to predict the occurrences and workload of these non routine maintenance tasks can have a profound impact on the accuracy of future workload estimations. Currently, the MRO uses the average hours of the historical non-routine workloads to account for any future unscheduled work. This could be improved through prediction algorithms that estimate the non routine workload based directly on the corresponding routine work.

Finally, the maintenance planning may be further optimized through an integrated approach. The MRO has divided it maintenance workload over different check types, A-, B-, C-checks, depending on the nature of the maintenance tasks involved. As such these checks are also separately planned. Although this works well for routine preventive maintenance tasks, this can result in an inefficient allocation of modification tasks.

The current limitations of the MRO's maintenance task planning may be summarized in the following problem statement:

The MRO's current tactical maintenance planning does not provide the desired insight to efficiently and effectively plan the complete maintenance workload and subsequently the MRO cannot optimally anticipate on the resource capacity required to accommodate these workloads. These limitations can negatively affect the availability and profitability of the airline's fleet.

In addition to the practical need for the present research, the literature study discussed more elaborately in Chapter 2 also shows the academic relevancy of the present study. The study contributes to the literature on applied task scheduling and resource allocation models. The present research adds to the current state-of-the-art by considering the 'complete' workload of base maintenance tasks, including modification tasks. Furthermore, the emphasis of the present research on resource capacity efficiency yields a completely different model formulation compared to the models in conventional task scheduling literature.

1.3. Project Objective

Based on the aforementioned problem statement, the following project objective has been defined:

To provide greater insight into the required maintenance production capacity for the complete workload of an airline's wide-body fleet over a tactical planning horizon by developing a resource and task scheduling optimization model

Provided that this objective is not just limited to the present research, but rather refers to ongoing improvements to the current maintenance planning at the MRO, it is referred to as the project objective, rather than the research objective.

1.4. Research Scope

As stated in the research objective, this research focuses on optimizing the task scheduling for a large European airline's wide-body fleet. To date this fleet consists of thirteen aircraft.
An important contribution of the study relates to the fact that it includes the "complete workload" of base

maintenance tasks. As such, it includes routine maintenance tasks (block tasks and out-of-phase tasks), non routine workload and modification tasks. Moreover, the study includes all major base maintenance check types, namely the A-, B- and C-checks, as per the MRO's definitions of these checks. A more elaborate explanation of these checks is provided in Chapter 3. It is important to note that although routine and non routine tasks are included for the A- and B-checks, a simplification was introduced in the model with regards to the C-checks for the sake of time. Thus, the routine and non routine tasks of C-checks have not been included in the scope of the present research.
All routine and non routine task information has been taken directly from the MRO's maintenance information system, MIS. Information regarding the modification has been extracted from the dedicated Data Exchange Platform used for tracking and planning modification within the MRO. All maintenance tasks and checks smaller than A-checks are not included in the model, e.g. line maintenance slots and tasks, deferred defects, etc.

1.5. Research Questions

Based on the aforementioned research objective, the following research question and sub questions have been formulated:

What is the added value of the integrated and automated maintenance task scheduling optimization model applied to a tactical planning horizon for all base maintenance of the wide-body fleet of a large European airline?

- What insights does the task scheduling model provide in terms of the required and available resource capacity?
- How accurate is the planning provided by the task scheduling model, in terms of:
  - ...the accuracy of the occurrence of maintenance tasks?
  - ...the accuracy of the scheduled labor hours of the maintenance tasks?
  - ...the accuracy of task scheduling?
- How efficient is the planning provided by the task scheduling model, in terms of:
  - ...the efficiency of using the available resource capacity?
  - ...the efficiency of using the available task interval?
- How is the performance of the task scheduling model, in terms of:
  - ...the optimality of its solution?
  - ...its robustness?
  - ...its run times?

1.6. Research Approach

The present research was approached as an operational research problem. The model developed in the study has been formulated as a deterministic, mixed integer linear programming model. An iterative loop has been built around the optimization model. As such, the model may be classed as a heuristic model, with an approximate optimal solution. A detailed explanation of the model formulation may be found in Chapter 4.
Based on the literature study performed (see Chapter 2), the model formulation is inspired after resource constrained project scheduling problems with some fundamental differences. The model considers the available capacity of maintenance checks as its resources, with the maintenance tasks as the activities or tasks to be allocated, each with a specified required capacity. Contrary to conventional resource constrained project scheduling problems, the model optimizes for maximum resource efficiency rather than a minimization of the maximum completion time of the tasks. As such, the model does not consider precedence relations between tasks. Furthermore, the resource capacity is modelled as a soft constraint, contrary to the contemporary hard constrained periodic resource capacity. The formulation, assumptions and limitations of the model are discussed in greater detail in Chapters 4 to 6.
The model has been formulated in Python 3.5 AMPL with CPLEX 12.7.1 as solver.

1.7. Setup of the Report

The remainder of the report is structured as follows. Chapter 2 will provide a summary of the literature study which has been conducted as part of this research. It elaborates on the state-of-the-art of task scheduling optimization problems and identifies the academic relevancy of this study in greater detail. Chapter 3 provides some additional background information specific to task scheduling and planning at the MRO. A detailed explanation of the model formulation, including the mathematical formulation and its application to the specific case for the MRO is provided in Chapter 4. Chapters 5 provides a detailed description of the formulation of the penalties included in the model formulation. Moreover, Chapter 6 provides a detailed description of the data preparation process which prepares the inputs for the optimization. It elaborates on the assumptions which have been made for the various maintenance tasks and checks and how these relate to best practices at the MRO. The model results are presented and discussed in Chapter 7. The validation of the model is subsequently discussed in Chapter 8. The sensitivity study results and discussion may be found in Chapter 9. The answers to the research questions as presented in this Chapter are discussed at the conclusion of the report in Chapter 10. Finally, the recommendations based on the present research are provided in Chapter 11.

2
Literature Study

A literature study was conducted to discover the state-of-the-art in the field of task scheduling optimization and its applications. This Chapter will first provide a brief introduction of the wider field of aircraft maintenance optimization, which includes the various sub-problems related to aircraft maintenance optimization research. A more detailed discussion on the literature related to task scheduling optimization is discussed in Section 2.2. Section 8.4 provides an elaborate discussion on the differences of the present study with respect to the current literature and as such the expected academic contribution of this research.

2.1. Aircraft Maintenance Optimization

As mentioned in the Introduction, "aircraft maintenance has evolved from a necessary into a strategic concern" [12]. Aircraft availability has to be maximized for airlines to be profitable and competitive in today's market. With aircraft base maintenance as one of the primary drivers of aircraft unavailability, efficient and effective maintenance programs directed towards maximum fleet availability are essential. "Air carriers are constantly striving to achieve high standards of safety and simultaneously to attain an increased level of availability performance at minimal cost. This needs to be supported through an effective maintenance program which has a major impact on the availability performance and which ultimately can enhance the aircraft's capability to meet market demands at the lowest possible cost" [13].

2.1.1. Sub-problems in Aircraft Maintenance Optimization

There is much literature related to aircraft maintenance optimization. According to Dinis et al., the literature may be divided into three main 'sub-problems' [14]:

1. **Capacity Planning**
2. **Spare Part Forecasting and Inventory Management**
3. **Task Scheduling and Resource Allocation**

Capacity planning problems involve the organization of a workforce to meet the future labor demand [14]. Provided the cost of personnel in any industry, including aviation and aircraft maintenance, the objective is to minimize the size of the workforce [14][15]. "Allocating the right amount of workforce in the right combination of skills and experience levels to eradicate any shortfalls or surpluses is essential for the continuation of efficient and effective maintenance in any organization" [16]. Determining the optimal workforce composition and size is complicated due to fluctuating demands and skill-,license-, and training requirements [15].

Spare part planning problems concern when and how many spare parts are required to ensure that the maintenance demand is met, whilst minimizing the involved cost. "An excess of spare parts inventory leads to a high holding cost and impedes cash flows, wheras inadequate spare parts can result in costly flight cancellations or delays with a negative impact on airline performance" [17]. Spare part problems are further complicated due to the unpredictable demand resulting from part failures, finite warehouse space, limited budgets and vendor agreements [17].

Dinis et al. provide numerous examples from literature in each of these three sub-problems [14]. Moreover, Van den Bergh et al. also provide a comprehensive review on aircraft maintenance operations [18]. Provided the emphasis of the present study on task allocation, a more elaborate literature review has been conducted on this sub-problem, which will be discussed in 2.2.
In addition to these three sub-problems, there is substantial literature related to other aircraft maintenance optimization problems, including fleet assignment problems and aircraft routing [18]. These problems consider the optimization of flight networks subject to maintenance demands and are considered outside of the scope of the present study.

2.1.2. Mathematical Model Formulations

Within the field of aircraft maintenance optimization literature, and within so-called operational research problems in general, there are different types of model formulations. The three most common types of model formulations in literature are:

1. **Linear Programming:** consists of an objective function, decision variables and constraints and required that all mathematical functions (objective function and constraints) are strictly linear [19].
2. **Non Linear Programming:** although for many problems linear relations suffice in modeling, non linear programming is applied when not all of the mathematical functions can be described linearly [19].
3. **Dynamic Programming:** is defined very differently compared to linear and non-linear programming. Dynamic programming problems consists of stages, states, policy decisions and cost functions. A problem is divided into stages, e.g. time periods [20]. The states "represents the possible conditions in which a system might be at a given stage" [19]. The idea is to determine the optimal policy stage by stage through a recursive formulation, such that one obtains the optimal policy ofr the whole problem [19].

The above mentioned modeling methods are all very different. The present study uses linear programming, as will be further presented in Chapter 4. The choice for linear programming has been based on the fact that it is a widely accepted modeling method and based on the logic involved which was considered more intuitive to the author.

As stated by Hillier and Lieberman, linear programming is a widely accepted approach in operational research problems and it is most commonly applied to "the general problem of allocating limited resources among competing activities in a best possible way" [19]. Thus, linear programming is commonly used for problems similar to the present study. Moreover, provided that most problems may be described by linear relations, non-linear programming was not considered any further for this research. Line linear programming, dynamic programming is also a commonly used and efficient method in resource allocation problems. However, the logic involved with linear programming appeared more intuitive to the author when compared with dynamic programming. Provided that the author had no knowledge of operational research models prior to this study and provided the time constraint on the study, the more intuitive and widely accepted linear programming approach was considered the wiser choice.

2.2. Task Scheduling Optimization

This Section will elaborate on literature related to task scheduling optimization. "Scheduling may broadly be defined as the allocation of resources to tasks over time in such a way that a predefined performance measure is optimized. From the viewpoint of production scheduling, the resources and tasks are commonly referred to as machines and jobs and the commonly used performance measure is the completion times of jobs" [1]. Three of the most common task scheduling problems are addressed:

1. **Job shop scheduling problems**
2. **Resource constrained project scheduling problems**
3. **Resource allocation problems**

The literature on these three categories of task scheduling provides relevant information on problem and model formulations for the present study.

2.2.1. Job Shop Scheduling Problems

The job shop scheduling problem (JSSP) represents one of the three categories in task scheduling optimization problems. "The job shop scheduling problem is to determine a schedule of jobs that have pre-specified operation sequences in a multimachine environment" [21].

The conventional JSSP consists of *M* machines and *N* jobs, where each of these jobs is composed of a number of operations or subtasks. The sequence in which the operations of each job are to be executed has been predefined and is known. Each of these operations has to be performed by a specific machine. Moreover, the process time of each operation is also predefined and known. No preemption of operations is allowed, meaning that once an operation has been started, it may not be interrupted. After a machine has completed an operation, it becomes available again for processing other operations. Moreover, a machine may only process one operation at a time. Thus, the problem of a JSSP is in what order to allocate the operations of the various jobs to the machines, such that the maximum completion time of the jobs (makespan) is minimized.

Considered as a "classical combinational optimization problem in real-world manufacturing systems", JSSPs have all sorts of different applications [22]. Thornblad et al. use a JSSP-based model to produce a production schedule to help reduce the number of delays and lead times for a manufacturing plant. Alternatively, Ahire et al. applied JSSP to formulate an aircraft maintenance task allocation model which optimizes the allocation of tasks to personnel with different skill sets such that the makespan of the tasks is minimized.

The application to real-world problems also results in variation of JSSP formulations. The manufacturing plant model of Thornblad et al. includes multi-purpose machines [23]. As a result of this the not only determines the optimal order in which the operations (subtasks) should be scheduled, but also determines which operation to allocate to which machine. This variation of the JSSP is known as the flexible JSSP, or FJSSP, and is common in literature as it more closely represents real-world scheduling problems. Ozguven et al. further expanded their FJSSP model to include different options of sequences of operations for each job. They refer to this as an FJJSP with process plan flexibility, or FJSSP-FPP [1].

An example of an optimized schedule for one of the problems solved by the model of Ozguven et al. is shown in Figure 2.1. This particular problem consists of five jobs with two optionable process plans (sequences) per job, at most six operations per job. The problem includes five machines and each operation had two machines to which it could be allocated [1]. Looking at Figure 2.1, the first job on machine four (M4), O_{524}, indicates that the first operation for M4 is operation 4 of process plan 2 of job 5. This particular process plan for job five consists of operation 4 followed by operation 6. As can be seen on M1, denoted by O_{526}, machine 1 completes job 5 by operation 6 after M4 has completed operation 4. As can be seen from the completion time on M5, the makespan for this particular problem instance is 193 hours [1].

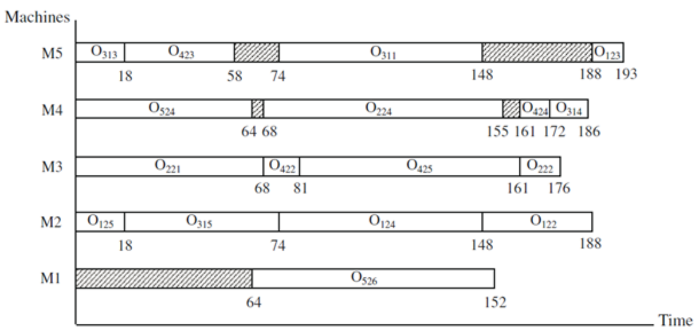

Figure 2.1: Example of one of the schedules of the FJSSP-FPP model of Ozguven et al. in which each of the jobs consisted of several operations (subtasks) which in turn had various optional sequences in which they could be executed. The allocation of the operations to the machines is optimized for minimum makespan [1]

Although the minimization of the makespan is the most common objective, JSSPs may also have other objectives [24]. Yang et al. formulated a JSSP model in which the objective function consists of two penalty terms which are to be minimized [24]. In their work Yang et al. assume that each job has both a due date and a deadline. The due dates are the "desired completion dates of jobs given by the customer" and as such have

been modeled as a soft constraint which may be violated [24]. The deadlines on the other hand are modeled as hard constraints which cannot be violated as it is assumed that beyond these dates "the manufacturer will lose the customer" [24]. To discourage the model from completing jobs passed their respective due dates, the model of Yang et al. incurs a tardiness penalty proportional to the amount by which the due date is exceeded. Furthermore, the model also incurs an earliness penalty in case a job is completed earlier than its due date. This latter penalty discourages potential inventory costs. Thus, the objective of the model of Yang et al. is to minimize the earliness and tardiness penalties of all jobs [24].
As aforementioned, the model of Thornblad et al. minimizes the tardiness and throughout times of all jobs, such as to avoid production delays and maximize machine utilization [23].

The various linear programming models used for JSSPs in literature have things in common as well as great differences. Many models are either mixed integer linear programming models (MILP) or integer linear programming models (ILP). Often these models have a set of binary decision variables, for example for the allocation of operations to machines in case of FJSSPs [23][21]. Furthermore, a set of integer decision variables is often used to define the starting times and completion times of operations and jobs [21][24][1]. These decision variables are subsequently either directly or indirectly part of the objective function of the models for which minimize the makespan. In the model of Yang et al. two additional integer decision variables are formulated for the earliness and tardiness penalties, which are in turn dependent on the value of the integer decision variables for the completion time.
In the model of Ozguven et al. an additional set of binary decision variables is added for the decision of a process plan (sequence of operations) for each job [1]. Furthermore, Ozguven et al. defined another set of bivnary decision variables to define precedence relations between tasks. A similar FJSSP model defined by Fattahi et al. also uses a set of binary decision variables for precedence relations between tasks [21]. However, contrary to the model of Ozguven et al. where the binary decision variables simply state whether or not an operation precedes another operation, Fattahi et al. formulated the binary decision variables to indicate whether or not an operation has a certain priority position on a machine or not. Ozguven et al. showed that their simpler formulation greatly reduces the number of decision variables of the model which consequently corresponds to a shorter model runtime [1]. In general it may be concluded that the great diversity of problems which are approached as JSSPs results in an equal diversity of model formulations.

In addition to the diversity of model formulations, literature shows a great diversity in the solvers which have been used for JSSPs. Thornblad et al. and Ozguven et al. use CPLEX to derive an exact solution to the problems [23][1]. Yang et al. compare the approximate optimization of 180 problem instances using an enhanced genetic algorithm (EGA) with a MILP model which is solved exactly by the commercial solver FICO Xpress [24]. Similarly, Fattahi et al. compare the performance of an exact branch and bound solution algorithm (Lingo) with a range of different heuristic algorithms based on simulated annealing and tabu search algorithms [21].
In general, the exact solution algorithms require substantially more computational time for larger and more complex problems to derive the optimal solution to a problem. Often times, the time limitations imposed on these exact solvers prevents them from finding an optimal solution for larger scale or more complex problems. Approximate solution methods on the other hand are able to derive nearly optimal solutions even for very large problem sizes within less computational time. However, these approximate solutions do show a reduction in solution accuracy when compared with the optimal solution. Ahire et al. demonstrate that their evolutionary strategy algorithm is able "to solve large-scale complex problems within realistic computational efforts", including as many as 500 jobs and a variety of workforce skills, sizes and other constraints [25]. The work of Yang et al. also concludes with the superiority of their EGA over the performance of the exact solution algorithm for all problem instances considered. The heuristic algorithms of Fattahi et al. produced approximate solutions for all problem instances within substantially shorter computational times and with acceptable deviations from the optimal solution.

2.2.2. Resource Constrained Project Scheduling Problems

Like the JSSPs, the resource constrained project scheduling problems (RCPSPs) represent a common type of task scheduling optimization problems [26][27]. In fact, "the job-shop scheduling problem is a special type of RCPSP" known as machine scheduling, in which the resources correspond to machines [28][2].

Like the JSSPs, RCPSPs concern scheduling problems in which a set of dependent activities has to be allocated to a finite number of resources in such a way that some performance parameter is optimized, e.g. makespan, throughput time, tardiness penalties etc. [28][29]. Furthermore, like the JSSPs, the tasks included in RCPSPs generally have precedence relations and predefined processing times which need to be satisfied in the schedules. However, unlike the JSSPs, the tasks in RCPSPs also have a finite required resource capacity which is predefined and known and constant over time. Thus, "activity i must be processed for p_i time units. During this time period a constant amount of r_{ik} units of resource k is occupied" [2]. Similar to the activities which require a finite resource capacity of one or multiple resources, the resources have a finite available capacity per time period which is also predefined and known. Moreover, the resources are considered renewable in the standard RCPSPs, such that their capacity is renewed with each time period. Resource constraints arise due to the fact that the demand for resources arising out of the activities exceeds the per-period availability of the resources which results in having to schedule activities to a later starting time. "The objective of the RCPSP is to schedule the activities such that precedence and resource constraints are obeyed and the makespan of the project is minimized" [27].

Figure 2.2 shows a feasible and optimal schedule for an RCPSP consisting of 4 activities, two resources and one precedence relation where activity 3 cannot start until activity 2 has finished [2]. The time periods are shown on the x-axis and the resource capacity on the y-axis. The dotted lines indicate the maximum available capacity of the two resources. As can be seen from the Figure, although the left-hand side schedule is feasible, its makespan is larger than that of the optimal schedule shown on the right-hand side.

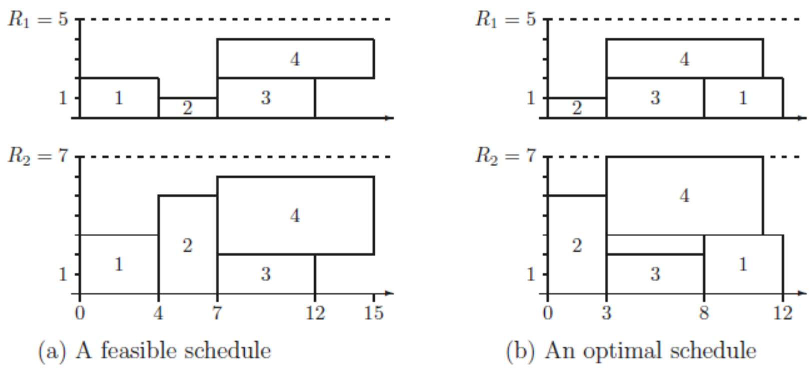

Figure 2.2: Example of a feasible (left) and optimal (right) schedule of an RCPSP consisting of four activities allocated over two resources [2]

The standard RCPSPs have so-called "cumulative resources", which implies that multiple activities may be allocated to a resource at any time period, provided that the resource capacity is not violated [30][2]. In JSSPs the machines are only ever allowed to be allocated one single operation (task) at any period in time [2]. Provided that "RCPSP has become a standard problem for project scheduling in literature", they have a wide variety of applications to different problems [26]. Consequently there is an equal diversity in RCPSP formulations. Variations include different types of precedence relations between activities, the inclusion of non-renewable resources, time-lags, flexible processing times and resource requirements through different modes, time-varying resource request or resource availabilities, etc. It is impossible and beyond the scope of this research to provide a complete overview of the various different RCPSPs. Rather, the literature study was used to provide the author with a basic understanding of task scheduling optimization models, including RCPSPs. The interested reader is referred to the works of Hartmann and Briskorn [26], Brucker and Knust [2] and Artigues and Neron [30] which provide a more complete review of RCPSPs and their variations.

As the formulation of RCPSPs is primarily driven by its applications, these are equally diverse. Some examples of applications include the allocation of ground handling tasks to airport personnel to minimize the aircraft turnaround time[2]; optimal scheduling of testing on experimental vehicles in the automotive industry to minimize the number of required experimental vehicles [31], the allocation of internal and external employees to IT-projects such that the labor costs are minimized [32], etc. As indicated by Hartmann, similar applications of RCPSPs may be found for the pharmaceutical industry, telecommunications, aircraft maintenance, port operations, steel manufacturing, etc [26]. One such application is addressed in the work of Cavalcante et al., who used RCPSP to formulate a MILP model to optimize the allocation of incoming IT problems to IT-helpdesk employees [33]. Different resource capacities of employees represented the level of experience of the employees with multitasking. A tardiness penalty is incurred whenever due dates are exceeded and another penalty is imposed whenever multiple tasks are allocated to one employee at the same time. The latter penalty represents the extra effort required by an employee who has to multitask. The model's objective is to find a schedule of tasks which minimizes the total penalties with limited resource capacity. Cavalcante conclude that the model is a "promising tool to help performance analyses of dispatching activities inside service delivery pools" [33].

The great diversity of RCPSP formulations also yields an equally diverse range of mathematical model formulation and optimization algorithms. Contrary to the standard RCPSPs the model formulation of Cavalcante et al. does not include precedence relations between the tasks, as these simply do not occur in the problem being addressed. However, similar to the FJSSPs discussed in the previous Section, the model of Cavalcante et al. allows for a task to be allocated to either of a variety of resources rather than specifying the required resources in advance. As such, the set of binary decision variables ordinarily used for the allocation of the starting time or completion time to a time period also concerns the allocation to one of the different resource types [33]. Nadjafi et al. formulated a mode-identity RCPSP, in which tasks are clustered to specific subsets, where each subset of tasks is to be executed by the same resources [34]. The problem represents a typical assignment-based scheduling problem, where the optimal division of work over various resources needs to be found which will minimize the makespan of the project [34]. Thus, in the model formulation of Nadjafi et al. the set of binary decision variables which allocates a task to be completed in a certain time period, also assigns that task to a specific mode.

In addition to the RCPSP applications, there is substantial literature dedicated to the improvement of solution algorithms for RCPSPs. In their work Nadjafi et al. compare the performance of a genetic algorithm (GA) expanded with a knowledge-based local search algorithm to an exact branch-and-bound algorithm [34]. The local-search algorithm has been added to improve the likelihood that the genetic algorithm converges towards the global optimum rather than a local optimum [34]. The GA has a significantly lower average computational time of just 0.263 seconds, compared with the 339.8 seconds of the exact method [34].

Much of literature contains a detailed sensitivity analysis of the models and solvers. Commonly the effect of the number of activities, the number of precedence relations between activities, the available resource capacity and the tightness of the due dates on the computational time or problem instances solved is investigated. For example, in the work of Wang et al. the performance of two bi-objective MILP formulations which are solved optimally using CPLEX are compared for a RCPSP of at most 20 jobs and 4 resources [35]. The models are optimized for minimum makespan and tardiness. One of the MILPs has a conventional set of binary decision variables, inspired after the formulation given by Pritsker, to allocate the jobs to a starting time [35][29]. In the other MILP the decision variables of a job equal one for all time periods for which the the job has been completed, as inspired by Klein [36]. The experiments conducted by Wang et al. conclude that the formulation of the second MILP "requires less computational time" [35].
The comparison of the RCPSP solved approximately using a GA and exactly with a branch-and-bound algorithm in the work of Nadjafi et al. has also been based on the results of extensive sensitivity analysis based on parameters similar to the once which have been mentioned [34].

2.2.3. Resource Allocation Problems

The resource allocation problem (RAP) comprises the third prominent category in task scheduling optimization problems. As with the RCPSPs, the problems focus on the minimization of the makespan of a set of activities (requests) and a set of resources with a finite available capacity per time period. However, "compared to the [RCPSP], where emphasis is given on sequencing tasks for a certain number of resources, considering or not their capacity constraints, resource allocation problems reverse somehow the issue, focusing on the distribution of resources by tasks to be performed" [14]. Thus, despite the familiar setup, the RAPs formulate the problem from the perspective of the resources rather than the activities. As such, RAPs include precedence relations between resources for a specific activity which requires multiple resources in a specific order. Moreover, each activity has a predefined processing time by a certain resource type in addition to the required capacity for each resource type. Thus, RAPs are concerned with which resource to allocate to what activities such that the overall processing time of all activities is minimized, subject to the precedence constraints between resources of activities and subject to the capacity constraints of the resources. Please note that there may also be precedence constraints between activities or due date constraints which have to be taken into consideration. As for RCPSPs and JSSPs, literature provides numerous examples of RAP applications to real world scheduling problems. Two examples are briefly discussed.

Bertsimas et al. consider the aircraft maintenance problem "as an important dynamic resource allocation problem" [37]. They consider a hangar in which multiple aircraft require maintenance simultaneously. Each aircraft has a predefined set of preventive maintenance tasks which need to be executed. A multi-skilled workforce and different types of equipment make up the available resources, each which a finite available capacity per time period. The order in which maintenance tasks have to be executed enforces precedence relations between resources for each task. Moreover, workers with different levels of experience and equipment with different status are incorporated in the model. These "alternative resources" affect the processing times of the activities [37]. Bertsimas et al. define a set of binary decision variables that equal one when a resource has been allocated to an aircraft by a certain time period [37]. The objective of the model is to minimize the makespan by left-shifting the starting times of all resources for all aircraft as much as possible. The model of Bertsimas et al. is solved exactly using Gurobi as a solver and a branch-and-bound method.
As has been explained previously, it is interesting to emphasize again how the RAP differs in this particular example from an RCPSP or JSSP. Whereas in a JSSP and RCPSP the order in which the activities are to be executed is among the primary objectives of the task scheduling, in an RAP this order has yet been determined through a prioritization of activities and instead an optimal allocation (and order) of resources to these tasks has to be determined.

A similar example of an RAP applied to equipment maintenance of a steel factory is provided in the work of Safaei et al [38]. The problem considers both preventive maintenance tasks as well as repairs and potential failure tasks. Moreover, similar to Bertsimas et al. the RAP includes a multi-skilled workforce, where each task requires one or more trades for completion. However, contrary to the formulation of the standard RAP, Safaei et al. do not specify the resource capacity of each trade in advance. Rather, they have formulated the problem to bi-objective mixed integer linear programming model in which the objectives are the minimization of the total throughput time of all maintenance tasks and the minimization of the size of the workforce [38]. Moreover, Safaei et al. included weights for task based on their effect on the equipment availability. The model is solved exactly and provides a Pareto front which shows the trade-off between workforce size and maintenance induced downtime of equipment. Thus, the work of Safaei et al. provides the manufacturing plant with valuable insight into the costs of labor in comparison with equipment unavailability costs.

The model of Safaei et al. uses an alternative approach for the model formulation. Safaei et al. use a so-called 'network flow formulation' in which the workforce allocated to one task may also be allocated for subsequent tasks [38]. This approach eliminates the time dependency of the decision variables as is the case in the work of Bertsimas et al. The alternative formulation have a profound effect on the model performance. Whereas the model of Bertsimas shows an increase in the computational time for an increase in the number of resources, the model of Safaei et al. actually shows a increase in computational time for a reduction in the number of resources. In the model of Bertsimas et al. the number of decision variables scales with the number of resources. Moreover, the time dependency implies that the decision variables have to be evaluated for each time period. Although the network flow formulation of Safaei et al. is not dependent on time, it also necessitates the formulation of alternative decision variables. One such is a binary decision variable that controls the allocation of resources from one task to another after the former has been completed. With fewer resources available, more tasks are scheduled in sequence, which also increases the number of these binary decision variables to be evaluated an subsequently the computational time. This shows again how the model formulation affects the model performance.

This Section has provided some basic knowledge on the formulation of RAPs and how it differs from the standard RCPSPs and subsequently the model formulations. This has been further explained using two examples from literature which have also shown the applicability of RAPs to describe real-world task scheduling problems. Finally, it has been shown again how the different model formulations affect the model performance. As for JSSPs and RCPSPs, a substantial amount of literature related to RAPs is dedicated to comparisons of model formulations, performance of different solvers both exact and approximate. Provided the interest in model applications and the literature discussed in the previous Sections, these have been purposely left out of this Section.

2.3. Discussion

The previous Section has discussed the three most prominent categories in tasks scheduling optimization models, JSSPs, RCPSPs and RAPs. As becomes apparent when comparing the three categories, the three problem types have much resemblance, but are also distinctly different. The type of scheduling problem is dependent on the problem formulation. The objective, constraints and decision variables, as well as the known and unknown input parameters affect the type of scheduling problem.

Additionally, task scheduling literature may be divided into two main categories. The first involves the formulations of new optimization methods and is focused on improving the current state-of-the-art in model formulations, solution algorithms and the performance of such models and solvers with respect to established methods. The second category in literature focuses on applications of task scheduling model to real-world problems. For such literature the emphasis is placed on bridging the gap between theory and practice and formulating a model which can provide decision support in task and resource allocation. The present model and study fall into this last category of literature. Section 2.3.1 will discuss how the present study is expected to add to the current state-of-the-art in applications of task scheduling models. Following this, Section 2.3.2 will discuss how the model formulation differs from the standard task scheduling problems introduced in Section 2.2. Lastly, a brief discussion will be provided on the solver choice.

2.3.1. Literature Gap

Section 2.2 has illustrated some examples of scheduling models applied to (aircraft) maintenance, including the work of Ahire et al., Bertsimas et al. and Safaei et al. [25][37][38]. There is vast amount of literature available on task scheduling optimization models applied to aircraft maintenance. Based on the performed literature study, the present study is expected to contribute on various aspects to this field.

Contrary to most literature on aircraft maintenance in general and more specifically literature on task scheduling models related to aircraft maintenance, the current study will include the complete workload of base maintenance tasks. As noted in the introduction, the project objective is *"To provide insight into the required maintenance production capacity for the complete workload of an airline's wide-body fleet over a tactical planning horizon by developing a resource and task scheduling optimization model"*. Section 1.4 explained that the complete workload of the wide-body fleet of the airline in question consists of routine (preventive) maintenance tasks, non routine (corrective) maintenance tasks and modifications tasks. A more detailed explanation of these tasks is provided in Chapter 3.
As mentioned in Section 2.2, the aircraft maintenance problem of Bertsimas et al. was limited to preventive maintenance tasks only [37]. The study of Safaei et al. included preventive maintenance tasks as well as corrective maintenance tasks in the form of failure repair and potential failure tasks [38]. However, to the knowledge of the author, literature on aircraft maintenance problems does not include modification tasks. With current workloads on the airline's wide-body fleet in the order of several hundreds of hours and with a continuous influx of new modifications, modifications represent a substantial part of the overall workload on the aircraft and therefore must be considered. Moreover, existing literature on aircraft maintenance would appear to be restrict itself to the set of preventive maintenance tasks which are scheduled to routine checks. However, there is a set of preventive maintenance tasks, referred to as out-of-phase tasks in this work, which are not fixed to any specific maintenance check and as such require scheduling. The present work will include both modification and out-of-phase tasks in addition to the other preventive maintenance tasks and corrective maintenance tasks. As such, the present work provides a much more complete representation of the actual base maintenance workload encountered on modern day aircraft. The addition of modifications and out-of-phase tasks is considered particularly valuable from a scheduling perspective as both task types are not "fixed" to any specific maintenance check (resource) and add scheduling complexity through additional constraints which have to be considered in producing an optimal task schedule.

As noted in the project objective, the present study will consider a tactical planning horizon which encompasses multiple years. As such, the model will include all the routine letter check types (A-, B- C-checks) of the airline's wide-body fleet. This will provide the MRO with the desired insight into complete workload as distributed over all check types. Thus, the present study concerns the optimal allocation of the complete workload of maintenance tasks to all maintenance checks within the specified planning horizon. This is different from most aircraft maintenance literature which is often limited to an operational level and as such only considers the optimization of a single maintenance check. Bertsimas et al. consider a multi-aircraft resource allocation problem (hence also multiple maintenance checks) [37]. However, the problem formulation of Bertsimas et al. considers how to best allocate the finite resources to the various aircraft, provided that they all undergo maintenance simultaneously. Thus, it too is restricted to an operational level only, rather than the tactical planning horizon including multiple maintenance checks addressed in the present study. Moreover, the inclusion of the different types of maintenance checks, modeled conform their diversity in practice, is considered another contribution to literature.

In addition to the allocation of tasks to resources, the present model also includes the allocation of resources to aircraft. Similar to the formulation of Bertsimas et al. the aircraft represent the owners of the various maintenance tasks and checks which "belong" to the respective aircraft registration [37]. As will be explained in greater detail in the Chapters 4 and 5, the allocation of extra maintenance checks has been introduced to the model to provide additional resource capacity when needed. Thus, the present model deviates from the conventional task scheduling models which have been discussed in this Chapter as these concern either task scheduling to resources or resource allocation to tasks. This too is considered a contribution to literature and exemplifies how real-world requirements affect the model formulation.

2.3.2. Problem Definition

As is apparent from the previous comments, the model in the present study will differ quite substantially from the conventional task scheduling allocation models. The primary difference is the objective of the current study compared with the conventional task scheduling models. Whereas for most JSSPs, RCPSPs and RAPs the problem's objective is to minimize the overall makespan or throughput time of all tasks subject to resource capacity and precedence constraints, the present study's objective is to determine the relation between the required and available resource capacity. The emphasis on resource capacity efficiency rather than makespan has a significant impact on the model formulation.

The conventional task scheduling models which minimize for makespan often include a soft due date constraint which allows for due dates to be exceeded. In these conventional task scheduling models, exceeding a due date simply implies that the task is allocated to a later time period in which the required resources have a replenished capacity to accommodate the task. Provided that this increases the makespan of the tasks, soft due date constraints are often accompanied with a tardiness penalty in the objective function. Examples of this have been shown for the works of Cavalcante, Yang, Wang, Ranjbar and Thornblad in the previous Section [33][24][35][39][23]. Provided that airworthiness regulations and manufacturers strictly specify the interval of tasks, extension of due dates is not considered in the present model.
Moreover, provided that the objective of the study is not to determine the minimal makespan of all tasks, the present study also does not consider precedence relations between tasks.

Similarly to the determination of the minimum makespan through soft due date constraints accompanied by tardiness penalties, the present study formulates a soft resource capacity constraint with accompanying penalties. In the work of Yang et al., earliness and tardiness penalties are imposed when the completion time of a task either precedes or exceeds the due date. Similarly, penalties are imposed for when the allocated resource capacity either exceeds or runs short of the available resource capacity. This will be discussed in more detail in Chapter 4. The soft resource constraint differentiates the present study further from the conventional task scheduling models. As discussed, Safaei et al. formulated a bi-objective model which considers both the minimization of the makespan and of the workforce size. However, most task scheduling literature imposes a predefined and known resource capacity. The formulation of a soft resource constraint more closely represents best practices in scheduling at the MRO. Furthermore, provided that the present study has a fixed planning horizon with a finite number of resources and the unknown relation between required capacity and available capacity, the present formulation better suits the objective of the present study.

Thus, the problem formulation is inspired by the conventional task scheduling problems. It most closely represents an RCPSP as tasks with a finite required capacity are allocated to resources with a 'finite' available resource capacity. However, the problem formulation differs as it includes soft resource capacity constraints

with accompanying penalties; hard due date constraints; no precedence relations between tasks; a fixed planning horizon with a finite number of resources; and it optimizes for resource capacity efficiency rather than makespan minimization. This Section has tried to explain how the present study fits within the established task scheduling literature. The mathematical formulation of the model is provided in Chapter 4.
One could argue that the present optimization study precedes the conventional task scheduling problems. This optimization will probide each maintenance check with a set of tasks optimized for resource capacity. Subsequently, a conventional task scheduling model could be used to optimize the makespan of this set of tasks for each maintenance check.

2.3.3. Solver Choice

A lot of literature is directed towards the performance heuristic solvers in comparison with exact solution methods. In general these comparisons show the superiority of approximate solution techniques in terms of computational time and tractability for large and complex problems, compared with exact methods. Moreover, most literature shows that with increasing problem sizes and complexity, approximate solution techniques do yield some loss in solution accuracy with respect to the exact solution which may or may not be acceptable. Examples of such comparisons on solver performance have been provided from the works of Nadjafi et al., Yang et al. and Fattahi et al. [34][24][21].

Despite the general superior performance of approximate solvers, the present study uses CPLEX as solver. This choice was based on a number of considerations. Firstly, as CPLEX licences are both available to the MRO and to Delft University of Technology, the solver is accessible for both the research study and potential implementation of the model later on. Secondly, CPLEX is a ready-to-go platform. It only requires one to provide the model inputs in a specific format for it to run the optimization. This makes CPLEX much easier to use and implement. This latter was particularly important, as the author had no prior knowledge of operational research models and solvers prior to this study. Hence, the relative ease of implementing CPLEX would save the extra time of becoming acquainted with heuristic solvers. Moreover, the ease of use of CPLEX is relevant for the MRO users in case the model is implemented for use. Finally, the model performance (in terms of computational time) was considered secondary to the insight provided by the model in the present study. As already mentioned, this study concerns the application of a task scheduling model to a real-world problem, rather than the improvement of a modelling or solver technique.

3

Background on the MRO

Following the literature study which has been discussed in the previous Chapter, this Chapter provides some more background information specific to the MRO. Firstly, it briefly describes the main differences between line maintenance and base maintenance, as the present research is limited to base maintenance only. Following this, some brief information related to the modification workload of the airline's wide-body fleet will be provided. Section 3.3 will provide an explanation of the organizational structure of the MRO's planning department: 'Planning Scheduling and Fleet Control'. Section 3.4 will provide an elaborate explanation of the current planning policies and processes at the MRO. Finally, the Chapter concludes with Section 3.5 which discusses the definitions and origins of the different types of maintenance tasks.

3.1. Line Maintenance and Base Maintenance

The MRO is a major MRO Service provider that provides both line maintenance and base maintenance services to the airline. The main difference between the two is that the former refers to maintenance which is performed to the aircraft in between flights at either the gate or an apron, while the latter refers to aircraft maintenance performed in a hangar. The distinction in location also implies a distinction in the nature of the maintenance work which is performed by either. Base maintenance generally involves heavier maintenance items, in terms of required labor hours and downtime of the aircraft, which are performed routinely at predefined intervals. Thus, at regular intervals the aircraft is taken out of operation and brought into the hangars for maintenance. Most of these routine maintenance tasks are structured into so-called letter checks, A-, B- and C-checks, as will be further explained in Section 3.5.2.

All maintenance tasks which are 'smaller' than the tasks scheduled in the routine A-checks is ordinarily performed by line maintenance. Here smaller refers to the labor hours and groundtime required, but also to the task interval. Thus, the line maintenance workload primarily consists of so-called 'time limited work' (TLW), which are routine maintenance tasks with a smaller interval than the tasks scheduled to the A-checks. In addition, much of the line maintenance work involves corrective maintenance tasks including so-called MSOs and DDs. MSOs refer to faults or failures on the aircraft which need to be addressed. Thus, MSOs are very much ad hoc. 'DD' refers to deferred defect. Deferred defects refer to faults on the aircraft which have not been resolved at the time that they were discovered. Rather, based on the nature of the fault, the corrective maintenance work required to resolve the fault is assigned a due date prior to which the work has to be performed. there are numerous reasons for delaying corrective maintenance, the most common reasons are avoiding operational disruptions and lack of available material or equipment at the time that the fault is discovered. As will be explained later, oftentimes some of the line maintenance workload is allocated to the base maintenance checks. The present study is restricted to base maintenance only.

3.2. The Modification Workload of the Airline's Wide-Body Fleet

This Section provides a brief introduction of the airline's wide-body fleet, as the subject of interest in the present study. The airline's wide-body fleet consists of thirteen aircraft. Although the aircraft is indeed a great addition to the airline's fleet and its operational efficiency, its design also presents unique and unexpected challenges. One of these challenges is the unexpectedly high workload of modifications to the aircraft.

Modifications include adjustments and improvements to the aircraft during its operational life. Following the launch of an aircraft, the need for design changes will arise during the aircraft's operational life. Such design changes are either issued as airworthiness directives (ADs) or service bulletins (SBs) for the purpose of improving the aircraft's operational performance, continued airworthiness, enhanced reliability of components, etc. ADs are issued by the regulatory authorities and are mandatory items for continued airworthiness [40]. Service bulletins are issued by the OAMs or OEMs as product improvements and are optional to the operator [40]. Furthermore, the operator may also request the MRO service provider to implement desired modifications, such as a cabin upgrade.

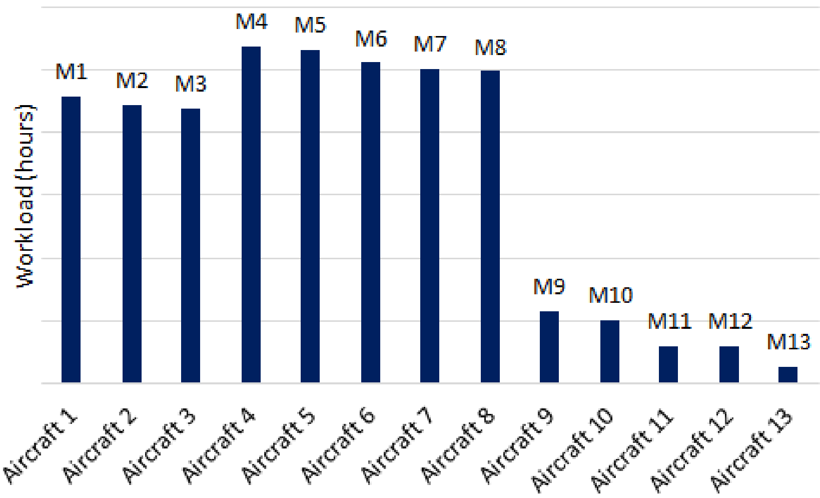

Figure 3.1: Overview of the modification workload on the airline's wide-body fleet

Figure 3.1 shows the modification workload for the airline's wide-body fleet per aircraft dated to early 2019. As can be seen from this Figure, the workload on the aircraft varies from M4 hours to M13 hours. The workloads on the 'oldest' eight aircraft registrations are largest provided that these aircraft still require some of the modifications which have yet been included in the designs of the later aircraft registrations by the OAM. The workloads on the first 3 aircraft registrations are slightly less compared to the following 5 aircraft as a result of the recent C-checks for these first three aircraft. However, despite these recent C-checks the workloads on these three aircraft are still substantial. Provided that a routine A-check maintenance slot caters for ANORM labor hours, in the most ideal world these workloads would be equivalent to at least five extra A-checks per aircraft. Furthermore, provided that each aircraft undergoes approximately four A-checks per year, the modification workload is indeed substantial.
Moreover, it should be noted that these workloads are based on the known estimated labor hours. For approximately PER% of the modification used in this analysis, the labor hours yet have to be estimated and reported and as such are not included in the values in Figure 3.1. Moreover, there is a continuous inflow of new modifications as desired by the operator or suggested by the OAM.

With the known workloads of modifications as high as described by Figure 3.1 and a finite available labor capacity, the MRO is faced with the challenge of allocating the modifications as efficiently as possible. Hence, the aim of the present study is to provide insight into the most efficient allocation of modification tasks and all other base maintenance tasks based on available resource capacity, over a tactical planning horizon.

3.3. Organizational Structure of PSFC

This Section provides an introduction of the organization structure of the "Planning Scheduling and Fleet Control" (PSFC) department of the MRO as the department which is responsible for all maintenance and flight planning of the airline's fleet.

3.3.1. Planning Scheduling and Fleet Control

In addition to the three primary maintenance units, the MRO also consists of various support units, including Central Engineering and PSFC. As has been explained, PSFC is responsible for planning and scheduling of all aircraft maintenance for the airline's fleet. According to the preparations and plans made by PSFC, the Airframe units (hangar or line maintenance support units) will execute the work. Thus, PSFC works in direct relation with both the operator and the Airframe units to determine when each aircraft will be receiving its required maintenance. In doing so, PSFC has to guarantee the airworthiness of the fleet whilst also optimizing fleet availability for the operator. The PSFC department may be divided into four sub-departments, as shown by the schematic in Figure 3.2.

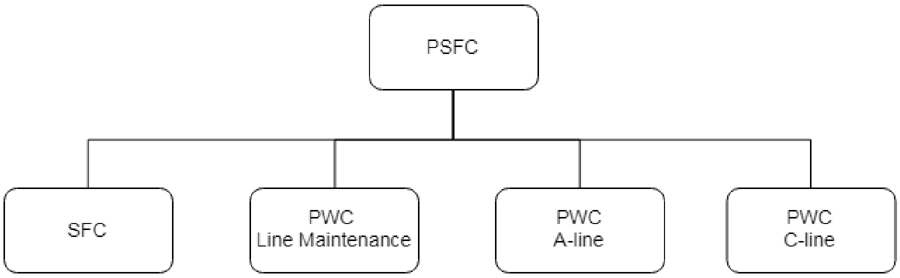

Figure 3.2: Organizational structure of Planning-, Scheduling and Fleet Control Department of the MRO

The Scheduling and Fleet Control department is concerned with the strategic and tactical planning horizons of aircraft maintenance. Based on the flight schedules provided by the operator, groundtime planners and maintenance planners make a schedule of the various maintenance checks for each aircraft based on its flight schedule. Historical data is used to estimate the utilization rate of the aircraft. With the utilization rate they amend the flight schedule such that it includes the maintenance checks of each aircraft conform the composition rules. These rules define the intervals of the routine maintenance checks and ensure that all maintenance tasks included in each check are performed within the intervals as specified by the OAMs and as endorsed by regulatory authorities. In allocating the aircraft registrations to the various available maintenance slots the groundtime and maintenance planners also ensure that the schedule does not exceed the capabilities of the hangars in which the maintenance is performed. Given that this schedule is planned over a tactical planning horizon, it is constantly susceptible to changes. As aforementioned, throughout these changes the groundtime planners and maintenance planners have to optimize for fleet availability whilst complying with the maintenance task intervals and the capacity and material availability constraints.

The three other departments are responsible for the Planning and Workload Control (PWC) of the line maintenance, A-checks and C-checks respectively. Once an aircraft registration has been committed to a maintenance slot by the groundtime and maintenance planners of the SFC, the responsible operational maintenance planners allocate the tasks to the maintenance check in question. Thus, the A-checks are prepared by the operational maintenance planners from the PWC department. The same is done by the operational maintenance planners of the PWC C-line department for the C-checks. A more detailed explanation of the preparation process of the A- and C-checks is provided in the next Section.

The line maintenance slots are not allocated by the groundtime planners and maintenance planners, but rather by the PWC Line Maintenance department. The groundtime and maintenance planners usually schedule the letter checks (A-, B, C-checks) such that there is some scheduling flexibility still available in the remaining interval of the maintenance tasks. However, the flexibility is limited to the due dates of the tasks involved. The line maintenance slots are inherently much more flexible and subject to the allocation of aircraft registrations to flights. Moreover, the line maintenance workload is also much more dynamic, as has been explained previously. Provided the dynamics involved, the workpackages and slots for line maintenance cannot be planned on a tactical horizon, but rather are allocated by the PWC Line Maintenance department which oversees the day-to-day operations. As such, this division may assign these workpackages to slots in between flights or to slots in the hangars. These hangars slots may be specifically for smaller workpackages such as line maintenance tasks, in which case these slots are referred to as H- or TO-slots. It is also customary within the MRO for line maintenance workpackages to be assigned to aircraft coming in for routine base maintenance (A-checks or C-checks). Since the aircraft is already coming in for maintenance, a small package of line maintenance work is allocated to a letter check.

3.4. The MRO's Planning Policy and Systems

This Section provides some background on the policies and systems which are used by maintenance planners at the MRO for planning and scheduling.

Like all MRO service providers, the MRO has a primary Maintenance Information System (MIS) which is used to support all business. For each aircraft registration of the airline's fleet it contains an overview of all routine maintenance tasks contained in the maintenance program, as well as all corrective maintenance tasks as issued by the production teams. It provides detailed information for all maintenance tasks including their upcoming due date; remaining interval days; job instruction cards; labor hours and skills required; access panel cards; zones on the aircraft; required material and equipment; references to instruction manuals, etc. As such, it is used by operational maintenance planners to prepare the workpackages for aircraft registrations with upcoming maintenance checks. Moreover, it is used by the production teams in the hangars to execute the workpackages. In the present study the MRO's MIS is used as the primary source of all task-related information required by the model.

3.4.1. Planning Policy: Tmin Structure

The routine letter checks are prepared according to the so-called Tmin-structure policy. A schematic visualization of this policy for the A-checks is shown in Figure 3.3 [3].
As can be seen from Figure 3.3, the Tmin policy identifies several checkpoints.

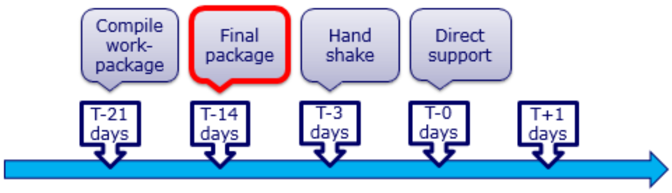

Figure 3.3: Schematic overview of the various steps of the Tmin structure policy for the preparation of the A-checks [3]

- The first checkpoint in the Tmin policy of the A-checks is at 21 days prior to the actual check date itself and marks the latest starting time for the operational maintenance planners to compile the workpackage of the check. Based on the due dates of the tasks and other criteria, the operational maintenance planners select from the MIS which tasks will be included in the upcoming A-check.
- At 14 days prior to the check the workpackage must be finalized. As such, no more work is added to the workpackage following this checkpoint. Based on this final workpackage, the material and equipment facilities of the MRO are informed and start with the preparation of the required items.
- The handshake at three days prior to the check involves a confirmation by all parties involved that all necessary preparations have been made. With this 'handshake' the responsibility is also transferred from PSFC to production.
- During the check itself a direct support team is available to the production teams to take care of any problems which arise during a check, including the need for addition materials, equipment, etc.
- Finally, the preparations and execution of the check are evaluated following its completion, to continuously improve these processes.

Similar to the A-check Tmin policy which has been shown and explained, there is a Tmin policy for the C-checks. It largely consists of the same checkpoints, however these are spread out over a much larger planning horizon, starting at 26 weeks prior to the check. The policies ensure that all 4Ms, according to the lean philosophy which is implemented at the MRO, are in place prior to the start of each maintenance check. These

4Ms represent the aircraft and required equipment (machines), people required to meet the demand of the scheduled workload (manpower), the approach to the scheduled work, e.g critical path and division of work to shifts (methods) and materials. Thus, compliance to this Tmin policies ensures that all necessary preparation are made prior to the start of the check.

3.4.2. Planning Policy: EO Order Process

As has been discussed, the modifications are an important part of the workload of the airline's wide-body fleet. To ensure an efficient implementation of new modifications on the fleet, the MRO has developed an order process to track and trace the implementation of new modifications from the initial proposal up to the first execution on an aircraft. This process is known as the EO Order Process.

The singular nature of the modifications implies that they generally require much more elaborate preparations compared with the routine maintenance tasks. For example, modifications may require specialized materials which need to be ordered; negotiations with vendors on warranty and price may be required; labor hours and turnaround times need to be assessed; instructions have to made or reviewed by engineers; special conditions may restrict the possible maintenance checks in which modifications may be executed; etc. Provided their inherent uniqueness and the large volume of modifications, the EO order process brings all relevant departments of the MRO together to enable the efficient preparation of the execution of all modifications.

The EO order process consists of four primary steps:

- **Evaluation Phase:** Proposed modifications are evaluated by the operator and the MRO based on their desirability, priority and the available budget. A so-called ORP, operator requirement proposal, is created for each potential modification.
- **Plan Phase:** Inquiries are made by all relevant MRO departments with regards to the potential execution of the modification. These include estimations of the required labor hours and turnaround times; lead times and delivery dates of materials or components; the need for specialized equipment or special conditions during task execution; the preferred check type to execute the task in, etc. Based on these preparations and the results from the evaluation phase, a final go- or no-go decision is made by the operator.
- **Plan Execution Phase:** Provided the approval by the operator, all preparations and plans of the previous phase are put into action: required materials or components are ordered and received; Central Engineering creates the task in the MIS; the first execution date is set, and so on.
- **Task Plan and Execution Phase:** The modification is selected by the operational maintenance planner from the MIS and allocated to the workpackage of an upcoming check, followed by a successful first execution of the task.

All relevant information regarding the ORPs throughout the EO order process is recorded and updated in a Data Exchange Platform by the various departments. Through continuous monitoring the progress of each ORP towards a successful first execution is ensured. This Data Exchange Platform is also used as the source of information related to the modification workload of the airline's wide-body fleet for the model in the present study.

It is important to note that the Data Exchange Platform is not a planning tool to replace the MIS. Rather, it is used to track the implementation progress of the modifications. Thus, as the modification is ready to be executed, the Central Engineering department will create an 'engineering order' (EO) in the MIS, which contains all relevant planning and execution data of the modification in question. This subsequently allows the operational maintenance planners to select EOs in the MIS when compiling the workpackages for an upcoming check. Please note that EOs refer to scheduled maintenance tasks which are ordinarily only executed once and which are issued by Central Engineering, e.g. modifications.

3.5. Definitions and Origins of Maintenance Tasks and Checks

This Section provides a general overview of the different types of maintenance tasks and checks as will be referred to in this report and in accordance with the definitions at the MRO and literature.

3.5.1. Maintenance Task Types

The overview of all maintenance tasks for the airline's wide-body fleet is contained in the MRO's Aircraft Maintenance Program (AMP). Such an Operator's Maintenance Program (OMP) is based on the Maintenance Planning Document (MPD) published by the OAM and subject to approval by regulatory authorities. The MPD is based on the Maintenance Review Board Report (MRBR); the Certification Maintenance Requirements (CMR); the Airworthiness Limitations (AWL); original aircraft manufacturer (OAM) and equipment manufacturer (OEM) recommendations or requirements; and ETOPS Configuration, Maintenance and Procedures Documentation (CMP) [41][42][43]. In addition to the MPD, the AMP also consists of non-MPD related tasks, including modifications originating from airworthiness directives (ADs) and service bulletins (SBs) [41][43]Å schematic overview of the origins of the OMP tasks is shown in Figure 3.4.

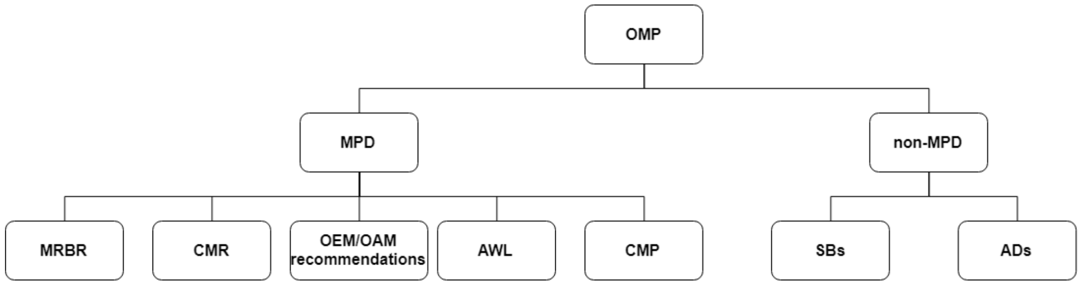

Figure 3.4: Schematic overview of the origins of tasks contained in an Operator's Maintenance Program

MPD Items

- **MRBR Items:** The Maintenance Review Board is responsible for the development of the "initial minimum maintenance requirements" [41]. The board consists of representatives from the respective OAM and OEMs, regulatory authorities, operators and other industry specialists [13]. The development of these items follows the structures and methods as outlined in ATA's Maintenance Steering Group (MSG) [44]. The latest aircraft are developed using the MSG-3 revision 2 analysis.

- **CMR Items:** "The CMR is a required scheduled maintenance task established during the design certification of the airplane systems as an operating limitation of the type certificate or supplemental type certificate" [45]. "The CMRs are intended to establish required tasks to detect safety-significant latent failures [and] impending wear-out of items whose failure would result in a hazardous or catastrophic failure condition" [45]. Unlike the maintenance tasks produced by the MSG-3 analysis, "that involve both preventive maintenance tasks, which are performed before failure occurs (and are intended to prevent failures), as well as failure-finding tasks, CMR's are failure-finding tasks only, and solely exist to limit the exposure to otherwise hidden failures" [46]. The tasks and intervals prescribed in the CMRs are mandatory [46].

- **AWL Items:** "AWLs are mandatory inspection items (with specified instructions and intervals) applicable to items which "the certification process has defined as critical from a fatigue or damage tolerance assessment" [46]. Thus, AWL items are "supplemental structural inspections", in addition to the MRBR's initial maintenance program [41].

- **CMP Items:** "The CMP standard specifies any additional configurations, maintenance or operational requirement that is uniquely applicable to ETOPS (Extended-range Twin-engine Operational Performance Standards)" [47]. Thus, the CMP details specific maintenance requirements and procedures for the aircraft-engine combination that must be complied for the aircraft to perform ETOPS operations [47]. "The requirements in the CMP are established by the FAA at the time of initial ETOPS type design approval of the airplane-engine combination" [47].

- **OEM/OAM Recommended Items:** In addition to the mandatory minimal requirement items, the MRBR also contains "non-safety" related items which are recommended by the OAM or OEMs, for example for economic reasons [41][46].

non-MPD Items

- **SBs and ADs:** In addition to the MPD-related items, modifications are also included as part of the AMP. As has been explained, a service bulletin (SB) is issued by OEMs or the OAM as an optional, but recommended improvement of the product. Airworthiness directives (ADs) are issued by the regulatory authorities (such as EASA and FAA) and are mandatory [40][48].

All MPD-related items are translated into so-called Maintenance Requirement Items, MRIs, by the MRO's Central Engineering department. All non-MPD related items are included in the AMP as Engineering Orders (EOs), as has been explained for the modifications. According to the MRO's documentation related to the AMP for the airline's wide-body fleet, "All mandatory requirements such as, but not limited to CMR, CMP, ALI, AWL, SCI, AD, MD, EASA Subpart M, etc. are implemented accordingly. All non MPD tasks and changed MPD tasks (if any) will be motivated and presented for approval to the operator and their authority (if required)".[8].

Each MRI in the MRO's AMP has the following properties:

- Task number: a unique task number based on ATA and sub-ATA chapters
- Description and references to AMM: a general description of the task including the type of work involved (general visual inspection, lubrication, etc) and ETOPS relevancy if applicable, as well as references to the Aircraft Maintenance Manual (AMM)
- Required labor hours and skills: a specification of the total scheduled labor hours, as well as the skill and number of people required to perform the task
- Interval: task interval in flight hours and/or flight cycles and/or calendar time (days/months/years)
- Zones: specification of the zone on the aircraft where the task is to be performed
- Access panels: specification of the access panels which will need to be opened and closed to access the component/system for which the MRI is intended
- Applicability: indication of the aircraft registrations to which the MRI is applicable
- Blocks: a specification of the various blocks in which the MRI reoccurs (see Section 3.5.2)
- Job cards: an overview of all job instruction cards of which the MRI is composed

Each MRI is composed of one or multiple job instruction cards, JICs. The JICs are the subtasks of which an MRI is composed. For example, one of the MRI's refers to the "Lubrication of the right main landing gear doors hinges, actuators, and mechanisms". This MRI consists of three JICs which refer to the deactivation of the nose and landing gear doors, the actual lubrication of the right main landing gear door hinges, actuators and mechanism, and the activation of the nose and landing gear doors [49]. Each job card contains much of the same information as is the case for an MRI, but only applicable to the specific 'subtask' described by the job card. As such, each JIC has a unique number, detailed task description (and references to the AMM), zone and access panel specifications, required material or equipment description and aircraft applicability. However, JICs do not have an interval or 'block' specification as these are specified only for the MRI to which the JIC belongs.

It is important to note that when a reference is made to a task this refers to an MRI and not to a JIC. Any references to JICs will be done so explicitly.

3.5.2. Maintenance Check Types

Based on the initial maintenance program contained in the MPD, it is the responsibility of each operator to develop their own maintenance program [13]. A common approach applied by operators is known as blocking, where the majority of the tasks in the MPD are clustered together based on their properties (required procedures, conditions) and interval into so-called blocks [41]. The strategy involved with blocking aims to minimize the number of times that the aircraft has to come in for maintenance by clustering the execution of maintenance tasks. These blocks are assigned to routine maintenance checks, commonly known as letter checks.

The MRO has used blocking strategy for the planning of most of the AMP of the airline's wide-body fleet. The MRO distinguishes three letter checks for the airline's wide-body fleet, the A-, B- and C-checks. Studies conducted by the Central Engineering department of the MRO resulted in the formulation of these letter checks and their respective intervals. Table 3.1 gives an overview of the different interval specifications of each of the letter checks, as per the MRO's definitions [8].

Table 3.1: Specification of the letter check intervals of the airline's wide-body fleet as per the MRO's definitions [8]

Interval Specification	A-check	B-check	C-check
Flight Hours	1500	-	15,000
Flight Cycles	200	-	2,000
Calender Time	105 days	15 months	3 years
Number of Blocks	24	2	15

Like the MRIs which are allocated to the letter checks, the intervals of the letter checks are defined in flight hours, flight cycles and calender time and are scheduled on a 'whichever comes first' policy. The smaller intervals of the A-checks imply that blocks consisting of tasks which have a higher scheduling frequency are allocated to the A-checks. As indicated in Table 3.1, the MRO has defined 24 unique A-blocks. MRIs with much larger intervals are allocated to the C-checks. The MRO has defined 15 unique C-checks. The B-checks, are also known as the Cabin-checks. Unlike the A- and C-checks, the MRO's B-checks do not include any MPD items, and as such do not have a hard interval. Rather, the Cabin checks were introduced per the request of the operator "to support company standards like; passenger and crew comfort, aesthetic, interior and exterior cleaning and ground handling service items to obtain an optimum in economical fleet performance" [8]. The workload involved has been distributed over two cabin check blocks.

MRIs recur in different blocks depending on their interval. For example, a '1A' MRI recurs in each of the 24 A-blocks. Similarly, a '3A' MRI may appear in the A3, A6, A9 blocks, etc. The same is true for the C-blocks. All MRIs which recur within the 24 A-blocks and 15 C-blocks are clustered as part of their respective blocks. These MRIs are also referred to as **block tasks** in the remainder of this report. The block structure is both sequential and cyclical, such that A1 is followed by A2, and after the completion of A24, the next block is A1 [8].
However, not all MPD MRIs are clustered into the blocks. The routine maintenance tasks which are not clustered into blocks are known as **out-of-phase tasks** or (OOP tasks). These notation will be used throughout the remainder of this report. As the name suggest some tasks are out-of-phase with the block structure as they do not recur within the 24 or 15 A- or C-block cycles. Furthermore, some MPD MRIs are purposely left out of the blocking structure to provide extra scheduling flexibility. Unlike the block MRIs which have some limited extensionability, these MRIs have strict interval limits. Thus, not clustering these MRIs into the blocks provides the scheduling flexibility to always ensure a timely execution of these tasks.

In addition to the OOP MRIs, the non-MPD tasks are also not clustered into blocks and as such not part of any specific letter check. However, unlike the recurring OOP tasks, these Engineering orders are non-recurrent tasks and therefore are generally only scheduled once.

"The advantage of these blocks is that the checks are very predictable content wise. The variations to the check content are the OOP task, Engineering orders/ORP and unscheduled maintenance" [8]. The following Chapter will discuss the formulation of the task scheduling optimization model, which aims to find the optimal allocation of the OOP and EO tasks to the available resource capacity of the letter checks, provided the resource capacity which has already been expended towards the block tasks and the corresponding unscheduled maintenance work.

I

Model Setup

4 Model Formulation

The following three Chapters provide an elaborate explanation of the setup of the task scheduling optimization model developed throughout the present study. This particular Chapter will provide a detailed introduction of the complete model. Chapter 5 elaborates on the various penalty terms of the objective function of the model, including their formulation and underlying reasoning. Chapter 6 discusses in greater detail the data preparation process preceding the actual optimization model and as such provides insight into the assumptions incorporated in some of the model inputs.
This Chapter presents the task scheduling optimization model developed during this study. Section 4.1 will provide a detailed description of the model formulation, including the logic and reasoning of the model's objective as well as the modeling of maintenance tasks and checks. The mathematical model formulation is subsequently presented and described in Section 4.2. The iterative solution method built around the optimization model is explained in greater detail in Section 4.3. Finally, a brief mention of the software and hardware involved with model is made in Section 4.4.

4.1. Model Description

The following Section provides an introduction of the task scheduling optimization model developed during the present study. The modeling of the maintenance tasks and checks, as well as the logic and reasoning of the model's objective will be described respectively.

As noted in the Introduction, the model focuses on the optimization of maintenance task planning from a resource capacity perspective. Thus, the two primary forces at play in the optimization are the required resource capacity of each of the maintenance tasks to be allocated, and the available resource capacity of the maintenance checks. The resource capacities of the maintenance checks and tasks are expressed in total available or required labor hours.

4.1.1. Task Description

The following four categories of maintenance tasks, previously introduced in Chapter 3, have been included in the model and compose the majority of the workload of the airline's wide-body fleet.

- Routine block tasks
- Routine out-of-phase (OOP) tasks
- Non routine tasks
- Modification tasks

A detailed explanation of the implementation of each of these task categories and their respective assumptions will be discussed in Chapter 6. The model divides these four categories of maintenance tasks into two primary groups, the 'fixed' and 'free' groups, as shown by the schematic in Figure 4.1. The 'fixed' tasks are assumed to be fixed to their respective maintenance checks and therefore are not allocated by the optimization model. As can be seen from the schematic in Figure 4.1, the routine block and non routine tasks are classed as fixed. As will be explained in greater detail in Chapter 6, it has been assumed that these tasks are always executed at their respective maintenance check and therefore need not be allocated by the optimization model.

Rather, these tasks have been allocated to the respective maintenance checks prior to the optimization itself. Contrary to the 'fixed' tasks, the tasks classed as 'free' are not restricted to any specific maintenance check and as such will be allocated by the optimization model. These 'free' tasks are the modification and out-of-phase tasks.

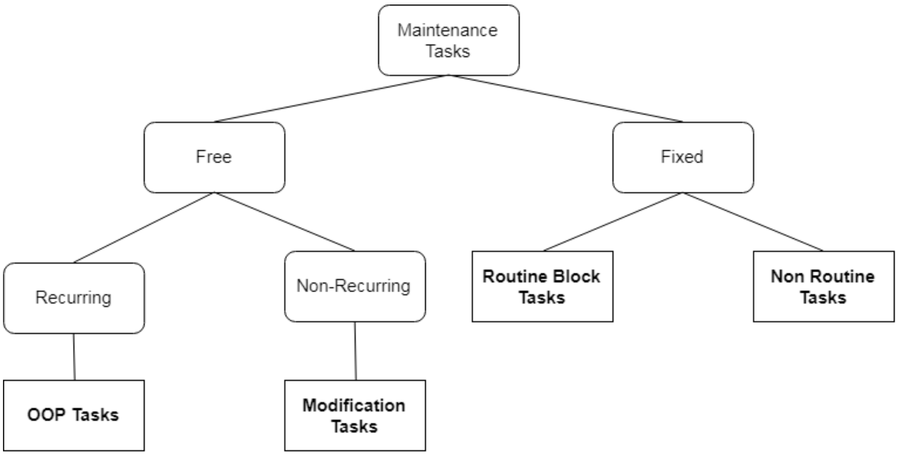

Figure 4.1: The distinction of two groups of maintenance tasks in the model: recurring and non recurring tasks

As can be seen from the schematic in Figure 4.1, the group of 'free' tasks is further divided into two subgroups; the 'recurring' and 'non recurring' tasks. The OOP tasks are classed as recurring tasks because these routine tasks are executed at specific intervals. Thus, an OOP task must be allocated as often as its interval repeats within the planning horizon of the model. Contrary to OOP tasks, modification tasks are classed as non-recurrent as these tasks are only executed once and therefore only need to be allocated once by the model. These distinctions are relevant to the formulation of the model and the primary reason for the iterative loop built around the model. A more detailed explanation is provided in Section 4.3.

Of the four categories of maintenance tasks included in the model, 'only' the modifications and OOP tasks are to be allocated by the model. However, the required workload of the routine block tasks and non routine tasks does affect the available capacity of the various maintenance checks to which the OOP and modification tasks are to be allocated, as shown by the illustration in Figure 4.2. In the Figure, 'NR' refers to the non routine workload and 'Blocks' to the routine block tasks. The workloads of the routine block tasks and non routine tasks of each of the maintenance checks within the planning horizon are computed prior to the optimization. This will be discussed in greater detail in Chapter 6.

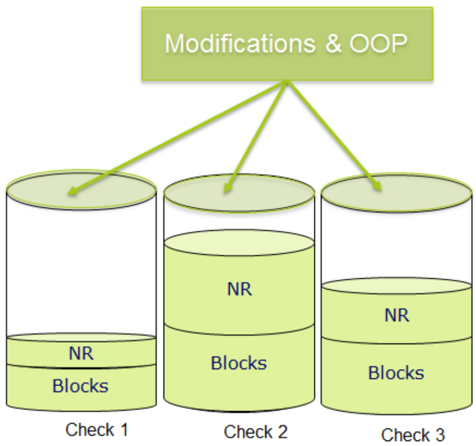

Figure 4.2: Schematic of the allocation of modification and OOP tasks to the maintenance checks, which have limited available capacity due to the routine block and non routine tasks that are fixed to the respective maintenance checks prior to the optimization

Each of the tasks in the model has the following key attributes:

- Aircraft registration
- Due date
- Labor hours required
- Type
- Number
- Interval
- Previous execution date

Each task is assigned to a specific aircraft registration, which may be classed as the 'owner' of the task. The due date of the task represent the latest calender date on which the task must be executed. The labor hours refer to the total hours of work involved with the completion of the task. The task type identifies the task as either a modification or OOP task, and thereby classes the tasks as either recurrent or not. The task number is either the ORP number or MRI code of the task and is used for task identification. The task interval and previous execution date are only applicable to the recurrent OOP tasks and denote the time window within which the tasks is to be executed and the calender date of the previous execution respectively.

4.1.2. Resource Description

The model's resources are the various maintenance checks over the planning horizon. The following maintenance checks have been included in the model:

- A-checks
- Cabin-checks (or B-checks)
- C-checks
- SB-checks (or extra checks)

Each of the letter checks has an available capacity, specific interval and consists of a set of routine maintenance tasks, the routine block tasks. In addition to the letter checks, a fourth category of checks is included in the model, known as the SB-checks, or extra maintenance checks. Unlike the letter checks, these checks are scheduled adhoc when additional capacity is required to execute the workload on a particular aircraft. Thus, a distinction is made in the model between the regular, letter checks and the extra, SB-checks as shown in the schematic in Figure 4.3. A more detailed explanation of these two groups is provided in Chapter 6. The distinction is relevant to the formulation of the model. An important modeling difference between the letter checks and extra checks, is that where the former are assigned to an aircraft registration prior to the optimization, the extra checks are not. Rather a set of extra checks is available throughout the planning horizon of the model. Thus, extra checks may be allocated by the model to an aircraft depending on the need for additional capacity.

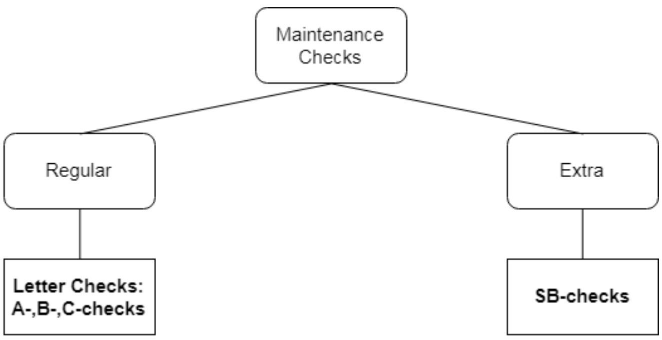

Figure 4.3: Schematic of the distinction between the regular, letter checks and extra, SB-checks

The maintenance checks have the following key attributes:

- Aircraft registration
- Available labor capacity

- Check date
- Number

As aforementioned, the aircraft registration is only appointed to the letter checks. The extra checks are not appointed to any specific aircraft registration prior to the optimization. The available labor capacity indicates the total available capacity of a maintenance check in hours. The check date has been computed prior to the optimization for both letter checks and extra checks. The check number is simply used for identification and tractability.

The planning horizon of the model has been set to four years. Although this value may be altered per the users input, the four year horizon ensures that at least one C-check (3 year interval) is included for all aircraft registrations and thus the inclusion of all letter checks.
Although the majority of all the maintenance tasks and checks are included in the present model, it is important to note that not all maintenance tasks and checks have been included. The following categories of tasks or checks have not been included in the model:

- Maintenance tasks and checks with an interval smaller than A-checks
- Maintenance tasks which are generally executed in H-checks
- Stochastic tasks and checks

The first category includes any line maintenance related items, such as deferred defects, daily and pre-flight maintenance tasks, etc. Deferred defects are corrective tasks which have been not been solved at the time of the finding, but rather have received an interval within which these faults or failures have to be resolved. The model assumes that all corrective maintenance is executed in the maintenance check in which the fault is found, hence it does not include deferred defects.
H-checks, or hangar checks, are relatively short and small maintenance checks performed in the hangars. Although performed in the hangars, these checks are scheduled adhoc for an aircraft when needed. These checks are used for various reasons, including resolving deferred defects, engine and APU changes, engine water washes, and drop-out items. Provided the stochastic nature of these checks and the tasks scheduled in these checks, these tasks and checks have also not been included in the model.
Failures and faults, other than those which are related to the routine maintenance tasks (OOP and block tasks) are not included in the model. The model includes a prediction algorithm which determines the workload of corrective maintenance related to routine maintenance tasks. This estimated workload is assumed to represent all corrective maintenance work and therefore implemented as a deterministic number in the model. The model does not take into consideration any stochastic maintenance task or check occurrences.

4.1.3. Objective Description

The objective of the task scheduling optimization model is to determine the optimal allocation of maintenance tasks over the maintenance checks within a predefined planning horizon. The optimality of the task allocation is driven by efficiency, which has a twofold application. First and foremost the model aims to use the available resource capacity as efficiently as possible. As such, the model addresses a problem comparable with the classical resource constrained project scheduling problem (RCPSP). However, unlike the classical RCPSP, the resource capacity is modeled as a soft constraint. An illustration of the way this has been implemented in the model is shown in Figure 4.4.

As is shown in Figure 4.4, each of the maintenance checks in the model has some finite available resource capacity, identified as the norm. Thus, the norm represents an agreed upon number of labor hours that are acceptable to be allocated during each maintenance check, based on the agreed upon available workforce. Hence, the norm is subsequently used by maintenance planners at the MRO for the allocation of tasks. Although deviations from the norm are allowed and even common in practice, ideally these variations are minimized, such as to control the workloads and thereby the likely on time completion of the maintenance check. A similar philosophy has been implemented into the model.
Furthermore, at the start of the research, the ratio between the total workload of an aircraft and the available capacity of the various maintenance checks over the planning horizon was unknown. Provided this unknown relationship, rather than restricting the available resource capacity in advance, a flexible, soft resource 'constraint' was formulated. As shown in the Figure, the soft resource constraints implies that the allocated workload of maintenance tasks may either be lower or higher than the norm. The former case is referred to

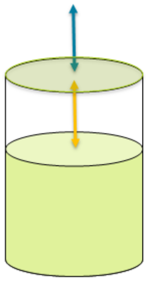

Figure 4.4: Illustration of the definition of shortage and excess hours with respect to the norm of the available resource capacity of a maintenance check

as a shortage, since there is a shortage of workload allocated to the maintenance check. Similarly, when the allocated workload exceeds the norm, this is referred to as an excess.
Thus, unlike the classical RCPSP where the most efficient allocation of tasks to resources is to minimize the makespan of all activities, the present model aims to find the most efficient allocation of tasks to resources, such as to maximize the efficient use of resources. Both the excess and shortage workloads of all maintenance checks are to be minimized by the model. Thus, the model aims to allocate the maintenance tasks such that the available resource capacity is used most efficiently.

In addition to the minimization of excess and shortage workloads of maintenance checks, the model also maximizes the efficiency of using the task intervals. As for the previous objective, this second objective is based on common practice in maintenance planning at the MRO. By maximizing the use of task intervals, the number of recurrences of recurrent tasks are minimized, which is equivalent to less maintenance and hence less cost. Thus, the model aims to allocate maintenance tasks as closely as possible to their respective due dates, whilst ensuring that these are never violated.

It is important to note that contrary to a classic RCPSP, neither time nor precedence relations between tasks are taken into consideration in the present model. The calendar dates of the maintenance checks are fixed and thus maintenance checks are modeled simply as points in time, without physical duration. As indicated by Figure 4.2, the model only considers how 'full' the bins (maintenance checks) are, not their width. In other words, groundtime is not considered in the model.

4.2. Mathematical Formulation

This Section presents the mathematical formulation of the model followed by a brief explanation of the formulation of the decision variables, objective function and constraints. The model may be classified as a deterministic mixed integer linear programming model (MILP), consisting of two sets of binary decision variables and three sets of continuous decision variables. Built around the optimization is an iterative loop, resulting in consecutive optimization runs of the deterministic MILP model. Each run utilizes the output of the previous run as input for the next optimization run. As such, the model may also be classified as a heuristic model with an approximate optimal solution. This iterative solution technique will be discussed in greater detail in Section 4.3.

Indices and Sets

A	set of aircraft
C	set of maintenance checks
T	set of maintenance tasks
a	aircraft *($a \in A$)*
c	general notation for a check (regardless of type) *($c \in C$)*
c_r	regular check (letter check) *($c_r \in C$)*
c_e	extra check (SB-check) *($c_e \in C$)*
t	task *($t \in T$)*

Parameters

$dd_{a,t}$	due date of task t of aircraft a
$ped_{a,t}$	previous execution date of task t of aircraft a
$mh_{a,t}$	required capacity hours of task t of aircraft a
$int_{a,t}$	interval of task t of aircraft a
$cd_{a,c}$	check date of check c of aircraft a
$AMH_{a,c}$	available capacity hours of check c of aircraft a
$SEH_{a,c}$	excess hours already scheduled in previous optimization runs
C_{lb}	labor costs per hour
C_{gt}	groundtime costs per hour
C_{sh}	cost factor for unused capacity (shortage)
C_{in}	cost factor for remaining interval
C_y	cost factor for using an extra check
$EH_{1,2,3}$	breakpoints of piecewise linear function for excess penalty

Decision Variables

$$x_{a,c,t} \quad \begin{cases} 1; \text{if task } t \text{ of aircraft } a \text{ is scheduled to check } c \text{ of aircraft } a \\ 0; \text{otherwise} \end{cases}$$

$$y_{a,c} \quad \begin{cases} 1; \text{if check } c \text{ will be used for aircraft } a \\ 0; \text{otherwise} \end{cases}$$

$EH_{a,c}$	excess hours of check c of aircraft a
$SH_{a,c}$	shortage hours of check c of aircraft a
$EP_{a,c}$	excess penalty of check c of aircraft a

Objective Function

$$min. \sum_a \sum_c EP_{a,c} + \sum_a \sum_c C_{sh} \cdot SH_{a,c} + \sum_a \sum_c \sum_t x_{a,c,t} \cdot C_{lb} \cdot mh_{a,t} \cdot \left(\frac{dd_{a,t} - cd_{a,c}}{int_{a,t}} \right) + \sum_a \sum_c C_y \cdot y_{a,c} \tag{4.1}$$

Subject to

$$\sum_t x_{a,c,t} \cdot mh_{a,t} + SH_{a,c} \geq AMH_{a,c} \cdot y_{a,c} \quad \forall\ c \in C, \forall\ a \in A \tag{4.2}$$

$$\sum_t x_{a,c,t} \cdot mh_{a,t} - EH_{a,c} \leq AMH_{a,c} - SEH_{a,c} \quad \forall\ c \in C, \forall\ a \in A \tag{4.3}$$

$$EP_{a,c} = C_{lb} \cdot EH_{a,c} + C_{step} \quad \forall\ c \in C, \forall\ a \in A \tag{4.4}$$

where:

$C_{step} = 0$ for $0 \leq EH_{a,c} < EH_1$
$C_{step} = C_{gt}$ for $EH_1 \leq EH_{a,c} < EH_2$
$C_{step} = 2 \cdot C_{gt}$ for $EH_2 \leq EH_{a,c} < EH_3$

$$x_{a,c,t} \cdot (dd_{a,t} - cd_{a,c}) \geq 0 \quad \forall\ t \in T, \forall\ c \in C, \forall\ a \in A \tag{4.5}$$

$$x_{a,c,t} \cdot (cd_{a,c} - ped_{a,t} - 1) \geq 0 \quad \forall\ t \in T, \forall\ c \in C, \forall\ a \in A \tag{4.6}$$

$$\sum_c x_{a,c,t} = 1 \quad \forall\ t \in T, \forall\ a \in A \tag{4.7}$$

$$\sum_{c} y_{a,c} \leq 1 \quad \forall\ c_e \in C, \forall\ a \in A \tag{4.8}$$

$$y_{a,c} = 1 \quad \forall\ c_r \in C, \forall\ a \in A \tag{4.9}$$

$$P \cdot x_{a,c,t} - Q \cdot y_{a,c} + R \leq P \quad \forall\ t \in T, \forall\ c \in C, \forall\ a \in A \tag{4.10}$$

where:
$P > Q > R$ are constants

Domain Definition

$x_{a,c,t}$	$\in \{0,1\}$	$\forall\ t \in T, \forall\ c \in C, \forall\ a \in A$
$y_{a,c}$	$\in \{0,1\}$	$\forall\ c \in C, \forall\ a \in A$
$EH_{a,c}$	≥ 0	$\forall\ c \in C, \forall\ a \in A$
$SH_{a,c}$	≥ 0	$\forall\ c \in C, \forall\ a \in A$
$EP_{a,c}$	≥ 0	$\forall\ c \in C, \forall\ a \in A$

4.2.1. Decision Variables

Five sets of decision variables have been formulated, two sets of binary decision variables and three sets of continuous variables. The set of $x_{a,c,t}$ is a set of binary decision variables $\in A, C, T$ that equal one when task t of aircraft a is scheduled to check c of aircraft a. As such this set of binary decision variables controls the allocation of a maintenance task to a maintenance check. The second set of binary decision variables, $y_{a,c}$ $\in A, C$, equal one when check c is allocated to aircraft a, and zero otherwise. Thus, this second set of binary decision variables controls the allocation of maintenance checks to aircraft. As will become clear from the constraints, this set of decision variables is particularly significant for the allocation of the extra maintenance checks.

The set of decision variables denoted as $EH_{a,c}$ is a set of continuous, non-negative decision variables $\in A,C$, and represents the excess hours of check c of aircraft a as a result of the excessive task allocation. Similarly, the set of decision variables denoted as $SH_{a,c}$ is also a set of continuous, non-negative decision variables $\in$ A,C for the shortage hours of check c of aircraft a resulting from a shortage of allocated tasks. The third set of continuous, non-negative decision variables, $EP_{a,c} \in A,C$, defines the excess penalty, which controls the penalties corresponding to the excess hours of the maintenance checks. A more detailed explanation will be provided hereafter in 4.2.2 as well as in Chapter 5.

4.2.2. Objective Function

The objective function formulated by Equation 4.1 is a cost minimization function which consists of four elements. The first element of the objective function identifies the excess penalty, which is the penalty corresponding to the excess hours of check c of aircraft a. The relation between the excess hours and the excess penalty is controlled by the set of constraints shown in Equation 4.4, which will be discussed hereafter. The second element of the objective function composes the penalty for a shortage of scheduled hours for check c of aircraft a, formulated as the sum of the product of the penalty factor for shortage hours, C_{sh}, and the corresponding binary decision variable $x_{a,c,t}$. An increase in shortage hours and excess hours results in a increase of the objective function value. Thus, the formulation of the excess and shortage penalties ensures that the model will optimize for maximum efficiency of the use of available resource capacity of maintenance checks as explained in 4.1.3.

The third element of the objective function concerns the penalty for unused task interval. As explained in 4.1.3, in addition to maximizing the efficient use of resources, the model also aims to maximize the use of task interval. The penalty for unused interval is proportional to the ratio of unused interval over the interval of the respective task. The unused interval is defined as the days between the due date of task t of aircraft a and the scheduled check date of check c of aircraft a to which that task has been allocated. This ratio is subsequently scaled by the required capacity hours of the corresponding task, which in turn is multiplied by a cost factor for the labor costs per hour. Thus, the unused interval penalty is equivalent to the additional labor cost incurred by wasting interval days.

The last term in the objective function is the penalty for the use of extra checks. Although available to be used when extra resource capacity is required, the extra checks are to be minimized as they imply downtime on the aircraft and hence operational costs. The penalty consists of the product of the binary decision variable $y_{a,c}$ and a cost factor, C_y.

As all four terms of the objective function are expressed in monetary terms, the objective of the model is to minimize the combined cost of the various penalty terms. The penalties are correlated through their formulations; less excess penalty may require a higher extra check penalty and interval penalty for instance. The ratio of the weights of the various penalty factors may be used to control the frequency of a penalty with respect to the others and as such may be used to control the task allocation behavior of the model. The computation of the various cost factors involved with the model and weights of the penalty terms are the primary topics of Chapter 5.

4.2.3. Constraints

The following provides a brief explanation of the formulation of the various sets of constraints in the model.

Set of Shortage Hour Constraints

The set of constraints described by Equation 4.2 ensures the correct computation of the values of the shortage hour decision variables. The left-hand side of the constraints are the sum of the labor hours of the allocated tasks and the shortage hours. The right-hand side are the available labor hours of the maintenance check. The product with the $y_{a,c}$ decision variables ensures that the unused capacity of maintenance checks which are not allocated by the model are not included. As illustrated by the schematic in Figure 4.5 and as per the definition in Equation 4.2 the shortage hours must be greater or equal to the difference between the available resource capacity and the labor hours of the allocated tasks. As the objective forces the shortage hours to a minimum, the hours will always be equal to the lowest feasible value. In case of an excess, according to the formulation of the constraint the shortage hours would equate to a negative number, in which case the shortage hours are set to zero, as per the domain definition.

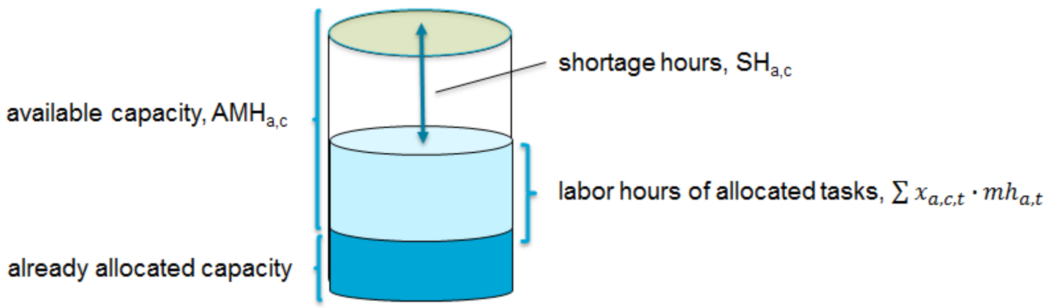

Figure 4.5: Illustration of the computation of the shortage hours as per the constraints in Equation 4.2

Set of Excess Hour Constraints

Equation 4.3 describes the set of constraints that controls the proper computation of the excess hours of the maintenance checks. It's formulation is very similar to the previous set of constraints for the shortage hours. Essentially, the formulation implies that the excess hours must at least be equal to the difference between the allocated workload and the available capacity. The right-hand side of Equation 4.3 has the additional term $SEH_{a,c}$. As has briefly been mentioned, an iterative loop has been build around the optimization model, such that the outcome of the previous run is used as an input for the next. The term $SEH_{a,c}$ represents the scheduled excess hours which may have accumulated over past runs, ensuring a continuity of the excess hours over the various consecutive optimization runs. Thus, if any additional tasks are allocated during the present optimization run, the corresponding labor hours are added as excess hours to excess hours accumulated over previous runs.

Set of Excess Penalty Constraints

As noted in 4.2.1, Equation 4.4 describes the relation between the excess hours and the corresponding penalty. A piecewise linear function may be used to describe the relation between the excess hours and excess penalty. A simplified representation is shown in Figure 4.6.

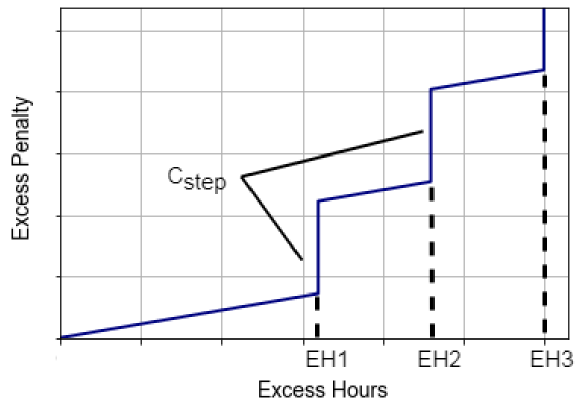

Figure 4.6: A simplified representation of the piecewise linear function that described the relation between the excess hours and excess penalty

The slope of the piecewise linear function has been set equal to the labor cost, C_{lb}. The points EH_1, EH_2 and EH_3 are the breakpoints of the piecewise linear function and as such represent the excess hours at which a step or jump in the function has been positioned. The stepsize scales with the cost of groundtime. A more detailed explanation of the formulation of the piecewise linear function and the cost is provided in Chapter 5.

Set of Due Date Constraints

The set of constraints described by Equation 4.5 ensure that the due dates of tasks are never violated. The formulation requires that the difference between the due date of task t of aircraft a and the check date of maintenance check c of aircraft a to which task t has been allocated is at least zero. Although in practice a due date may be violated depending on the task, it is assumed that this is never planned for. Thus, the model enforces the prohibition of the violation of due dates as a hard constraint.

Set of Previous Execution Date Constraints

Constraints 4.6 are known as the "previous execution date" constraints. As has been explained, the model distinguishes between recurrent and non-recurrent tasks. The OOP tasks are recurrent as these tasks must be scheduled routinely within a specified interval. The present set of constraints prevents the model from allocating a future occurrence of an OOP task to a maintenance check prior to the previous occurrence of the task in question. Thus, the set of constraints introduces some form of precedence for the recurring OOP tasks.

Set of Task Occurrence Constraints

The next set of constraints formulated by Equation 4.7 ensures that each task is only scheduled once to a maintenance check.

Set of Extra Check Constraints

Similar to the previous set of constraints, constraints 4.8 ensure that each extra maintenance check is scheduled to at most one aircraft. As aforementioned, unlike the letter checks which have been assigned to a specific aircraft registration prior to the optimization, the extra checks are free to be allocated by the model depending on the need for extra resource capacity.

Set of Regular Check Constraints

The binary decision variables $y_{a,c}$ control the allocation of maintenance checks to aircraft registrations. Provided that the letter checks are regular checks, rather than the optional extra checks, the set of constraints represented by Equation 4.9 ensure that for the letter checks the decision variables $y_{a,c}$ are always set to one. As such these constraints ensure that the model 'recognizes' the availability of the letter checks for task allocation.

Set of X-Y Correlation Constraints

The last set of constraints is described by Equation 4.10 and ensures the proper correlation between the two sets of binary decision variables. Consequently, they prevent maintenance tasks from being allocated to a maintenance check which is not allocated to an aircraft. This set of constraints is therefore primarily relevant for the extra checks. P, Q and R are randomly selected constants, where P must be larger than Q and Q must be larger than R.

4.3. Iterative Solution Technique

The previous Sections have explained the setup of the MILP model. As aforementioned, an algorithm has been built around the MILP optimization model, resulting in an iterative loop of consecutive optimization runs where the output of the former run is used as an input for the next. Hence, the model may be classed as a heuristic model, with an approximate optimal solution.

The iterative solution technique was introduced to the MILP model to deal with the allocation of the recurring OOP tasks. The periodicity of these tasks over a specified time interval implies that the allocation of such a task to a maintenance check affects the due dates of all future occurrences of the task in question as well, and subsequently where these future occurrences can and cannot be allocated. The iterative solution technique breaks the 'chain' of occurrences of recurrent OOP tasks into individual occurrences. Thus, during each optimization run, only a single occurrence of an OOP task is allocated by the model. Following the allocation, the due date of the subsequent occurrence of that task is computed using the check date of the maintenance check to which the former occurrence has just been allocated and the task interval. This new due date is subsequently fed as an input for the next occurrence of the OOP task in the next optimization run. This process keeps repeating itself until the occurrences of the tasks fall outside of the planning horizon. A schematic representation of this iterative solution technique and its interaction with the MILP model is depicted in Figure 4.7.

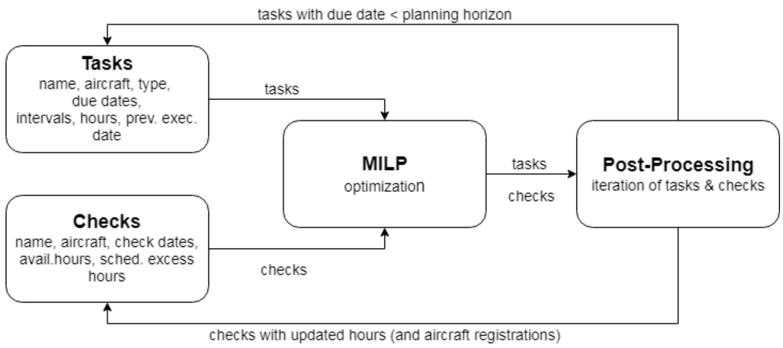

Figure 4.7: Schematic of the iterative solution technique built around the optimization run

The module indicated as 'Post-Processing' filters between the non-recurrent modification tasks and the recurrent OOP tasks. Provided that the modifications are non-recurrent, these tasks are all allocated during the initial run only. For the OOP tasks the module determines the due dates for the next optimization run based on the check date of the latest execution of the OOP task and its respective interval. As aforementioned, a task is passed onto the next optimization run, until its due date falls outside of the planning horizon. More specifically, a task is passed onto the next optimization run, until the due date falls outside of the planning horizon *plus* the interval of one A-check. It has been assumed that when the due date exceeds this threshold, the occurrence of an OOP task may be scheduled to one the regular maintenance checks outside of the planning horizon and therefore is no longer considered in the subsequent optimization runs of the model. Thus, the list of tasks gradually decreases over the subsequent runs. Once there are no more tasks to be allocated, the iterative solution algorithm stops, and the optimization is completed.

Additionally, the 'Post-Processing' module also updates the available resource capacities of the maintenance checks (denoted as $AMH_{a,c}$ in the mathematical formulation) based on the tasks which have been allocated in the latest optimization run. In case of excess, the available resource capacity is set to zero and the scheduled excess hours ($SEH_{a,c}$) are set to the accumulated excess hours of all previous optimization runs. This ensures that the up-to-date available capacity, or excess hours, of all maintenance checks are used in each subsequent optimization run.

In case an extra check has been allocated to an aircraft during the optimization run, the 'Post-Processing' module appoints the extra check to its allocated aircraft registration for the subsequent optimization runs. Similar to the regular letter checks, this extra check will have a 'pre-appointed' aircraft registration for all subsequent optimization runs.

4.4. Software and Hardware

All model algorithms were coded in Python 3.5. CPLEX version 12.7.1 was used to solve the formulated MILP. The choice for Python was primarily driven by the fact Python is the preferred programming language at the MRO. Moreover, the version of Python was driven by the compatibility with the used version of CPLEX. The choice for CPLEX was primarily driven by the availability both at the MRO and at Delft University of Technology. The experiments were run on a PC with an Intel Core i5 2.3 GHz processor and 4 GB of RAM, equipped with Windows 7 Professional.

5

Penalty Formulation and Computation

Following the introduction of the model in Chapter 4, this Chapter elaborates specifically on the formulation, computation and reasoning of the various penalty terms in the objective function of the model. First, Section 5.1 will provide a detailed description of the formulation of the piecewise linear functions that constitute the excess penalties corresponding to the excess hours. A detailed description of the need for extra maintenance checks and their respective penalties is provided in Section 5.2. Subsequently, Sections 5.3 and 5.4 discuss the shortage and interval penalties respectively. Finally, Section 5.5 discusses some additional penalty factors which have been included in the task scheduling optimization model.

5.1. Excess Hour Penalty

This Section describes the excess penalty imposed by the model when the required resource capacity of the allocated maintenance tasks exceeds the available resource capacity of the maintenance check. As described in Chapter 4 the relation between the excess hours and the excess penalty is described by piecewise linear functions. This Section describes the logic and formulation of the piecewise linear relations. Firstly, a study on the primary costs of maintenance is briefly discussed. This study provides some background to the excess penalty formulation. This formulation is discussed in detail in Section 5.1.2.

5.1.1. Maintenance Costs

A distinction can be made between direct and indirect maintenance cost. In the book "Reliability, Maintenance and Logistic Support - A Life Cycle Approach", Kumar et al. state that "the direct cost of aircraft maintenance is composed of cost for resources, which consists of spare parts, materials, personnel, tools and equipment, facilities and technical data" [50]. Papakostas et al. distinguish "equipment and facility costs, supplies and logistic costs, personnel costs and overhead" as the primary categories that make up "the costs related to aircraft maintenance" [51]. Over the year 2016, IATA reported that airlines around the world spent an average of 9.5% of the total operational expenses on MRO services [4]. It is important to note that this figure refers to the direct cost of maintenance (DMC) only. A simplified breakdown of the primary DMC as reported by IATA is shown in Figure 5.1.
As can be seen from Figure 5.1, IATA breaks down the DMC in labor costs, material costs, life limited part (LLP) costs and the costs of subcontracting or outsourcing. The labor cost are estimated at 15% of the total DMC. By far the largest reported costs are due to subcontracting. The right piechart in Figure 5.1 shows that approximately 15% of all DMC is spent on Airframe base maintenance. Again, it is important to note that these figures only provide an indication of the DMC, and do not apply to the MRO specifically.

In addition to DMC, the indirect cost of maintenance (IMC) are significant and must be considered. Saltoglu et al. refer to the importance of including the cost of aircraft downtime as an indirect cost to be considered in the total cost of aircraft maintenance [10]. Referring to downtime, they state that "this cost is defined as the cost of lost revenue"[10]. In referring to downtime costs, Saranga adds that "cost of lost revenue (CLR) is an inevitable cost, and is common to any system that has to be taken off operation, whether planned or unplanned"[11]. As stated by Saltoglu et al. and Saranga, the cost of aircraft downtime, is the loss of revenue

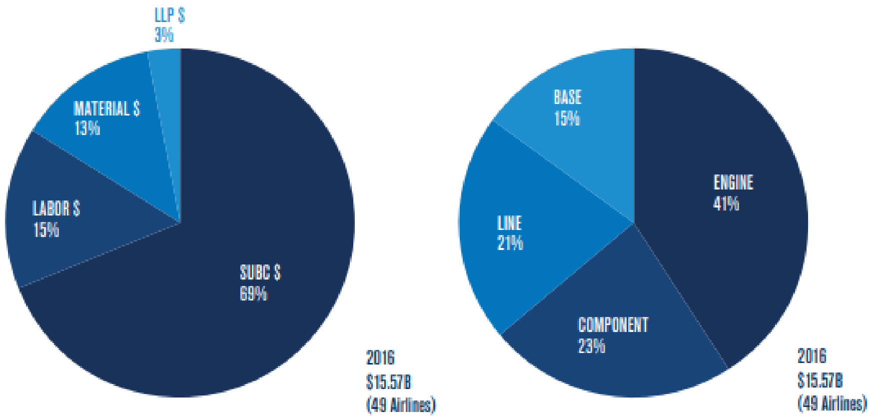

Figure 5.1: Breakdown of the primary direct maintenance costs (DMC) of aircraft maintenance (left) and breakdown of maintenance cost per segment (right) according to IATA's 2016 Airline Maintenance Costs Analysis [4]

which could have been generated if the aircraft would have been available. Ultimately, the MRO services are performed to ensure the reliability and availability of the operator's aircraft fleet. Hence, any (maintenance) delay or additional required maintenance downtime, may be considered as a loss of aircraft availability. In turn, these losses in aircraft availability have a negative impact on flight operations. However, even the loss of aircraft availability without any direct operational disruptions such as delays or cancellations may of itself be considered as a cost factor. A lower aircraft availability simply implies less flight time. Subsequently, loss of aircraft availability reduces the number of flights which can be scheduled over a prolonged period of time and as such potential revenues. Thus, the loss of aircraft availability as a result of maintenance delays or additional maintenance groundtime has a negative impact on the profitability of the aircraft and should be considered as an additional cost of aircraft maintenance.

5.1.2. Modeling of Excess Hour Costs

Ultimately, excessive workloads will result in maintenance delays because the available labor capacity simply cannot finish the workload within the allotted time slot of the maintenance check. Although more maintenance production workers may be allocated from other maintenance checks to the maintenance check with an excessive workload, this extra available maintenance workforce is finite. More importantly, precedence relations between the maintenance tasks also imply a finite number of maintenance tasks which can be executed at any point in time. Thus, excessive maintenance workloads can contribute to maintenance delays and ultimately to disruptions in the flight network. Although precedence relations between tasks are not explicitly taken into consideration in the present task scheduling model, the positive correlation between excessive maintenance workloads and maintenance delays should be taken into consideration in the penalty factor for excess hours. Moreover, this relation may be considered as non linear as an increasing number of excess hours will induce increasingly larger maintenance cost, starting with merely the costs of additional labor scheduled to the aircraft, and increasing over time with the additional costs of delays.

Piecewise linear functions were used to approximate the non linear relation between the excess hours and the corresponding cost or penalty. Of the various maintenance costs, two primary costs were included in the model: the labor cost of additionally allocated maintenance personnel and the cost of lost availability of the aircraft. The cost of flight delays were also considered as an alternative to the cost of lost aircraft availability. However, maintenance delays do not always incur a disruption to the flight network. Although in most cases a flight will be scheduled in a mere couple of hours after the release time of the aircraft from its maintenance slot, this small window provides a buffer which may prevent flight disruptions in case of a maintenance delay. Moreover, some airlines have an additional in the form of an actual buffer aircraft which, if available, will be used to prevent any flight disruptions in case of a severe maintenance delay. Thus, modeling the costs of flight delays appears more complicated to model because of the various dependencies involved. Instead, the costs of loss of aircraft availability is more comprehensible and straightforward to model.

The labor cost corresponds to the cost of extra allocated maintenance personnel to the aircraft with an excessive workload, which would have otherwise worked on another aircraft or maintenance check. Thus, the additional cost are not the cost of hiring additional maintenance workers, but of allocating these maintenance workers from one maintenance check to the maintenance check with excess scheduled hours. The annual cost of one maintenance engineer for the MRO has been estimated at ANNLAB. Furthermore, the annual hours worked by an engineer have been estimated to HLAB labor hours. These numbers equate to approximately CLAB per maintenance engineer per hour. This has been assumed as the cost of labor in the model.

The cost of lost aircraft availability refers to the loss of potential profit through a loss of operational availability of the aircraft, for example because of maintenance downtime. During this study no specific data was found on the hourly profitability of the airline's wide-body aircraft. The 2017-2018 Airline Economic Analysis performed by Oliver Wyman provides insight into the revenue per availalbe seat mile (RASM) and cost per available seat mile (CASM) of various airlines based in North America [9]. The following figures have been taken from Oliver Wyman study. The original values of the RASM and CASM in dollar cents as reported by Oliver Wyman are listed on the left.

Table 5.1: Revenues, costs and profit margins per available seat mile in dollar cents [9]

	RASM	CASM	Margin
Delta	16.4	13.3	3.1
United	14.8	12.7	2.1
American	15.5	13.3	2.1
Hawaiian	14.2	11.1	3.2
Virgin/Alaska	13.5	10	3.4
Average	14.88	12.08	2.78

It has been assumed that the airline in question in this research is a legacy network carrier comparable with Delta, United, American and Hawaiian Airlines which have been included in the study of Oliver Wyman. As such, the averages of the RASM and CASM have been assumed as the estimated RASM and CASM values for the airline's wide-body fleet. Thus, the profit margin per available seat mile ($PASM$) is 2.78 dollar cents. To derive the potential profit of one wide-body aircraft per hour, the following equation has been used:

$$PROFIT = \frac{PASM \cdot S \cdot LF \cdot V}{100} \tag{5.1}$$

Here, S represents the total number of seat onboard of an aircraft; LF represents the load factor of the aircraft; and V represents the average velocity of the aircraft. Furthermore, the division by 100 converts the cents to dollars.
The airline's wide-body aircraft have approximately XSEAT seats in their present configuration. A load factor of LF% has been assumed. The average flight velocity has been estimated at VCR miles per hour. Inserting these values for the parameters results in a profit per hour of approximately LOSTP dollars. Thus, one wide-body aircraft has a potential profit per hour of LOSTP dollars. Equivalently, one hour lost in aircraft availability equates to a potential profit loss of LOSTP dollars. Although this estimation is very simple and crude, it provides a estimate of the scale of the costs of lost aircraft availability and this rough assumption is considered to be sufficient for the purposes in the model. A critical view on the cost factors will be discussed in the sensitivity analysis in Chapter 9.

Provided the estimates for the cost of labor and the cost of lost aircraft availability, next to determine is how to implement these costs into the model. A schematic of the design of the piecewise linear function which has been implemented in the model is shown in Figure 5.2. The PWLF implemented in the model approximates a discontinuous PWLF, with large steps at the two breakpoints. However, for the ease of modeling, the PWLF has been modeled as a continuous PWLF, where the steps are modeled as segments with large slopes over very small stepsizes in the x-domain.

The following assumptions have been made in the design of the piecewise linear function:

1. A small number of maintenance workers with a finite number of available labor hours is available at all maintenance checks and may be allocated to the maintenance check with excess scheduled hours. The

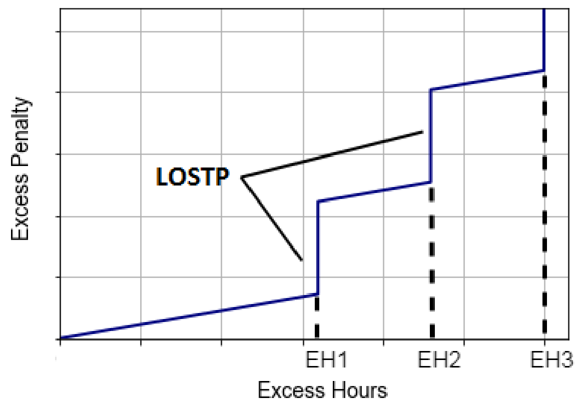

Figure 5.2: Piecewise linear function of the excess penalty versus excess hours as used in the model

use of this extra workforce is penalized by the labor cost. These costs represent the first linear segment of the piecewise linear function shown in Figure 5.2.

2. If the excess hours exceed the available capacity of the extra workforce, it is assumed that the maintenance check will be delayed by one hour. This will incur a one hour loss of aircraft availability, which is imposed instantly, as shown by the step in the function. Furthermore, it is assumed that during this hour a complete shift of maintenance personnel is allocated to the maintenance check. This additional workforce is again penalized by the labor cost and is represented by the second linear segment in the piecewise linear function.

3. Should one additional hour of an entire shift of maintenance personnel still proof insufficient additional resource capacity, the maintenance check is once more delayed by another hour, imposing again the costs of lost aircraft availability and the subsequent additional labor cost.

4. After a two hour delay of the maintenance check, it is assumed that the delayed maintenance check will incur severe disruptions to the flight network, necessitating flight delays or cancellations or the need for a buffer aircraft. These costs have not been quantified. Rather, the post-slope of the PWLF has been set equal to the slopes of the steps, resulting in significant penalties for any remaining excess hours.

It has been assumed that there is a small, finite workforce available at each maintenance check, which can be allocated to the maintenance check with excess scheduled hours. The MRO's planning policy for the hangar in which the wide-body fleet of the airline is maintained prescribes the allocation of available maintenance capacity to the various maintenance checks for each day. According to this schedule, in addition the A-checks of a the wide-body aircraft, there are CORS hours available for small corrective maintenance checks, also referred to as H- or TO-checks. Moreover, the schedule prescribes that there is a team of WAOG maintenance workers available at all times to be allocated in case of an AOG event, or aircraft-on-ground. An AOG event applies when a fault or failure has been found on an aircraft which deems it unfit to fly and thus requires immediate maintenance action. Of course if there is no AOG event, these maintenance workers may be allocated to either the A-check or TO/H-checks. Moreover, the hangar policy prescribes that in case of a "heavy A-check" (max. AHEAV hours), resources are shifted from the H/TO checks to the A-check.

Based on the hangar policy it has been assumed that this small number of WAOG additional maintenance workers is available in case of a "heavier" maintenance check, or a maintenance check with excess scheduled hours. Moreover, it has been assumed that these additional maintenance workers are initially only allocated for one shift, resulting in EH1 additional available hours, which are only penalized by the cost of labor. If the excess hours are larger than EH1 hours, the maintenance check will be delayed by one hour, incurring the additional excess penalty corresponding to the loss of aircraft availability. Thus, the first breakpoint in the piecewise linear function is located at EH1 hours.
The hangar policy also specifies that at least WSHIFT maintenance workers must be available during each maintenance shift. Hence, in case of more than EH1 excess hours, a shift of WSHIFT maintenance workers is assumed to be allocated to the maintenance check, in addition to the previously allocated additional WAOG

maintenance workers. This is represented by the second linear segment in the piecewise linear function, ranging from EH1 to EH2 hours. A second one hour delay is incurred if the excess hours exceed EH2 hours, as shown by the second step in the piecewise linear function. As for the first delay, again WSHIFT plus WAOG labor hours are assumed to be available during this second hour delay as indicated by the third linear segment in the piecewise linear function. If the excess hours exceed EH3 hours, it is assumed that the maintenance delay incurs a disruption to the flight network, modeled by the steep post-slope of the piecewise linear function.

Thus, the breakpoints of the PWLF are located at EH1, EH2 and EH3 excess hours respectively. The slopes of the linear segments are equal to the cost for labor. The stepsize located at the breakpoint equal the cost of lost aircraft availability. The slopes, locations of breakpoints, number of breakpoint and stepsizes are all modeled as parameters and may be changed as desired.

5.2. Extra Maintenance Checks

In Chapter 4 a distinction was made between two types of maintenance checks in the model: regular and extra maintenance checks. The most important difference between these two groups is that the regular, letter checks have already been appointed to an aircraft registration prior to the optimization, whereas the extra checks are not. Rather, the extra checks may be allocated by the model based on the need for extra resource capacity. This Section elaborates on the logic behind these extra maintenance checks from the perspective of maintenance planning at the MRO for the airline's wide-body fleet. Moreover, this Section also elaborates on the penalty which is imposed when such an extra maintenance check is allocated.

5.2.1. The Need for Extra Checks

In addition to the regular maintenance slots for the letter checks (A, B, C-checks), additional maintenance checks may be required. These extra maintenance checks may be required when there is a fault or failure on the aircraft which requires corrective maintenance. At the MRO additional maintenance checks may also be scheduled to reduce the workload of upcoming maintenance checks. These additional maintenance checks are also referred to as SB-slots, short for "service bulletin slot", as they are generally used to reduce the workload of modifications (also referred to as service bulletins) on the aircraft. With a greater need for modifications than anticipated at the introduction of the wide-body aircraft to the airline's fleet, the workload of the modification currently exceeds the available capacity of the regular letter checks. In addition to the excessive workload of modifications, there are other considerations that also contribute to the need for these extra maintenance checks. Some modifications require specific configurations of the aircraft, such as "power-off" conditions in which all electrical power has to be switched up during the execution of the task. These conditions can disrupt the regular maintenance inspection tasks of the letter checks, making it highly undesirable to schedule such modifications in any of the regular letter checks. Hence, for such modifications these extra checks provide an alternative and less disruptive maintenance opportunity. However, it is considered in general that modifications disrupt the regular maintenance checks, regardless of the need for special conditions. Contrary to routine preventive maintenance tasks which have a regular frequency with which they are executed, modifications are generally only executed once on each applicable aircraft. This "one off" characteristic of a modification task implies that the maintenance engineers will likely have to spend more time to "get the job right", studying the requirements and instructions of the modifications, compared to the other routine maintenance tasks with which they are more familiar. Moreover, once started on the execution of the modification task, it may be discovered that more work is required than initially planned for. These are some of the potential risks with modifications which may severely disrupt the otherwise routine maintenance checks. All of these considerations contribute to the need for extra maintenance checks within the MRO.

The SB-checks are unused A-check slots which are ordinarily allocated to an aircraft by the groundtime and maintenance planners. The MRO's current groundtime planning offers two A-checks slots per week for the airline's wide-body fleet, but not all of these slots are used. Remaining available slots may therefore be scheduled as SB-checks. An overview of the number of A- and B-checks over the planning horizon from mid 2018 to mid 2022 for the wide-body fleet as computed prior to the optimization in the Data Preparation Process is shown in Table 5.2.

As can be seen from Table 5.2, the number of A- and B-checks is spread approximately evenly over the various years and aircraft. Based on the analysis in Table 5.2 it may be concluded that the annual total number of

Table 5.2: Annual overview of the number of A- and B-checks on the airline's wide-body fleet over the planning horizon from mid 2018 till mid 2022 per aircraft (ac)

	ac 1	**ac 2**	**ac 3**	**ac 4**	**ac 5**	**ac 6**	**ac 7**	**ac 8**	**ac 9**	**ac 10**	**ac 11**	**ac 12**	**ac 13**	**Total**
2018	3	2	2	2	2	3	3	3	4	4	2	3	1	34
2019	5	5	5	6	5	4	5	5	4	4	5	5	5	63
2020	5	5	5	5	6	5	5	5	5	5	5	5	4	65
2021	5	5	5	5	5	5	5	5	5	5	6	5	5	66
2022	3	3	3	2	2	3	3	3	4	4	2	3	4	39
Total	21	20	20	20	20	20	21	21	22	22	20	21	19	267

checks ranges between 60 and 65 checks. The years 2018 and 2022 deviate from the other years in the planning horizon since the planning horizon used for the analysis shown by Table 5.2 runs from mid 2018 to mid 2022.

Based on this analysis it has been assumed that at most SBN extra maintenance checks are available for the wide-body fleet per month over the entire planning horizon. Furthermore, it has been assumed that each of these maintenance checks has the same norm of available resource capacity as a regular A-check, which is equivalent to ANORM labor hours. Since the current model does not include an optimization of the maintenance check allocation, an assumption has been made that these extra checks are always scheduled one the first and 15th day of each month within the planning horizon, regardless of when the other regular maintenance checks are scheduled. These dates may be altered as desired. Finally, it has been assumed that in addition to modifications, out-of-phase tasks may also be scheduled to these extra maintenance checks.

With regard to the iterative solution technique built around the model, it is important to note that in the current setup of the model, the extra checks are only permitted to be allocated during the initial optimization run of the model. In all subsequent optimization runs, any remaining extra checks which were not allocated on the initial run are not available to the model. Earlier model runs showed that such a restriction was necessary to prevent extra checks from being allocated with very little tasks allocated to it. The restriction is implemented in the model as a simple adjustable parameter. As such, the number of optimization runs during which extra checks are available for allocation may be adjusted. However, in the basic setup of the model this is restricted to the initial run only. The adjustment of this parameter has been studied in the sensitivity analysis of the model. The interested reader is referred to Chapter 9.

5.2.2. Extra Maintenance Check Penalty

Despite the need for extra maintenance checks, their use must be limited as much as possible to maximize the aircraft's availability. Thus, whenever an extra check is allocated to one of the aircraft registrations by the model, a penalty must be imposed. Equation 5.2 is part of the objective function of the model and is used to impose the penalty for extra checks. It consists of the product of the binary decision variable $y_{a,c}$ and the cost factor C_y, where $y_{a,c}$ is one if extra check c is allocated to aircraft a, and zero otherwise. It is important to note that this penalty term only applies to the extra maintenance checks and not the the regular maintenance checks. Thus, Equation 5.2 ensures that whenever an extra check is scheduled, a penalty is added to the objective function, which penalties are summed for all aircraft and all extra checks.

$$\sum_a \sum_c C_y \cdot y_{a,c} \tag{5.2}$$

The question that remains to be answered is what should C_y be such that the use of an extra maintenance check is sufficiently discouraged to maximize aircraft availability, whilst also ensuring that an extra check is allocated to avoid significant excess hours at the regular maintenance checks?
Based on the same logic as has been applied for the excess hour penalty, both labor costs and the cost of lost aircraft availability apply to the extra maintenance checks. In the previous Section the cost of one hour lost in aircraft availability was computed at LOSTP euros per hour. As aforementioned, an extra check is an empty A-check, which has a GTA hour groundtime. This would equate to a loss in aircraft availability of LOSS1 euros per extra maintenance check. Moreover, an A-check has ANORM available labor hours. Provided the cost of labor are set at CLAB euros per hour as derived in the previous Section, this would equate to LOSS2 euros of labor cost for an extra maintenance check. The combined cost of an extra check would equate

to LOSS1 plus LOSS2 euros. Converting these costs back to the equivalent excess hours equals over EH3 excess hours. From a modeling perspective, this would imply that the model would rather schedule over EH3 hours of excess in a regular maintenance check, before it would consider allocating an extra check. This would yield a very undesirable task schedule, in which large quantities of excess hours are scheduled in the regular maintenance checks. Moreover, such a maintenance planning would imply that one is scheduling for maintenance delays, which would not be considered a feasible maintenance planning. Considering only one of the two cost contributions, would still yield the same result. Hence, instead of using the same logic and cost contributions as for the excess penalty, the cost penalty for the extra maintenance checks has been based on the amount of excess hours which justifies the use of an extra maintenance check. The MRO's hangar policy stipulates that a heavy A-check contains a maximum of AHEAV required labor hours. Provided that the norm is set at ANORM hours, this is equivalent to XEH hours of extra labor required. Based on the formulation of the excess penalty, XEH hours of excess are equivalent to CEXTRA euros of extra labor cost. By setting the penalty factor for an extra check, C_y, equal to CEXTRA euros, more than XEH hours of excess on one of the regular letter checks justifies the need for an extra maintenance check. The reader is referred to Chapter 9 to read about the sensitivity analysis which was conducted on the penalty factor for extra maintenance checks.

5.3. Shortage Penalty

The shortage hour penalty is included to minimize unused available resource capacity. As aforementioned, the assumption is made that all available resource capacity may be used for scheduling of modification and OOP tasks. Thus, the model uses a greedy approach to optimally use all available resource capacity subject to the other constraints. In practice a similar greedy policy is used in task allocation by maintenance planners at the MRO, subject to feasibility constraints. As such maintenance planners have to consider precedence constraints between the tasks material and equipment availability, and so on. Nevertheless, a greedy scheduling approach is desirable as this implies that the available labor capacity is used most efficiently, which equates to less maintenance downtime and subsequently more aircraft availability.

$$\sum_a \sum_c C_{sh} \cdot SH_{a,c} \tag{5.3}$$

Hence, the shortage penalty endorses this greedy approach in the model. The shortage penalty term of the objective function of the model is shown again in Equation 5.3. The value of the shortage hour penalty factor, C_{sh}, is set to unity. As such it penalizes all shortage hours, $SH_{a,c}$, to ensure the efficient use of all available labor capacity, whilst also being a reference value relative to which all other penalty factors in the model are set.

5.4. Interval Penalty

As noted in Chapter 4 the model has two efficiency objectives: first the efficient use of resources and second the efficient use of task interval. The interval penalty ensures that the maintenance task interval is maximized. Ideally maintenance tasks are executed as closely as possible to their due date. Scheduling tasks early will result in a more frequent execution of the task, which ultimately will require more resource capacity and hence amount to higher cost. This same policy of maximizing the interval of maintenance tasks is applied by maintenance planners at the MRO, subject to feasibility and reasonability. In various cases it may be more worthwhile to sacrifice the interval for other benefits. For example, maintenance planners often schedule bigger maintenance items in the winter checks even when the tasks are not due till summer. This is done to "offload" the maintenance checks in the summer season, when the flight network is more demanding and maintenance delays potential for large disruptions. Although such scheduling reduces the risks of maintenance delays and network disruptions, it does imply a loss of usable interval and an earlier subsequent execution of these tasks.

$$\sum_a \sum_c \sum_t x_{a,c,t} \cdot C_{lb} \cdot mh_{a,t} \cdot \left(\frac{dd_{a,t} - cd_{a,c}}{int_{a,t}} \right) \tag{5.4}$$

The interval penalty as contained in the objective function of the model is shown again in Equation 5.4. The interval penalty consists of four terms; from left to right: the decision variable $x_{a,c,t}$, the labor cost C_{lb}, the workload of the task in question denoted by $mh_{a,t}$ and the fraction of unused task interval days over the total task interval expressed in days. It has been assumed that the percentage of unused task interval incurs the same percentage of the task's required labor hours as extra workload over time. This latter portion of a task's required labor hours are subsequently expressed in labor cost. Thus, if 10% of a task's interval remains unused after it has been allocated and the task requires 30 hours of labor in total, it is assumed that 3 additional labor hours will be required in time. Analogously, this would equate to a penalty of IP euros, provided that the labor cost are set to CLAB euros as derived in Section 5.1. This method has been inspired by the same method used by Coolen [53].

5.5. Additional Penalty Factors

In addition to the penalty terms which are explicitly part of the objective function, two other penalty factors have been included in the model: the seasonality penalty and check type penalty factors.

The seasonality penalty ensures that any excessive workloads of maintenance checks scheduled during the summer period are penalized more heavily compared to checks outside of the summer season. This penalty models a policy which is also employed by maintenance planners at the MRO. This was already briefly mentioned in an example in the previous Section. The higher flight density of the summer season requires a minimization of the risk of maintenance delays and potential disruptions. Hence, maintenance planners at the MRO will often allocate some bigger maintenance items much earlier on in the winter checks. The seasonality penalty factor aims to enforce this same principle in the model. In addition to being applied to the excess penalty of maintenance checks with excessive workloads in the summer period, it is also applied to the extra check penalty in case an extra check is allocated during the summer period. As the name suggest, it is a simple factor which amplifies the ordinary excess penalty or extra check penalty.
The MRO's 2018 summer period was defined as the period from March 27th till October 29th. This schedule has been adopted in the model and it has been assumed that this schedule does not change over the planning horizon.

The check type factor differentiates between the various check types and penalizes the excess hours on A-checks and extra checks more heavily than on the B- and C-checks. Compared to the B- and C-checks, the A-check is scheduled at a much higher frequency. Furthermore, the A-checks are a much "lighter" maintenance check compared with the B- and C-checks, both in terms of the both in terms of the quantity of tasks as well as the nature of the tasks to be performed. The heavier nature of these checks generally also implies that more buffer is built into the flight schedule following such a check. Provided that this is not the case for A-checks, and given their higher frequency, excessive workloads on A-checks are much more likely to cause disruptions to the flight network, than excessive workloads on the B- and C-checks. Furthermore, because of the extensive nature of the B- and C-checks, work scheduled to these checks can often profit from an efficiency gain. As a much larger number of access panels are opened during these checks compared with A-checks, this can often result in a time saving for a task which is added to these checks, which would not be experienced in an A-check. Provided that extra checks are scheduled in A-check slots, these considerations are also considered to apply to the extra checks. Thus, it is considered desirable to include these risks and benefits related to the check types into the model by adding the check type factor. As with the seasonality factor, it is a simple factor which amplifies the excess penalties of all A-checks and extra checks.

The penalty factors for both seasonality and check type are set at 1.25, increasing the excess penalty or extra check penalty by 25% if applicable. The values for these penalty factors was estimated using both engineering judgment and various iterations. Including the seasonality and check type factors in the original objective function introduced in Chapter 4, would result in the objective function shown in Equation 5.5, where C_{sea} denotes the seasonality penalty factor and C_{typ} denotes the check type penalty factor.

$$min. \sum_a \sum_c C_{sea} \cdot C_{type} \cdot EP_{a,c} + \sum_a \sum_c C_{sh} \cdot SH_{a,c} + \sum_a \sum_c \sum_t x_{a,c,t} \cdot C_{lb} \cdot mh_{a,t} \cdot \left(\frac{dd_{a,t} - cd_{a,c}}{int_{a,t}} \right) + \sum_a \sum_c C_{sea} \cdot C_y \cdot y_{a,c} \tag{5.5}$$

Moreover, the following notation would also need to be added to the mathematical model formulation to convey the applicability of the penalty factors:

Seasonality Penalty Factor

C_{sea} $\begin{cases} 1.25; \text{if start date of summer season} < cd_{a,c} < \text{end date of summer season} \\ 1.0; \text{otherwise} \end{cases}$

Check Type Penalty Factor

C_{typ} $\begin{cases} 1.25; \text{if check type is "A-check" or "SB-check"} \\ 1.0; \text{otherwise} \end{cases}$

6
Data Preparation Process

The preparation of the inputs for the model is also referred to as the 'Data Preparation Process'. Whereas a comprehensive flowchart and description of this process are provided in Appendix A, this Chapter will provide a detailed description of the computation and assumptions involved with the inputs of the task scheduling optimization model. Firstly, Sections 6.1 and 6.2 discuss the preparatory process related to the routine block and routine out-of-phase (OOP) tasks respectively. Subsequently, Sections 6.3 and 6.4 discuss the modification and non routine tasks respectively. Finally, Section 6.5 discusses the routine letter checks. The extra checks have already been discussed in Chapter 5.

6.1. Routine Block Tasks

As has been explained in Chapter 3, the routine block tasks refer to routine maintenance tasks which have been clustered together into blocks of tasks, which are allocated to a specific letter check. As mentioned in Chapter 3, for the airline's wide-body fleet there are 24 unique A-blocks, 2 unique B-blocks and 15 unique C-blocks [8].

It is assumed that all routine block tasks are executed in their designated letter checks. This assumption will be explained in greater detail in Section 6.5. Although these block tasks are not allocated by the model, their required capacity affects the available capacity of the maintenance checks to which the modification and OOP tasks can be allocated. Subsequently, this Section briefly describes the computation of the labor hours and due dates of the block tasks. It is important to note that the notation "maintenance task" refers to an MRI.

6.1.1. Computation of Required Labor Hours of Block Tasks

For all aircraft, the total required resource capacity is determined for each of the routine blocks within the planning horizon. The total required capacity of an MRI is the sum of the labor hours of the various JICs it is composed of. However, provided that the blocks consists of numerous MRIs, there are some MRIs which share the same JICs. Hence, a set of unique JICs is compiled for each routine block of each aircraft. The hours of this unique list of JICs are the required scheduled labor hours to perform the MRIs in the block.

In addition to the hours required to perform the actual maintenance work, the access panel hours also need to be added. The access panel hours are the hours required to access the systems and components of the aircraft on which the maintenance is performed. Similarly to the JICs, a unique list of access panels is compiled for the block in question, of which the corresponding times are summed. The total required resource capacity of the routine blocks is the sum of the JIC hours, access panel hours and the corresponding non routine hours. This last category will be discussed in greater detail in Section 6.4.

6.1.2. Computation of Block Task Due Dates

The due dates of the routine block tasks are based on the previous execution date of each task, the task intervals and the utilization rate of the aircraft. For each task in a block the previous execution date is determined from historical records. The interval of the tasks may be expressed in flight hours, flight cycles and calendar time (days, months, years). Based on the utilization of the aircraft, the flight hour and flight cycle intervals are converted to interval time in days. The due date of each task is subsequently determined by adding the intervals of each task (in days) to the previous execution date. Provided that a task may have multiple intervals specified, the due date of the task corresponds to the earliest due date based on the intervals. The following assumptions are included in the due date computation:

1. If a task has no previous execution date (it will be executed for the first time on the aircraft in question), the manufacturing date of the aircraft is used as the previous execution date
2. The utilization of the aircraft is assumed to be constant over the planning horizon

The first assumption is purely practical. The algorithm involved with the computation of the due dates has been written such that it requires a previous execution date in order to determine the upcoming due date of a task. The manufacturing date denotes the day on which the aircraft is delivered from the aircraft manufacturer to the operator. Although in practice the delivery date and start date of actual operations may vary slightly from the manufacturing date, these deviations are sufficiently small to ignore their effect on the due date of the task.

The utilization of the aircraft is assumed to be constant throughout the year and throughout the planning horizon. In practice, for the due dates computed in the MRO's MIS a distinction is made between the summer and winter season utilization. These values are listed in Table 6.1.

Table 6.1: Overview of the airline's wide-body fleet utilization rates according to the MRO's MIS

Date	Rate	Unit
January 1st	WIFC	cycles
January 1st	WIFH	hours
March 27th	SUFC	cycles
March 27th	SUFH	hours
October 29th	WIFC	cycles
October 29th	WIFH	hours

Due to time constraints, the model currently assumes a fixed utilization all year round, and does not account for seasonality as is done in practice. Moreover, where the utilization is updated each year in the MRO's MIS based on the recorded actual utilization, the model assumes a constant utilization throughout the planning horizon. To be conservative, the slightly higher utilization for the summer period has been adopted in the model. For now this assumption is considered valid as the error introduced by the assumption is sufficiently small. A simple analysis was performed to evaluate these errors. The tasks with the smallest and largest intervals were selected for the analysis, with 500 and 3000 flight cycles; and 1500 and 16000 flight hours respectively. Furthermore, both the summer and winter utilization values were used for the comparison in the analysis. Table 6.2 shows the results of the analysis. The 'Delta Days' column shows the difference between the interval days based on the summer utilization and the winter utilization respectively. Thus the 'error' shown in Table 6.2 is an upper bound, since the actual error with respect to the combined winter and summer utilization interval will be smaller.

Table 6.2: Overview of the interval days according to the winter and summer utilization rates for the maintenance tasks with the lowest and highest number of flight cycles (FC) and flight hours (FH) respectively

Interval	Winter Utilization (days)	Summer Utilization (days)	Delta Days
500 FC	WUC1	SUC1	DIF1
3000 FC	WUC2	SUC2	DIF2
1500 FH	WUH1	SUH1	DIF3
16000 FH	WUH2	SUH2	DIF4

As can be seen from Table 6.2, the absolute differences vary between DIF3 days and DIF2 days. Using engineering judgment, DIF3 days out of SUH1 and WUH1 days as well as DIF2 days out of WUC2 and SUC2 days are considered sufficiently small errors. Thus, the fixed utilization rate is considered acceptable for the present model.

6.2. Routine OOP Tasks

As explained in Chapter 3, the out-of-phase (OOP) tasks are routine maintenance tasks contained in the AMP. However, unlike the block tasks, these tasks are not clustered together with other tasks in blocks and subsequently not fixed to a specific letter check. Hence, these tasks are scheduled by the optimization model. It is important to note that the model only includes A-check type OOP tasks. C-check type OOP tasks are not included in the model. The following criteria apply to A-check type maintenance tasks in general [8]:

- Minimum number of Flight Cycles: 200
- Maximum number of Flight Cycles: 2000
- Minimum number of Flight Hours: 1500
- Maximum number of Flight Hours: 18000
- Minimum number of calender days: 105
- Maximum number of calender months: 36
- Maximum number of calender years: 3

If a maintenance task complies to these requirements, it is considered an A-check type task, classes either as a block task or an OOP task. Provided that the routine block tasks are predefined and known, any maintenance task in the AMP which complies to these criteria and which is not part of the A-blocks, is classified as an OOP task.

It has been assumed that the workload of each OOP MRI consists of the labor hours of all its JICs, the corresponding access panel hours and an estimated non routine workload. It may occur that an OOP task is allocated to a check of which the routine block task share some of the same JICs or panels as the OOP task in question. Consequently, this assumption would result in duplicate JIC or access panel hours and hence in a slightly higher required resource capacity for this respective check. These errors are considered small and conservative and as such no further action was taken to correct for this in the present model.
The initial due dates of the OOP tasks are provided directly from MIS data. These initial due dates are used during the initial optimization run of the task scheduling model. As has been explained in Section 4.3 of Chapter 4, the iterative solution technique built around the model computes the due date of the next occurrence of each OOP task, which are used as inputs for each subsequent optimization run.

6.3. Modification Tasks

As explained in Chapter 3 modification tasks are generally non-recurring tasks which are issued as product improvements or for continued airworthiness by the regulatory authorities, manufacturers (OAMs), or at the request of the operator. As such, there is a continuous inflow of new modification tasks. Contrary to the routine tasks, modifications are also not restricted to any particular maintenance check and may have due dates which lie relatively far into the future (in the order of months and even years). As such, the modification tasks are allocated by the optimization model to the various maintenance checks over the planning horizon.

As explained in Chapter 3 the inputs for the modifications are taken from the MRO's dedicated Data Exchange Platform. All modifications classed as "Closed" or in the "Evaluation Phase" are excluded from the model. The "Closed" phase marks that a modification has been executed on all applicable aircraft and is therefore completed. The "Evaluation Phase" marks that a decision is yet to be made regarding the execution of the modification. As it is yet unclear in this phase whether or not the modification will be executed, these modifications are also excluded from the task planning model. All other modifications in the EO execute process are considered relevant and included in the model.

Following the selection of modifications from the Data Exchange Platform, the execution status of each modification (how much of each modification has been completed on each aircraft) is updated using MIS information. Provided that the Data Exchange Platform only tracks modifications up until the first execution, this aircraft specific information is not contained in the Data Exchange Platform. Provided that a modification may consist of subtasks which are separately executable, the status of a modification may range from fully active to partially completed to fully completed.

The primary reason for using the Data Exchange Platform as the source for modifications instead of the modification data contained in the MRO's MIS, is the more complete overview of modifications contained in the Data Exchange Platform. As explained, the Data Exchange Platform is used to by the MRO to assist in the preparation of all modifications up until the first execution. Thus, it also contains modifications which are not yet accounted for in the MIS, but for which the decision has yet been made that the modification will be executed. An analysis on a sample of modifications dated October 2018 showed that of the 574 modifications, 148 were not yet registered in the MRO's MIS at that time. This equates to 25% of the total number of modifications. The more complete overview of modifications in the Data Exchange Platform was therefore selected as the source of the modifications, whilst the MIS information is used to confirm the current execution status.

Several assumptions have been made with regards to the modifications:

1. The preferred check type of modifications is not considered in the allocation of tasks by the model.
2. If a modification has no preferred check type, it is assumed to have a preference for C-checks.
3. If a modification has no due date, a due date is assumed according to the preferred check type.
4. If a modification has no specified required labor hours, this is assumed according to the preferred check type.
5. All active subtasks of a modification are scheduled at once rather than separately.
6. Modifications with a preference for P-checks are ignored as these are assumed to be executed at line maintenance.
7. The interval of a modification is assumed to be the time between the 'creation date' in the Data Exchange Platform and the due date of the modification.

Each modification is designated a preferred check type indicating which check type would suit best with the modification. For example, software updates will generally have a preference for line maintenance slots as these updates usually require very little elapse time. Although in practice this information is useful in allocating a modification, it is not strictly followed. Currently, if a modification requires more than XH labor hours, it is assigned a preference for C-checks. This results in a large group restricted to the C-checks. Ordinarily a maintenance planner would see through this and allocate the task accordingly to a more favorable check. To enable the model to schedule optimally for resource capacity, the preferred check type is ignored in the allocation of tasks by the model. However, this will potentially introduce errors into the model, since the check preference may be defined based task conditions other than the required capacity, e.g. power-off conditions, large elapse times, etc. These conditions are currently not included in the scope of the model. Nevertheless, it is important to be aware of the implications of this assumption and the restricted scope of the model.

Although the preferred check type is not used for task allocation, it is used to make assumptions on the required labor capacity and due dates of modifications for which this information is yet to be determined. If the preferred check type itself is missing, a preference for C-checks is assumed. Generally, if the preferred check type is yet unknown, the due date and required labor capacity are also unknown. Provided that the preferred check type is used to estimate these parameters, assuming a preference for C-checks results in a conservative estimation of these two parameters as will be shown in the following.

When the due dates are missing, the due dates are estimated based on the preferred check type. For modifications with preferred check types 'A' or 'H' or 'P', the due date is scheduled TA months from the date it was created in the Data Exchange Platform. The modifications are 'created' in the Data Exchange Platform at the start of the EO process. For modifications with preferred check types 'B' or 'C' or 'Special', the due dates are set to TC months from the creation date. These intervals are conservative as they exceed the actual check intervals. As such, they allow for several check date options to allocate the modification task to, which provides scheduling flexibility. The intervals have been set after consulting with employees of the MRO's Planning, Scheduling and Fleet Control division.

Similar to the missing due dates, the labor hours of modifications for which this information has not yet been recorded in the Data Exchange Platform are based on the preferred check types. For the modifications with preferred check types 'A' or 'H' or 'P' the required labor hours are set to HAMOD hours. For the modifications with preferred checks 'B', 'C' or 'Special', the required labor hours are assumed to be HCMOD hours. These hours have been based on the medians of the labor hours of a sample set of modifications.

The sample set considered consisted of a collection of 'active' modifications and dates back to August 2018. The sample contained XA modifications with preferred checks A, H or P and XC with preferred check types B, C or Special. Figures 6.1 and 6.2 show the boxplots and histograms of the sample set. Due to the skewed distribution of the modification hours as shown in these Figures, the median was assumed to be a more accurate representation of the required hours compared to the average values.
Of the sample set corresponding to preferred check type A, H or P the largest labor hours are AH1, AH2, AH3 and AH4 hours respectively. For the type B, C or special the largest hours in the sample set are CH1, CH2, CH3 and CH4 hours respectively. Removing these 'outliers' from the analyses would have resulted in a median of HCMOD2 for the C-type modifications, which is slightly smaller than the median of HCMOD hours. The median of the A-type would remain the same. With 4 out of the XC C-type modifications classed as outliers, the probability that such outliers occur is considered relatively small. However, it is important to be aware of the possibility of such outliers and that the medians are only an approximation.

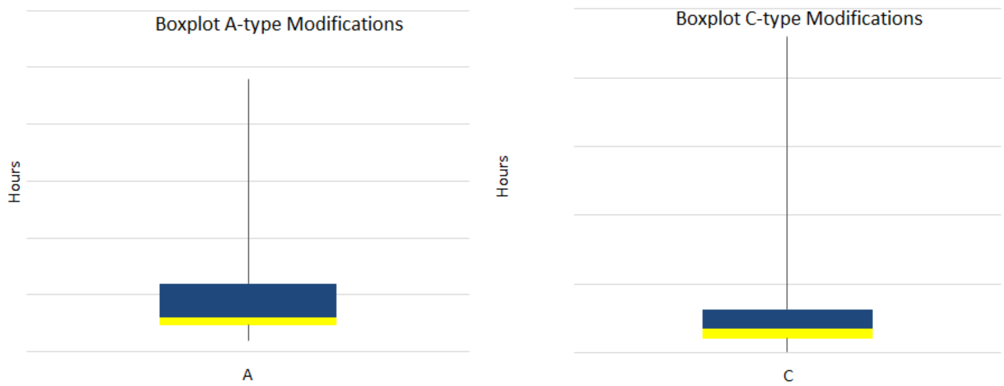

Figure 6.1: Boxplots of the scheduled hours of a sample of modifications; left: preferred check types A, H or P; right: preferred check types B, C, Special

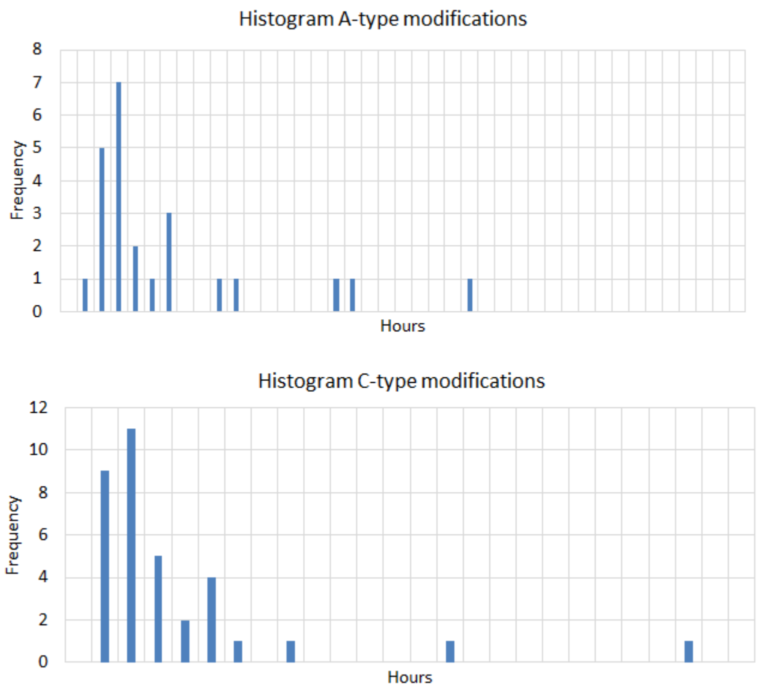

Figure 6.2: Histograms of the scheduled hours of a sample of modifications; top: preferred check types A, H or P; bottom: preferred check types B, C, Special

It has been assumed that all subtasks of a modification are executed at the same time in the model. As explained, a modification may consist of various subtasks. In practice, such subtasks allow for the modification to be executed in smaller parts, making it less disruptive to the routine maintenance work. Moreover, it may be necessary to schedule the modification in its subtasks because of limited component or kit availability, or because of other restrictions such as ETOPS regulations. The present model uses a simplified approach which assumes that all subtasks are scheduled at once. However, it should be noted that the model does take into consideration which subtasks have yet been completed. The subtasks which have been completed are no longer considered by the model. Rather, it scales the required labor hours of the modification by the fraction of the number of remaining active subtasks over the total number of subtasks. It has been assumed that the required labor hours scale linearly with the number of subtasks of the modification.

It has been assumed that the modifications with a preference for P-checks may be ignored. These modifications generally refer to software updates or other small maintenance tasks which are generally scheduled to line maintenance checks, rather than the base maintenance checks considered by the model.

A task interval has been assumed for the modifications. Modifications do not have an interval like the routine maintenance tasks, as they are generally non-recurrent. However, as explained in Section 5.4, the interval penalty in the objective function of the model is scaled by the interval of the maintenance tasks and therefore requires all tasks to have an interval. The interval of a modification has been assumed to be the number of days between its due date and the date it was first registered in the Data Exchange Platform.

6.4. Non Routine Tasks

For the routine block and OOP tasks the model takes into consideration an estimated non routine workload. The non routine tasks are the corrective maintenance tasks resulting from inspections during routine maintenance tasks. As such, their occurrence is stochastic and therefore difficult to anticipate. The MRO has developed an algorithm to predict the non routine tasks that are likely to occur given a workpackage of routine maintenance tasks. This algorithm is referred to as the Non Routine Predictor (NRP). Provided that this algorithm has already been developed and is currently undergoing further developed within the MRO, the NRP algorithm is used by the present model and improvements to the algorithm are considered outside of the scope of the present model. Figure 6.3 gives a schematic overview of the workflow of the NRP.

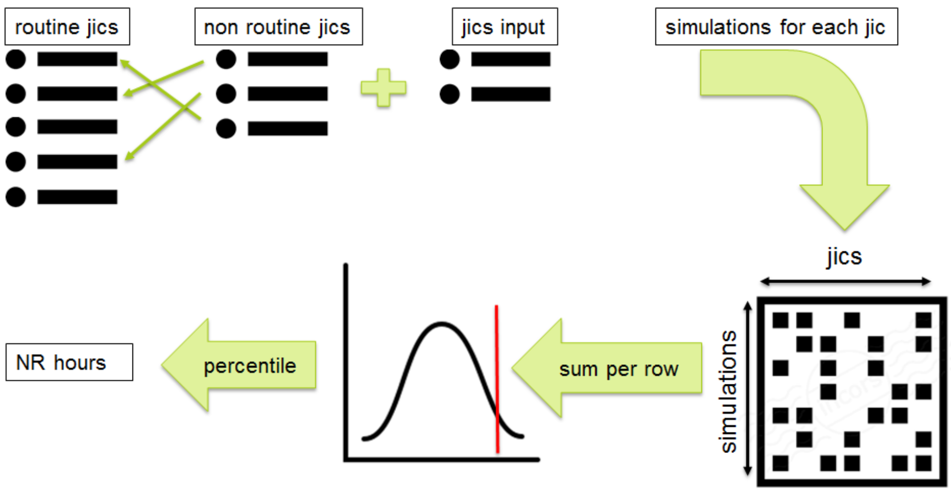

Figure 6.3: Schematic overview of the non routine predictor algorithm [5] [6] [7]

The NRP algorithm uses historical records of all maintenance tasks (routine and non routine) for the entire airline's wide-body fleet, and correlates the occurrence of a non routine maintenance task to the routine maintenance task which induced it. Provided an input of routine JICs as defined by the user, it subsequently runs one thousand Monte Carlo simulations for each of the input JICs. During each simulation it uses a random number generator to select one of the occurrences of the JIC under consideration from the historical records. If the selected JIC occurrence is related to a non routine task, it subsequently adopts the hours required for that non routine task to a matrix. If there are no non routine tasks related to the randomly selected JIC occurrence, it adopts zero to the matrix. It repeats this one thousand times for all JICs. After these simulations, it sums the non routine hours of each row in the matrix. The sum of each row represent the total non routine hours of all JICs under consideration for a specific simulation. By doing this for each row, a distribution of one thousand summed non routine hours is obtained for the routine tasks under consideration. Provided a user desired percentile, a single value is obtained for the expected non routine hours. Thus, the result of the NRP is a statement that in x percent of the cases of executing the JICs in question, it is expected to have at most y hours of non routine tasks. Here x is the user-defined percentile and y the corresponding non routine hours. This setup allows for the user to decide on the probability (and hence the risk) involved with the number of non routine hours accounted for in a maintenance check; a lower percentile implies more capacity available to schedule other tasks, but a lower probability that the non routine hours planned for will match with the actual non routine hours encountered during the check.

An example of a "distribution" of non routine hours for the A01 block is shown in Figure 6.4. The histograms has bins of XBIN hours. The frequencies on the y-axis indicate how often the total non routine hours of a simulation fell in each of the respective bins. The distribution is largely clustered between HBIN1 and HBIN2 hours and appears to peak at HBIN3 hours. A few smaller frequencies occur around HBIN4 hours, with a final single occurrence at HBIN5 hours which is not visible in the Figure.

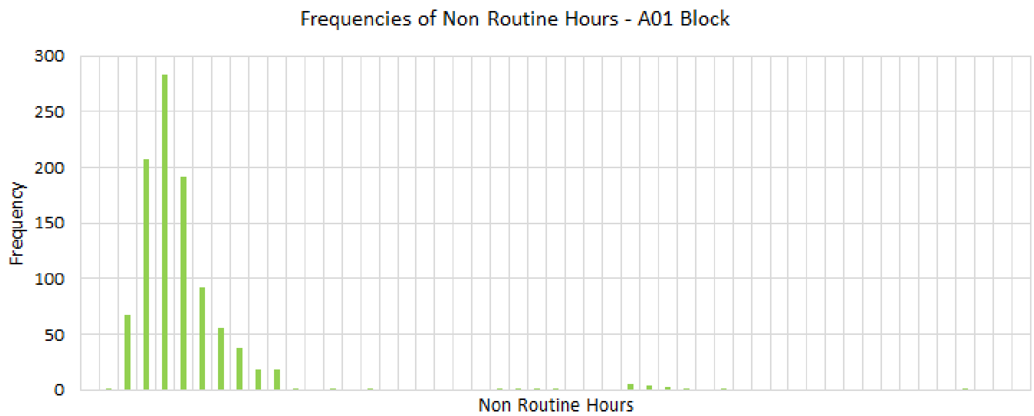

Figure 6.4: Histogram of the non routine hours of the A01 block based on 1000 Monte Carlo simulations

The present NRP algorithm provides an estimate of what the non routine hours could be for a given set of input JICs. As such, it is used to estimate the non routine hours of the routine block and OOP tasks in the model. It has been assumed that the modifications do not lead to any corrective maintenance tasks. Thus, for the routine block hours and routine OOP tasks, the corresponding JICs are fed to the NRP to estimate the non routine hours of each block and OOP task. These non routine hours are subsequently added to the required resource capacity of the block and OOP tasks respectively. The combined total of the scheduled resource capacity and unscheduled or non routine resource capacity is assumed to be the total required resource capacity of the blocks and of the OOP tasks respectively.

The present NRP algorithm has several limitations. Firstly, the historical data is not filtered. There is a substantial portion of records of which the reported required labor hours for the non routine tasks are faulty. The required non routine hours of corrective maintenance tasks are recorded by the maintenance personnel during a maintenance check. These hours are not always recorded properly. Without filters on the historical data records, these erroneous hours are also considered in the estimation of the non routine hours which degrades the quality of the NRP estimate.

Moreover, the current simulation method adopts the hours of the non routines directly from the occurrences. An alternative method would be to make a probability distribution of the non routine hours per JIC, based on the empirical data and simulations. Subsequenly, this distribution could be used to predict the non routine hours per JIC; the total non routine hours of the input JICs would simply be the sum of the non routine hours at a given percentile of each JIC. The present method is limited to the empirical data only.

As aforementioned, the NRP is currently undergoing further development within the MRO. The present version of the NRP used in the model represent an initial trial version. Nevertheless, it is important to beware of the inaccuracy of the NRP algorithm in its current state.
Lastly, it is important to note that the present model only adopts the non routine hours as predicted by the NRP. It does not include specific non routine tasks.

6.5. Letter Checks

Unlike the extra maintenance checks discussed in Chapter 5, the letter checks are routine checks and as such not scheduled by the optimization model. Rather, the dates and resource availability of these maintenance checks are determined for each aircraft prior to the optimization and are subsequently fed as inputs to the optimization model. As denoted by constraints 4.9 in Chapter 4, the binary decision variables $y_{a,c}$ for these regular maintenance checks are set to one. This ensures that the regular maintenance checks are always available for task scheduling. The following two Sections describe how the check dates and the resource availability of the regular maintenance checks have been determined.

6.5.1. Computation of Letter Check Dates

Section 6.1 has explained how the due dates of the routine block tasks are determined based on the previous execution dates, aircraft utilization and the task intervals. After the due dates of all routine block tasks are determined, the check date of the respective letter check is set equal to the earliest due date of the tasks in its block. This process is repeated for all maintenance checks for which the check dates fall within the planning horizon. The following assumptions are included in the check date computation:

1. The check date of a letter check is set equal to the earliest due date of the tasks in the block corresponding to that letter check
2. Check dates of different aircraft may overlap
3. C-check dates are set based on the check intervals, rather than task intervals

It is assumed that the letter check date is equal to the earliest due date of the routine block tasks. In practice the due dates of the block tasks are indeed the primary driver for the maintenance check dates. However, presently maintenance ground time planners take into consideration a small scheduling margin to avoid exceeding the task due dates. Provided the dynamics of aircraft operations, there are various factors which can prevent an aircraft from coming in for maintenance on the scheduled date. Hence, this margin allows for some scheduling flexibility in these scenarios. This has been neglected in the present model.

The second assumption allows for aircraft checks to be scheduled on the same date. Although the hangar has sufficient capacity to physically accommodate two wide-body aircraft, in practice this would not be desirable. However, the present setup of the algorithm may result in checks with the same check date. The algorithm involved with computing the check dates loops over each of the aircraft consecutively. This sequential method may result in two different aircraft with a coinciding check date. To avoid the model from scheduling maintenance checks on the same dates, an optimization algorithm would be required which aims to find the optimal maintenance slot for all aircraft in the fleet. A similar optimization model for A-check allocation has already been developed and is already being used at the MRO. However, provided that this model is limited to an operational planning horizon and A-checks only, its scope is not sufficient enough to be used as an input in the present model. Hence, the model uses the crude approximation method as has been explained. Ideally, an extension of these existing models would enable it to provide an accurate and optimal maintenance check allocation to be used by the present model. However, such extensions were considered out of the scope of the present model and research.

The last assumption implies that the due date calculation method explained in this Section applies to the A- and B- checks but not for the C-checks. Provided that the C-check routine tasks are excluded from the scope of the present research, the C-check dates cannot be estimated according to the same task-based approximation method used for the A- and B-checks. From a study conducted by the MRO in anticipation of the airline's first wide-body aircraft, the Engineering Department of the MRO derived the following C-check intervals for the airline's wide-body fleet [8]:

- 18,000 Flight Hours
- 2,000 Flight Cycles
- 3 years (1095 Days)

Based on these intervals, the C-check blocks have been designed. The due date of the initial C-checks are computed using these intervals and the manufacturing date. All subsequent C-checks are scheduled with respect to the previous C-check date and these intervals.

6.5.2. Computation of Letter Check Available Capacity

As mentioned in Chapter 4, the routine block tasks and corresponding non routine tasks are assumed fixed to their respective letter checks. Thus, the required capacity of the block tasks affects the available resource capacity of the letter checks to which the modification and OOP tasks may be allocated by the model. This remaining available capacity is computed by subtracting the workload of the routine block tasks and non routine tasks from the norm of the letter check. The norm is a predefined standard for the available resource capacity per check type. The following assumptions are made in the computation of the resource capacity:

1. It is assumed that the total available resource capacity is exactly equal to the norm for all maintenance check, at all times.
2. The routine block tasks and corresponding non routine tasks are all executed in the designated maintenance check, there are no task extensions.
3. A small portion of the available resource capacity of all maintenance checks is reserved for line maintenance tasks, referred to as the P-package.
4. It is assumed that all of the remaining available resource capacity may be used to schedule modifications and OOP tasks.
5. The available resource capacity of the C-check is estimated to be CNORM hours.

The following norms have been used in the computation of the resource availability of the maintenance check types. These norms have been checked with officials from the Planning Scheduling and Fleet Control division of the MRO:

- A-checks: ANORM available hours
- B01-checks: B1NORM available hours
- B02-checks: B2NORM available hours
- C-checks: CNORM available hours
- P-package: PNORM required hours

It is assumed that each check has the resource capacity available prescribed by the norm of the corresponding check type. The norm represents an agreed upon labor capacity which the hangar can cater for during a certain check type. Subsequently, the norm is used by maintenance planners and hangar staff to schedule maintenance tasks. Although in practice deviations from the norm are common, such variations are minimized. The purpose of the norm is to ensure that the workloads of maintenance checks do not exceed the labor capacity available. This in turn ensures a steady operation of maintenance checks, such that maintenance causes minimal disruption to the flight operations of the airline. However, in practice shortage of personnel (sickness), and last minute changes in the workload present difficult challenges. In the model however, it is assumed that the available capacity of the maintenance checks always equals the norm. It does not take incidental changes in available maintenance personnel into consideration. This assumptions is justified by the fact that the maintenance planners also use the norms when compiling the workpackages of upcoming maintenance checks.

It is assumed that the routine block tasks and their corresponding non routine tasks are executed in their respective letter check. In practice it is possible to extend the due date of due maintenance tasks within the limits set by the airworthiness regulations. The extended maintenance task must be completed within a predefined extension interval. Such extensions may be useful in case materials required for replacement or repair are not available at the present time and have substantial lead times that would result in severe delays. However, extensions of routine maintenance tasks, also referred to as drop-out items, are not ideal. A drop-out item will require the aircraft to come in for maintenance a second time within a relatively short time window following the original maintenance check. If a drop-out item does occur, the policy is to schedule the next execution of the maintenance tasks with respect to the original check date rather than the actual execution date of the drop-out item. This prevents block tasks from permanently running out of sync with the rest of the block. Non routine tasks which are extended are referred to as deferred defects. Deferred defects are more common, provided that the defect does not affect the airworthiness of the aircraft. Despite the possibility of extending due maintenance tasks, the MRO's policy is never to plan for such an extension. Thus, the model assumes all routine block and corresponding non routine maintenance tasks to be completed in the designated letter check.

In addition to the routine block and non routine maintenance task hours, a so-called P-package, "platform package", is scheduled in each maintenance check. This P-package is reserved for line maintenance tasks. The current MRO's policy prescribes for at most PNORM hours to be reserved during letter checks. Provided that the aircraft comes in for routine base maintenance anyway, some of the smaller maintenance items ordinarily performed by line maintenance can be scheduled to this P-package. As explained in Chapter 3, these tasks generally concern fault corrections and as such are generally scheduled for only a few days prior to the maintenance check. Provided the substantially longer planning horizon considered in the model, these tasks cannot be taken into consideration by the present model. Thus, the model only takes into consideration the agreed upon standard of PNORM hours.

The available resource capacity of a regular (letter) maintenance check equals the difference between the norm of the check and the total hours of all block tasks, all non routine tasks and the P-package hours. It is assumed that all remaining available resource capacity may be used for allocation of OOP tasks and modification tasks. In practice, it may not always be desirable to match the required resource capacity exactly to the available capacity as this also increases the risk of delays. However, the general policy of maintenance planners at the MRO is also to deliver the aircraft as 'clean' as possible, meaning that the planners aim to schedule as much work as possible to the various maintenance checks, within the limits of feasibility. Given that the purpose of this model is to provide insight into the available resource capacity compared with the required resource capacity, this greedy scheduling approach is considered valid.

As aforementioned, the routine C-check tasks are not included in the present scope of the model. As such, the remaining available capacity of the C-checks also has to be estimated. It has been assumed that the resource capacity of the C-checks, available for OOP and modification tasks, equals CNORM hours. This estimate is equal to the resource capacity which was dedicated for modifications in the C-checks of the airline's first three wide-body aircraft scheduled near the end of 2018. Moreover, it is assumed that the C-check resource availability excludes the aforementioned P-package hours.

II

Results

7
Introduction of the Results

This Chapter introduces the results of the task scheduling model. The results of the model are the allocated workload of all maintenance tasks over the maintenance checks within the planning horizon for the airline's wide-body fleet. The first Section of this Chapter provides a visual representation of the results and highlights the various observations. Section 7.2 provides a more detailed analysis of the most prominent observations of the results. Subsequently, Section 7.3 briefly discusses the model performance. Lastly, Section 7.4 provides a discussion on the potential value of the task scheduling model results for the MRO.
During the analysis of the results an anomaly was detected. This anomaly was removed and the model was run again. As such, all results presented in this Chapter are the corrected results only. A summary of the anomaly, its effect on the results and model performance, and the correction of the anomaly are briefly discussed in Section 7.5.

Following the introduction of the results provided in this Chapter, the following two Chapters will discuss the model validation and sensitivity analysis respectively.

7.1. Task Scheduling Results

This Section will provide a visual and numerical representation of the results of the task scheduling model and will highlight the most prominent observations. First, in Subsection 7.1.1 an overview of the complete workload distribution for a single aircraft will be presented. Following this, in Subsection 7.1.2 the workload distributions of the entire fleet will be presented per check type. Lastly, Subsection 7.1.3 provides a brief general overview of some additional model results.

7.1.1. Workload Distribution - Single Aircraft

A visualization of the task scheduling results for Aircraft 1 is shown in Figure 7.1. The Figure shows the workload of the aircraft for a planning horizon of four years, starting with the first check (A11) on the 11th of June 2018 and finishing with the last check (A04) on the 19th of August 2022. The bar segments represent the required capacity of the various task types which have been allocated to each maintenance check. From top to bottom, the bar segments represent the modifications; OOP tasks; non routine hours; routine block hours; P-package hours; and the excess hours on the negative y-axis. The dates and names of the various maintenance checks are shown on the x-axis. The extra checks are denoted by the notation 'SB' and a corresponding check number.
Table 7.1 provides an overview of various output parameters of the workload distribution of Aircraft 1. Please note the following in relation to the parameters in Table 7.2:

- The percentage of checks with excess hours and the percentage of checks with more than XEH hours of excess are both relative to the total number of checks.
- The average excess hours are based only on those checks that have excess.

- The average deviation refers to the deviation from the capacity norm, which may either be positive (shortage scheduled) or negative (excess scheduled). Thus, the average deviation has been computed based on these non absolute values.
- The maximum deviation refers to the maximum shortage hours, or unused capacity hours.
- The minimum deviation equals the maximum excess hours, see 'max. excess hours'.
- All non-integer values have been rounded to one decimal.

Table 7.1: Overview of various output parameters for the workload distribution of Aircraft 1 over the four year planning horizon

Parameter	**Value**
total checks	28
num. checks with excess	SNEX
percentage of checks with excess	SEPER
num. checks with >XEH hrs excess	SNXEH
percentage of check with >XEH hrs excess	SEPERX
max. excess hours	SMEH
avg. excess hours	SAEH
sum of excess hours first half of planning horizon	SUM1
sum of excess hours second half of planning horizon	SUM2
avg. deviation	SADEV
max. deviation	SMDEV
total OOP hours	STOOP
total modification hours	STMOD

Of the 28 checks which have been allocated to Aircraft 1, 18 check are A-checks; 3 Cabin checks; 2 C-checks; and 5 extra checks. Of these 28 checks, SNEX checks have excess hours. Moreover, of the SNEX checks with excess hours, SNXEH checks have more than XEH hours of excess scheduled, equivalent to almost two-thirds of all checks. It is apparent that the excess hours do not exceed more than SMEH hours. Of the SNEX checks with excess hours, the average excess hours are SAEH hours. It is important to note that this includes all different check types. The analysis per check type of the entire fleet of aircraft will be provided in the Subsection hereafter.

It is apparent from Figure 7.1 that the maintenance checks in the second half of the planning horizon have more excess hours compared with the first half of the planning horizon. If the A19 on the 31st of May 2020 is considered as the halfway point of the planning horizon, the total excess hours scheduled in the first half of the planning horizon are SUM1 hours compared to SUM2 hours in the second half of the planning horizon.
As can be also seen from Figure 7.1, the five extra checks are all scheduled in the first half of the planning horizon, with the last extra check scheduled on the 1st of January 2020. Two out of the five extra checks have excess hours scheduled. One of the extra checks is scheduled at exactly full capacity (ANORM hours). On the other hand, there is one extra check with only SFEW hours allocated to it. As such it has SMUCH hours of unused capacity. The maximum of unused capacity occurs for the second C-check of Aircraft 1, it has SMDEV hours of available capacity remaining. On average the checks of Aircraft 1 have SADEV hours of available capacity remaining, as indicated by in Table 7.1.
A total of STMOD hours of modification tasks and STOOP hours of OOP tasks have been allocated by the model.

7.1.2. Workload Distribution - Fleet

Table 7.2 provides an overview of various output parameters of the model. Furthermore, a visual representation of the workload distribution per check type is provided by Figures 7.2 to 7.5. Please note that in all Figures, the checks are in chronological order on the x-axis and the allocated labor hours of the various task types on the y-axis. The initial observations from Table 7.2 and Figures 7.2 to 7.5 will be briefly discussed for each check type.

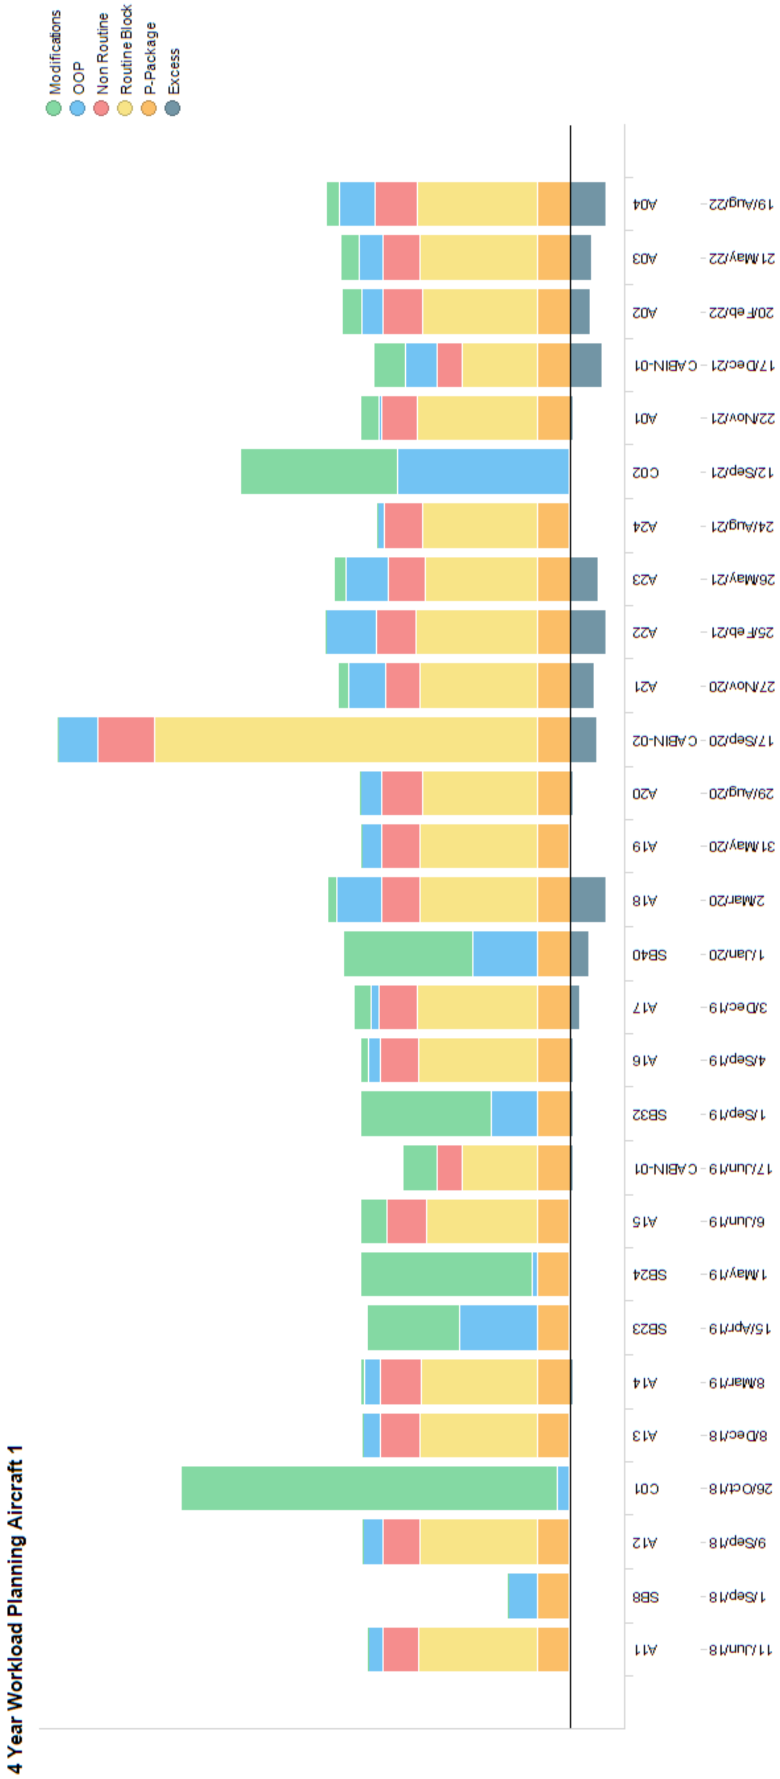

Figure 7.1: Visualization of the workload of Aircraft 1 for the four year planning horizon. From top to bottom the bar segments represent the allocated workload per task type of the modifications, OOP tasks, non routine tasks, routine block tasks, P-package hours and the excess hours (negative)

Table 7.2: Overview of some of the task scheduling optimization model output parameters for the airline's wide-body fleet over the four year planning horizon

	A	**B**	**C**	**SB**	**Total**
number of checks	TAC	TBC	TCC	TSBC	TC
num. checks with excess	NEA	NEB	NEC	NESB	NET
percentage of checks with excess	EXPERA	EXPERB	EXPERC	EXPERSB	EXPERT
num. checks with >XEH hrs excess	NEHXA	NEHXB	NEHXC	NEHXSB	NEXHT
percentage of checks with >XEH hrs excess	PEHXA	PEHXB	PEHXC	PEHXSB	PEHXT
max. excess hours	MEA	MEB	MEC	MESB	MET
avg. excess hours	AEA	AEB	AEC	AESB	AET
avg. deviation	-ADA	-ADB	ADC	ADSB	ADT
max. deviation	MDA	MDB	MDC	MDSB	MDT
avg. OOP hours	OOPA	OOPB	OOPC	OOPSB	OOPT
percentage of allocated OOP hours of total workload	PROOPA	PROOPB	PROOPC	PROOPSB	-
avg. modification hours	MODA	MODB	MODC	MODSB	MODT
percentage of allocated modification hours of total workload	PRMODA	PRMODB	PRMODC	PRMODSB	-

A-Check Distribution

The workload of the A-checks for the airline's wide-body fleet over the four year planning horizon is displayed in Figure 7.2. It is apparent that over the first 1/3 of the planning horizon the excess hours of the A-checks are substantially lower than in the two-thirds that follow. A similar observation was made for the integral workload distribution of the single aircraft in 7.1.1.

As can be seen from Figure 7.2, the workloads of the routine A-block tasks and corresponding non routine tasks are relatively constant over the entire fleet and planning horizon. As noted in Table 7.2, EXPERA percent of all A-checks have excess hours. Moreover, PEHXA percent of all A-checks have more than XEH hours of excess. A maximum of MEA excess hours are scheduled. As was observed for the single aircraft, the maximum excess hours appear to be just above EH1 hours. Of the A-checks with excess hours, the average excess is AEA hours. Moreover, the average deviation of all A-checks from the capacity norm is -ADA hours, where the negative sign indicates excess rather than shortage hours. Thus, on average all A-checks exceed the norm by ADA hours. The norm for the A-checks, of ANORM hours, as implemented in the model is also visualized in Figure 7.2. The maximum deviation from the norm is MDA hours, which represent unused labor capacity. On average OOPA hours of OOP tasks and MODA hours of modifications are scheduled to A-checks respectively. Overall, the sum of all OOP hours scheduled to A-checks equates to PROOPA% of the total OOP task workload. Similarly, approximately PRMODA% of the total workload of modifications is scheduled to the A-checks.

B-Check Distribution

The workload distribution of the B-checks, or cabin checks, is shown in Figure 7.3. Both the Cabin-01 and Cabin-02 checks are shown. Please note the different norms for the two Cabin checks, indicated by the horizontal lines in Figure 7.3.

Of the TBC cabin checks, NEB checks have excess hours, of which NEHXB checks have more than XEH hours of excess scheduled. The maximum excess hours are MEB hours. Compared to the other types of maintenance checks, the cabin checks have the highest average of excess hours per check. Furthermore, at PEHXB, the cabin checks also have the highest percentage of checks with more than XEH hours of excess. The negative average deviation of ADB hours further emphasizes that on average all cabin checks have approximately ADB hours of excess. As shown in Table 7.2, the average hours of the OOP and modification tasks which have been allocated to the cabin checks are slightly higher compared to the A-checks. However, as also indicated by the percentages of allocated OOP and modifications tasks in Table 7.2, a smaller number of OOP and modification tasks has been allocated to the cabin checks compared to the A-checks. Thus, compared to A-checks, the modification and OOP tasks allocated to the cabin checks were generally tasks with larger required labor

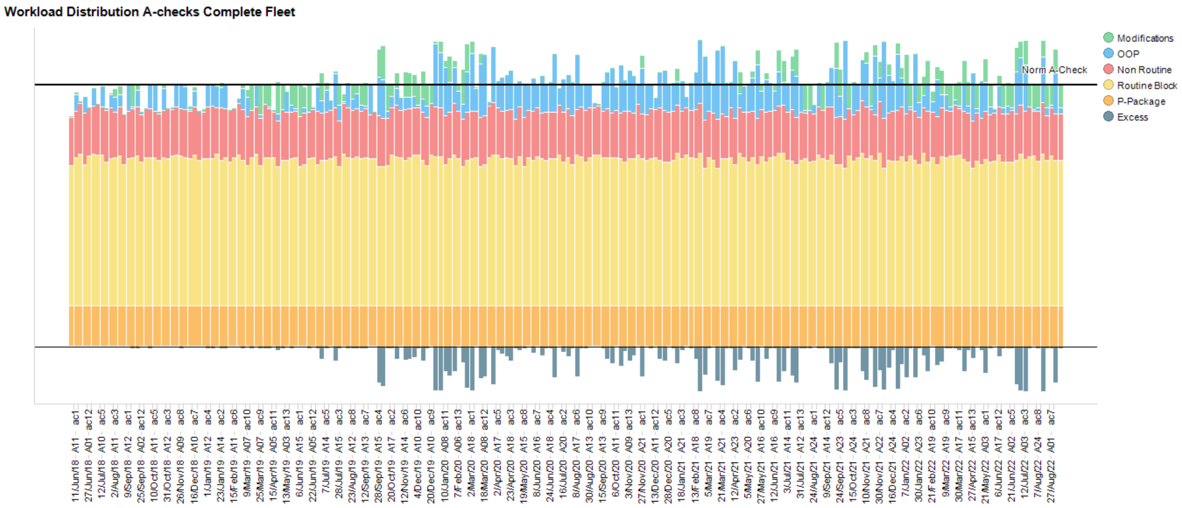

Figure 7.2: Workload distribution of the A-checks for the airline's wide-body fleet over a four year planning horizon. From top to bottom the bar segments are the allocated capacity in required labor hours of the modifications, OOP, non routine tasks, routine block tasks, p-packages. The excess hours are shown on the negative y-axis.

hours.

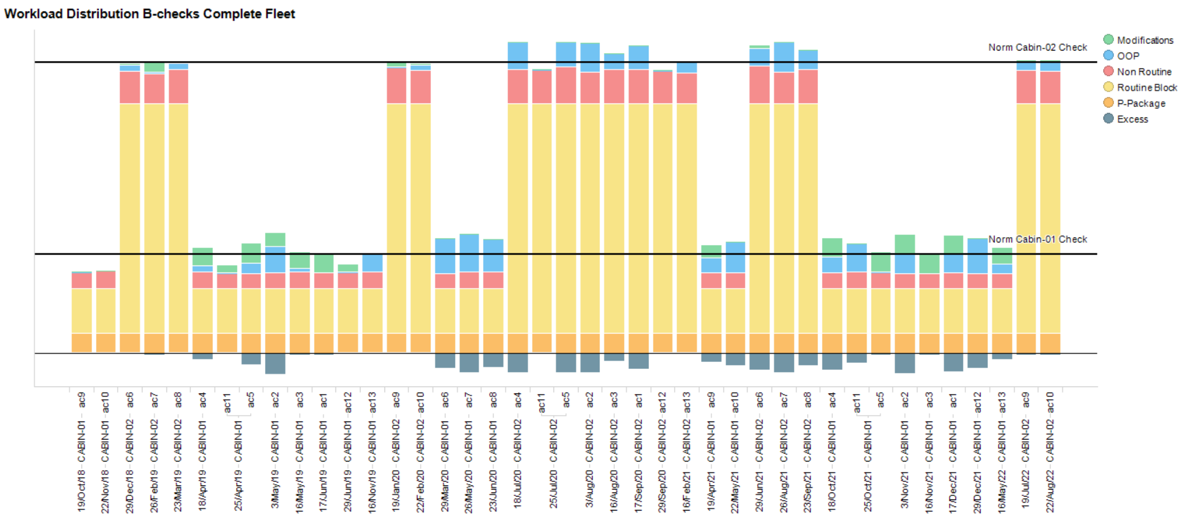

Figure 7.3: Workload distribution of the Cabin-checks for the the airline's wide-body fleet over a four year planning horizon. From top to bottom the bar segments are the allocated capacity in required labor hours of the modifications, OOP, non routine tasks, routine block tasks, p-packages. The excess hours are shown on the negative y-axis.

C-Check Distribution

The workload distribution of the C-checks is depicted in Figure 7.4. It is apparent that the workload of the C-checks is highest for the first eight C-checks. These C-checks correspond to the eight oldest aircraft in the airline's wide-body fleet. Furthermore, it is apparent from Figure 7.4, that the workload of these C-checks consists predominantly of modifications. In addition to the C01-checks of the entire fleet, the C02 checks of the eight oldest aircraft are also scheduled by the model within the four year planning horizon. It is interesting to note that the workloads of the C02-checks of these eight oldest registrations exceed the workloads of the C01-checks of the five 'youngest' registrations both in terms of OOP tasks and modification task workloads.

As can be seen from Figure 7.4 and as is listed in Table 7.2, only NEC of the TCC C-checks have excess hours, with a maximum of MEC hours. The majority of the C-checks has substantial amounts of unused resource capacity, especially within the second half of the planning horizon. On average the C-checks have ADC hours available, with a maximum of MDC hours (out of CNORM). PRMODC percent of the modifications' workload is scheduled to the C-checks. Moreover, just over PROOPC percent of the OOP task workload is scheduled to the C-checks. As can be seen from Table 7.2, the average modifications and OOP tasks hours per C-check are the highest of all check types. This indicates that the modifications and OOP tasks with the largest required labor hours are allocated to the C-checks.

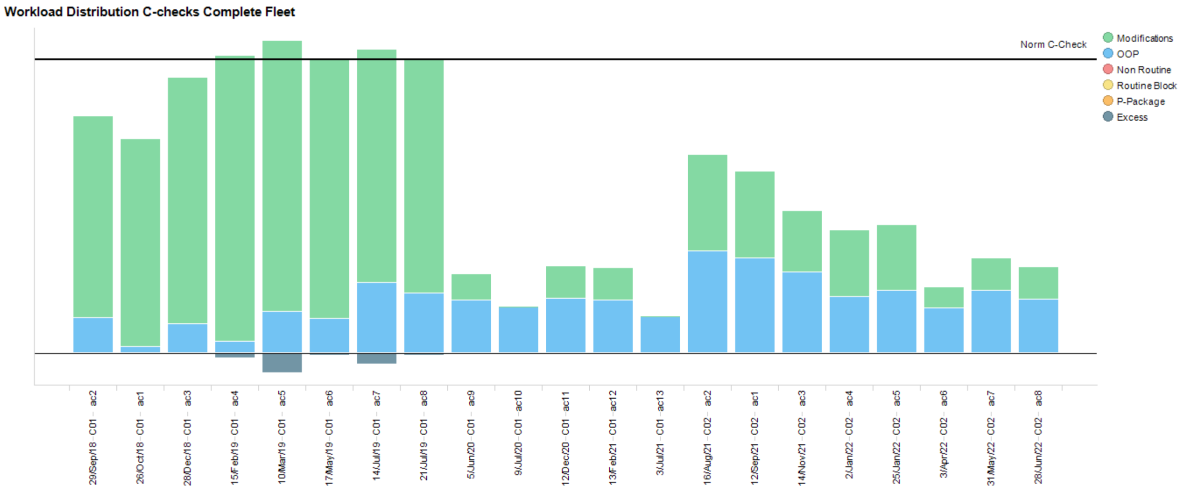

Figure 7.4: Workload distribution of the C-checks for the airline's wide-body fleet over a four year planning horizon. From top to bottom the bar segments are the allocated capacity in required labor hours of the modifications and OOP tasks. The excess hours are shown on the negative y-axis.

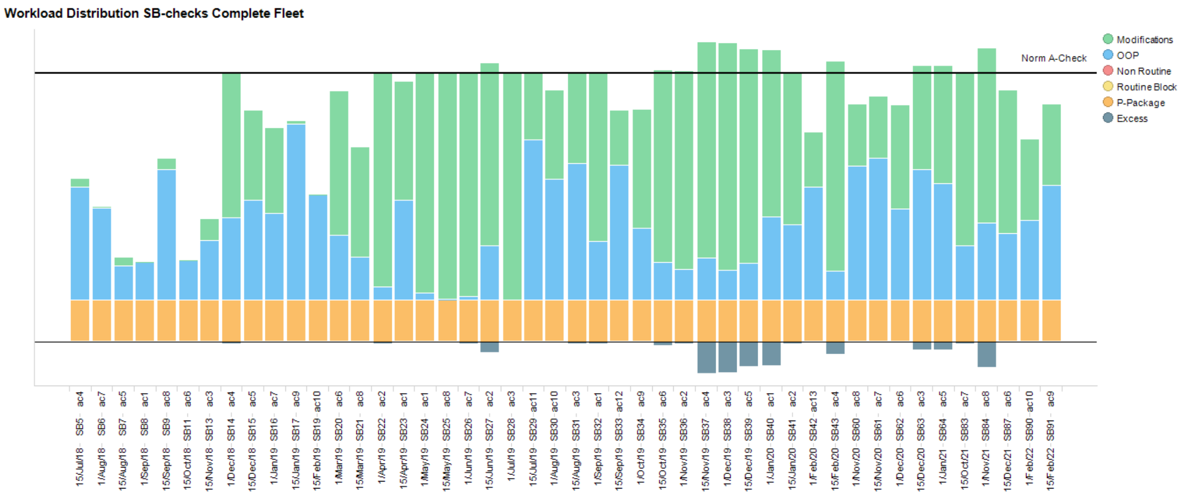

Figure 7.5: Workload distribution of the SB-checks for the airline's wide-body fleet over a four year planning horizon. From top to bottom the bar segments are the allocated capacity in required labor hours of the modifications, OOP tasks and P-packages. The excess hours are shown on the negative y-axis.

Extra Check Distribution

Figure 7.5 shows the workload of the extra checks, also referred to as the SB-checks, which have been allocated by the model. A total of TSBC extra checks have been allocated by the model. The norm for the SB-checks implemented in the model has been assumed to be the same as the norm for the A-checks, at ANORM hours. The norm is depicted by the horizontal line segment in Figure 7.5.

It is apparent from both Figure 7.5 and Table 7.2 that a large sum of the modification workload has been allocated to the SB-checks. More specifically, with PRMODSB percent, the quantity of modification hours allocated to the SB-checks is almost the same as for the C-checks. It would appear from Figure 7.5, that relatively few SB-checks have excess hours. As shown in Table 7.2 this is not necessarily true, since NESB SB-checks actually have excess hours. However, just NEHXSB of these extra checks have more than XEH hours of excess. Furthermore, it is apparent from Figure 7.5, that only SBSH of the TSBC SB-checks are scheduled in the second half of the planning horizon. Finally, Figure 7.5 shows that the first few extra checks, scheduled in the second half of 2018, have large quantities of unused available capacity. As noted in Table 7.2, a maximum of MDSB hours of available capacity is left unused in the extra checks.

7.1.3. General Results Overview

A combined overview of various performance and planning parameters for all maintenance checks is shown in Table 7.3. As can be seen from Table 7.3, of the total of TC maintenance checks, TSBC extra checks are scheduled by the model, equivalent to PSBC percent. The NET checks with excess hours equate to PET percent of all checks. The maximum observed excess hours of all check types is MET hours. As can be seen from Table 7.2, it is interesting to note that the maximum excess hours are around EH1 hours for all check types.

A comparison of the average and total excess and shortage hours in Table 7.3 shows that overall three times more capacity is left unused compared to the excessive use of resource capacity. A comparison of the total excess penalty with the objective function of the model shows that the excess penalty predominates the objective function value. More on this will be discussed in Section 7.3. Furthermore, of the NSUM maintenance checks which are scheduled during the summer season, ESUM maintenance checks have excess hours, equivalent to approximately PESUM percent. Furthermore, of the total of TASBC A- and extra checks, EASB checks have excess hours, which equates to PEASB percent.

Finally, Table 7.3 shows that of the TTS tasks which have been allocated by the model TID interval days were unused. This equates to an average of AVID days per task, with a maximum reported interval loss of MID days.

7.2. Further Analysis of the Results

The previous Section presented the results and various observations. This Section provides a more detailed discussion and analysis of the most prominent observations of the results.

7.2.1. Distribution of Excess Hours over the Planning Horizon

It has been observed in the results that the excess hours generally occur in the second half of the planning horizon. The results for both the single aircraft and fleet have shown both a higher frequency and higher quantity of maintenance checks with excess hours in the second half of the planning horizon, compared with the first half. Of the TEH excess hours in total, SHE are allocated in the second half of the planning horizon. A further analysis of the extra check allocation, task due dates and task workloads explains why this is the case. A summary of this analysis is depicted by the piecharts in Figure 7.6.

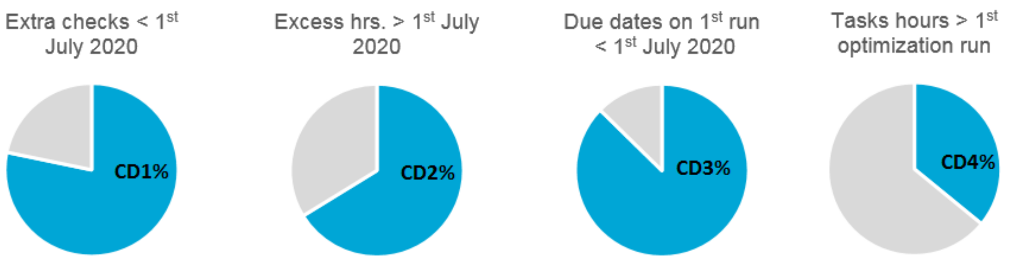

Figure 7.6: From left to right: piechart of the total number of excess hours scheduled in the second half of the planning horizon; piechart of the number of extra checks allocated to the first half of the planning horizon; piechart of the number of tasks in the initial optimization run with due dates in the first half of the planning horizon; and piechart of the workload of tasks to be allocated after the initial optimization run

Table 7.3: Combined overview of planning parameters for all maintenance checks for the optimization run presented in this Chapter.

Parameter	Value	Description
Total tasks scheduled (-)	TTS	Total number of tasks scheduled
Total checks (-)	TC	Total number of maintenance checks
Extra checks (%)	PSBC	Percentage of extra checks w.r.t. the total number of checks
Checks with excess (%)	PET	Percentage of checks with excess hours scheduled w.r.t. the total number of checks
Total excess hours (hrs)	TEH	Summed total of all excess hours of all checks
Avg. excess hours (hrs)	AVEX	Average excess hours of all checks
Max. excess hours (hrs)	MET	Maximum excess hours scheduled of all checks
Summer checks with excess (%)	PESUM	Percentage of checks scheduled during the summer season with excess hours
A- and extra checks with excess (%)	PEASB	Percentage of A-checks and extra checks with excess hours
Total shortage hours (hrs)	TSH	Summed total of all shortage hours of all checks
Avg. shortage hours (hrs)	AVSH	Average shortage hours of all checks
Max. shortage hours (hrs)	MDC	Maximum shortage hours scheduled of all checks
Objective Function Value (EUR.)	$1.47851 \cdot 1e5$	Objective function value of the last optimization run
Total excess penalty (EUR.)	$1.41406 \cdot 1e5$	Summed total of all excess penalties of the last optimization run
Max. excess penalty (EUR.)	MEP	Maximum excess penalty allocated of all checks
Total interval days remaining (days)	TID	Summed total of all interval days remaining of all tasks
Avg. interval days remaining (days)	AVID	Average interval days remaining of all tasks
Max. interval days remaining (days)	MID	Maximum interval days remaining of all tasks

As also observed in the previous Section, the majority of the extra maintenance checks are allocated to the first half of the planning horizon. In fact, XSB of the total TSBC extra checks which have been allocated by the model were allocated to the first half of the planning horizon (prior to July 2020). This equates to CD1% of all extra checks. Of the NT0 tasks allocated during the initial optimization run, DD1F tasks have a due date in the first half of the planning horizon, equivalent to CD3 percent. Provided that the model is restricted to only allocate extra checks during the initial optimization, these extra checks will therefore predominantly be scheduled to the first half of the planning horizon, since the majority of the tasks scheduled in the initial run has a due date in the first half of the planning horizon. However, another NTX tasks are to be allocated after the initial optimization run. Together these tasks have a workload of XTH hours which is equivalent to approximately CD4% of the total workload of TTH hours of all tasks allocated by the model. Thus, although a substantial portion of the workload has yet been allocated on the initial run, approximately a third still has to be allocated during all subsequent runs. Without the possibility of allocating extra capacity, this workload can only be allocated to the available capacity of the maintenance checks which remained after the initial run. This subsequently results in excess hours. Provided that the tasks allocated in later optimization runs also have later due dates, the majority of these excess hours are predominantly allocated to maintenance checks in the second half of the planning horizon.

7.2.2. Maximum Excess Hours

Another prominent observation from the results is that maximum excess hours on the various check types is just above EH1 hours. In fact, the maximum excess hours of all check types, except the extra checks, are approximately EH1 hours. Figure 7.7 shows a histogram of the excess hours of all maintenance checks with excess. As can be seen from the histogram, the majority of excess is smaller than XE hours. Furthermore, the piechart in Figure 7.7 shows that for 95% of all maintenance checks with excess, the excess is less than EH1 hours.

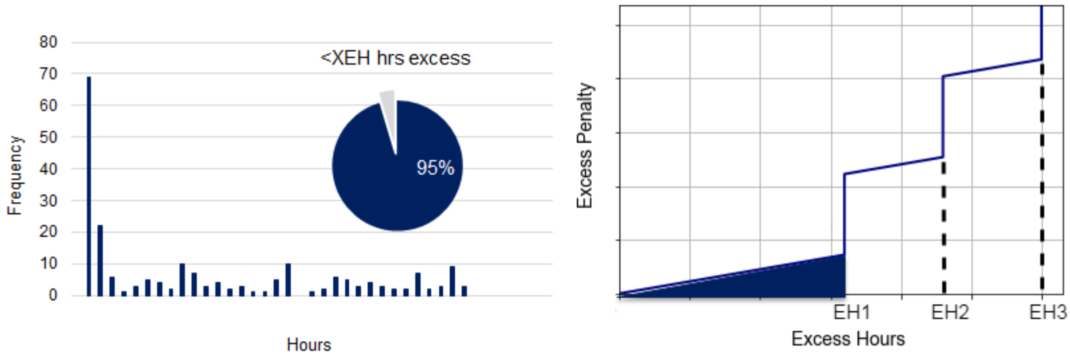

Figure 7.7: Left:Histogram and piechart of the excess hours, Right: piecewise linear function for excess penalty, the highlighted section indicates excess hours smaller than XEH hours

As shown by Figure 7.7, the reason for this observed behavior is simply because the excess penalty increases substantially beyond EH1 excess hours. As explained in Chapter 4 and Chapter 5, the relation between the excess hours and corresponding excess penalty is modeled by a piecewise linear function. The first breakpoint in the piecewise linear function in the current model setup is located at EH1 excess hours. Thus, it would appear that piecewise linear relation between the excess hours and excess penalty works effectively to discourage excess larger than EH1 hours.

However, not all maintenance checks with excess have less than EH1 hours. As can be seen from Figure 7.7, 5% of maintenance checks with excess has more than EH1 hours of excess. Although the reported overall maximum excess of MET hours lies only just above the EH1 hours, why was the model not able to keep the excess hours of these XMC maintenance checks below EH1 hours as well?

The primary reason for this behavior is the design of the model itself. More specifically, the iterative solution technique which has been built around the model in combination with the set of 'previous execution date' constraints.
For example, the A22 check of Aircraft 1 has EXAM excess hours allocated to it. During the first ten optimization runs the model allocates 5 maintenance tasks with a combined total workload of WL hours. Provided that this particular letter check had AMHA hours of available labor capacity prior to the allocation of modification and OOP tasks, this equates to SEH of scheduled excess hours. Although excess has been allocated, at this point the hours are still below EH1. However, on the eleventh optimization run, the model allocates a single task of AA hours to the A22 check, pushing its excess hours to just above EH1. A further observation of the task in question reveals that its was previously allocated to the A21 letter check, which precedes the A22 check. Furthermore, the due date of the task inhibits it from being allocated to the next maintenance check on the planning horizon of Aircraft 1. Thus, in order to comply with the due date constraint and previous execution date constraint, the model has to allocate this task to the A22 check, resulting in more than EH1 hours of excess.

The previous example highlights an important limitation of the model. The iterative solution technique has been introduced to deal with the recurrent nature of the OOP tasks. Although this approach enables the scheduling of all occurrences of the OOP tasks within the planning horizon and as such helps provide better insight into the complete workload of the fleet, the formulation is such that each run is singularly optimized. Hence, the allocation of the model for a single run may be optimal for that specific run, but may in fact result in a suboptimal result when combining all optimization runs.

Although this limitation of the model has been illustrated as an explanation for the XMC maintenance checks which excess hours are larger than EH1, the implications of this limitations affect the allocation of tasks to all maintenance checks. A simple example is illustrated in Figure 7.8. This Figure shows the last six maintenance checks on the planning horizon of the Aircraft 1. It is apparent that while all four A-checks and the Cabin-

01 check have excess hours, of the CNORM hours it has available only CH hours are allocated to the C02-check. Thus, the C02 check still has AMHC hours of remaining available capacity, while the other checks have excess hours. Similar to the previous example, the model has aimed to optimize each of its runs. With all checks having available capacity on the first few optimization runs, the model has allocated tasks as closely as possible to their respective due dates, to minimize the loss of interval. However, due to the non-integral nature of the model, it could not 'foresee' that it would have been better overall to sacrifice interval on the first few runs, for less excess hours in later runs.

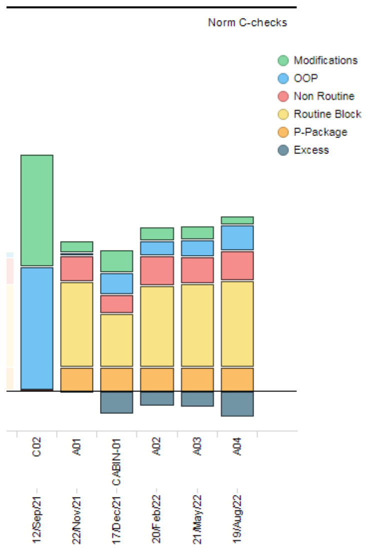

Figure 7.8: Overview of last year of the planning horizon of Aircraft 1; showing available capacity in the C02, followed by excess hours in the subsequent letter checks

Thus, it would appear that the iterative solution technique is a significant limitation to the optimality of the solution of the model. A single optimization run, in which the due dates of the OOP tasks are initially estimated and adjusted depending on their allocation would yield a more optimal solution, since the entire workload could all be considered at once, rather than in series. It is recommended that further research be conducted either to alter or replace the current iterative solution technique in the further development of the model.

7.2.3. Distribution of Shortage Hours

It has been observed that approximately 3 times as much shortage hours compared with the excess hours occur in the model results. At least some of this unused capacity originates from extra checks which have been allocated by the model. Although the extra checks are allocated to provide much needed extra capacity, it would appear from these results that not all of this capacity is required. As has been explained in Chapter 5, the penalty of the extra checks has been set equal to the the penalty of XEH excess hours. This has been based on the fact that a maintenance check of which the allocated workload exceeds the norm by more than XEH hours is considered a heavy check. Hence, the extra check provides the model with an alternative option whenever a maintenance check has more than XEH hours of excess. However, because the penalty was set equal to XEH excess hours, it may occur that an extra check is allocated only to remove these XEH hours from a letter check. For example, extra check 'SB8' for Aircraft 1 has only XSB hours allocated to it. Provided that the model is more heavily penalized for losing interval, no tasks are allocated to reduce the resulting unused available capacity of such extra checks.
In addition to the underuse of available resource capacity of extra checks, some of the shortage hours are also a result of the non-integrality of the model design as explained and illustrated in the previous Subsection.

7.2.4. Effectivity of the Seasonality and Check Type Factors

The results have shown that PESUM% of all checks in the summer season have excess hours. Moreover, PEASUM% of all A-checks and extra checks have excess hours (see Table 7.3). It would therefore appear that the seasonality and check type penalty factors are not very effective at their current settings. As explained in Chapter 5, these factors are imposed to maintenance checks with excess hours in the summer season and to A-checks and extra maintenance checks with excess hours respectively. They are to ensure that excess hours are less likely to occur in the busier summer season and to ensure a more on time performance of the regular A-checks (or extra checks) which have a relatively tight ground time. A sensitivity analysis has been conducted to assess the effectiveness of these penalty factors, the results of which are presented in Chapter 9.

7.3. Model Performance Results

This Section describes the model performance. Table 7.4 gives an overview of some of the performance indicators of the model for each of the consecutive optimization runs. The result presented in this Chapter correspond to 19 consecutive runs of the optimization model. From left to right the columns of Table 7.4 indicate the run; the number of decision variables; the number of constraints; the runtime in seconds; the solution status provided by CPLEX; the optimality gap in percentages; the number of nodes evaluated; the objective function value expressed in hundred thousand euros; and the number of tasks which have been allocated.

Table 7.4: Overview of model performance parameters for the consecutive optimization runs

Run	DVs (-)	Constr. (-)	Runtime (s)	Sol. Status (-)	Gap (%)	Nodes (-)	OF. Value (1e5)	Tasks Alloc. (-)
0	167,884	460,665	3905.63	107	3.947	15,870	1.10113	NT0
1	19,070	37,157	9.41	102	6.932e-03	2,534	0.41727	NT1
2	11,556	21,834	0.97	102	9.647e-03	0	0.66724	NT2
3	8,620	15,850	0.75	101	0.0	0	0.84232	NT3
4	7,325	13,210	0.36	102	9.680e-03	0	1.11213	NT4
5	4,992	8,452	1.78	102	4.974e-03	0	1.26608	NT5
6	3,226	4,854	0.25	101	0.0	0	1.32234	NT6
7	2,589	3,556	0.20	102	3.961e-03	0	1.30854	NT7
8	2,263	2,892	0.19	101	0.0	0	1.30825	NT8
9	2,138	2,637	0.22	101	0.0	0	1.31351	NT9
10	1,983	2,321	0.20	102	4.904e-03	0	1.35536	NT10
11	1,983	2,321	0.17	101	0.0	0	1.36104	NT11
12	1,983	2,321	0.14	101	0.0	0	1.36586	NT12
13	1,983	2,321	0.17	102	4.604e-03	0	1.37126	NT13
14	1,955	2,264	0.16	101	0.0	0	1.37622	NT14
15	1,876	2,103	0.20	101	0.0	0	1.42720	NT15
16	1,776	1,899	0.14	101	0.0	0	1.47675	NT16
17	1,574	1,487	0.22	101	0.0	0	1.47792	NT17
18	1,444	1,222	0.16	101	0.0	0	1.47851	NT18

As can be seen from Table 7.4, the number of tasks allocated per optimization run gradually decreases with the consecutive optimization runs. As explained in Chapter 4, following each optimization run the iterative solution algorithm built around the optimization model updates the due date of each OOP task for the next optimization run. Provided that a due date fall within a predefined threshold, the OOP task is fed as an input to the next optimization run. The threshold equals the calendar date which results from adding the planning horizon of four years plus one A-check interval to the first letter check date. OOP tasks for which the next due date falls outside of this threshold date are no longer included in the subsequent optimization runs. Thus, the number of tasks gradually decreases with subsequent optimization runs.

Table 7.4 clearly shows that the number of decision variables and constraints scales proportionally with the number of tasks to be allocated. The number of decision variables and constraints is by far the largest for the initial run. In addition to the NT0 tasks which are to allocated, the model also has to decide on the allocation of TNSB extra checks for the 13 aircraft during the initial run. The scale of the decisions to be made during this run is further reflected in both the number of nodes evaluated and the runtime, both of which are substantially larger for the initial run compared with the subsequent runs. Provided that the extra checks are only allocated in the initial run, for all subsequent runs the number of decision variables and constraints is proportional to the number of tasks only. As shown by the 0's in the "Nodes" column of Table 7.4, for all but the first two runs CPLEX only has to evaluate the root node to arrive at an optimal solution.

The objective function values of the consecutive runs show a drop between the initial run and run 1, after which the values gradually increase with each subsequent optimization run. The reason for this is that the objective function values shown in Table 7.4 correspond to the objective function of each respective optimization run only, rather than the accumulated value of all objective function values of all previous optimization runs.
As explained in Chapter 4, the shortage and excess hours of the maintenance checks are passed on from one optimization run to the next. This has to be done to ensure that each subsequent optimization run can only use the available capacity (or lack thereof) which remains after the task allocation of the previous runs. As such, the sum of the shortage and excess penalties are also accumulated with each optimization run. Hence, the gradual increase in the objective function values observed in Table 7.4 is caused by the accumulation of the shortage and excess penalties of previous runs. Furthermore, provided that the extra check penalty and interval penalties are not passed on, the drop observed between run 0 and run 1 is caused by the "loss" of the extra check penalty. As the extra checks are only allocated during the initial run, the corresponding penalty is also only considered in run 0 which is not considered in the objective function value of run 1.

The "Sol. Status" refers to the solution status of the optimization and is a parameter provided by CPLEX. The three numbers observed in Table 7.4 imply the following [54]:

- **101:** CPLEX has proven that the current solution is optimal.
- **102:** The current integer solution has not been proven to be optimal, but the solution lies within the tolerated optimality gap from the best feasible solution and therefore "appears optimal".
- **107:** The current integer solution is not the optimal solution. The model has stopped at the current solution as the predefined "time limit has been exceeded".

As can be seen from the solution status and optimality gap columns in Table 7.4, nearly all optimization runs yield either a proven optimal solution, or a nearly optimal solution as indicated by the small optimality gaps. The optimality gap, or MIP gap, indicates the difference between the best integer solution and the best feasible solution. Where the latter is the non-integer (linear-relaxation) solution of the tree and the former is currently the best integer solution found. CPLEX will keep iterating for a better solution either until it has exceeded the permissible time limit, or until the optimality gap meets some predefined criteria. In the present model the time limit for CPLEX was set to 3600 seconds and the tolerance of the MIP gap to 0.01%.

Contrary to the other optimization runs where CPLEX stopped because the gap between the integer and best feasible solution was less than the specified 0.01% or because the optimal solution had proven to be found, for the initial run CPLEX stopped after running for just over an hour, as specified by the predefined time limit. The gap between the best feasible and best integer solutions which had been found within that time frame is 3.947%. This equates to an absolute difference in the order of 4,000. The desired optimality gap of 0.01 percent would be equivalent to an absolute difference in the order of 10. To assess what the time limit should be in order for CPLEX to reach the (nearly) optimal solution on the initial optimization run, the convergence rate of the optimality gap was analyzed. The convergence rate of the optimality gap for the initial optimization run has been plotted over time in Figure 7.9.

As can be seen from Figure 7.9, the convergence rate is highest within the first 300 seconds. After 291 seconds the optimality gap has been reduced to 5.77%. During the remaining 3500 seconds the optimality gap is reduced by a further 1.82% to 3.95%. At the convergence rate of 1.82% in 3500 seconds, the desired optimality gap of 0.01% would be attained in a total runtime of approximately 3.5 hours. A rerun of the model was performed in which the time limit was set to ten hours, or 36,000 seconds. Unfortunately, this rerun was stopped due to a CPLEX memory error. Due to the time constraints on the present study, no further improvements could be conducted at the present time. Thus, the optimality gap of 3.95% was considered sufficiently optimal for the time being. However, it is highly recommended that further studies are conducted into the

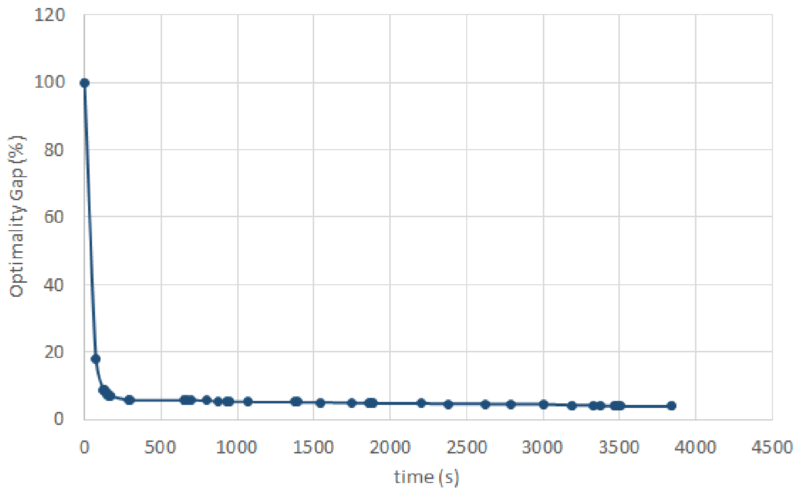

Figure 7.9: Convergence rate shown by the reduction of the optimality gap over time

improvement of the optimality of the model. A successful 'unbounded' time run in CPLEX would allow for a trade-off to be conducted between the value of the improved model optimality to the MRO's maintenance planning and the corresponding runtime.

7.4. Application of Results to the MRO

So far this Chapter has presented the results and highlighted and discussed some of the most prominent observations. This Section will briefly discuss some of the potential benefits of the task scheduling optimization model for the airline's wide-body fleet.

The model results presented in this Chapter show the potential of a task specific, integrated maintenance planning as modification and OOP tasks have been allocated to all check types and over the full planning horizon. Moreover, the planning provided by the model has been optimized based on the available capacity of all letter check within a four year planning horizon and the required capacity of the actual workload of all base maintenance tasks. As such it provides maintenance planners with a suggestion on where to allocate the various OOP and modification tasks. It is assumed that this planning is particularly helpful for the allocation of modifications, for which there are currently no sufficient means to plan due to the lack of insight into the available resource capacity on a tactical planning horizon. It is anticipated that the model's insight will contribute to a faster implementation of modifications, which will in turn yield operational and financial benefits. Moreover, the model may assist the MRO in providing the airline with a more firm commitment of when the various modification will be executed, which will yield a more satisfied customer.

However, perhaps the most valuable insight of the model is the ratio between the available capacity of the letter checks compared with the complete workload of the maintenance tasks over the coming four years. Given, the current norms for the maintenance checks and the known workload on the aircraft, the model shows that without any additional measures the available capacity of the maintenance checks is not sufficient for the capacity required by the modifications and OOP tasks within the four year time window. An overview of the available capacity compared with the required resource capacity per aircraft registration is shown in Figure 7.10. Here, the available capacity refers to the capacity of the letter checks only, excluding the extra maintenance checks, which has been computed by subtracting the routine block hours, p-package hours and non routine hours from the norms of the various check types. Moreover, the required capacity has been based on the total hours required by all modification and OOP tasks allocated within the planning horizon.

As shown by the Table, for the first nine aircraft registrations the workload exceeds the available resource capacity. The difference between the required and available capacity is largest for Aircraft 2 with UNDX more capacity required than available. The four latest aircraft have sufficient capacity for their respective workloads. This is primarily due to the fact that these aircraft have a smaller modification workload. Overall, the summed required capacity exceeds the summed available capacity by approximately UNDY hours.

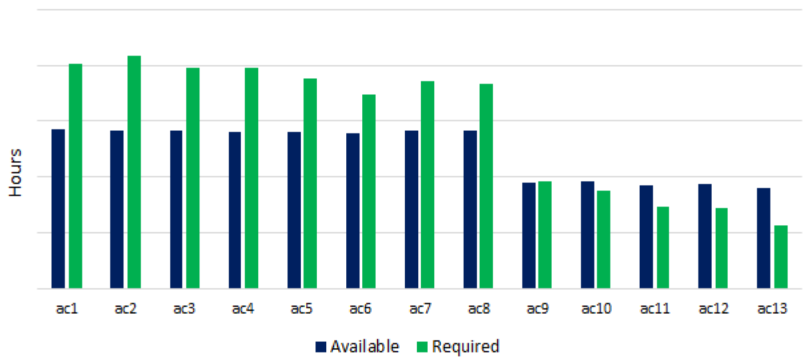

Figure 7.10: Comparison of the lump sums of the available and required capacities per aircraft registration (ac)

The comparison shown in Figure 7.10 only shows the lump sums of the required and available resource capacities. As such, it ignores the due date constraints and capacity constraints of the maintenance tasks and checks which further implicate the effective availability of the maintenance checks. Thus, the actual surplus of the required capacity in comparison with the effective available resource capacity will be even greater. As such the surplus shown in Figure 7.10 only highlight the need for more available resource capacity. Without additional measures, there simply are not enough available resources to accommodate the workload of all maintenance tasks. However, the fact that the model does include due dates constraints of tasks and the capacity constraints of maintenance checks implies that it can provide even more accurate insight into the effective available capacity. As such, the model can also provide much better insight into how much additional capacity is required and where to allocate any additional capacity. Alternatively, the norms of the various maintenance checks could also be increased. However, provided that in the model the norms have been assumed as true, the model schedules extra checks to compensate for the extra capacity which is required. On average, the model schedules AVSB extra checks per aircraft, equivalent to an average of SBH extra hours per aircraft over the four year planning horizon. As observed from the workloads per aircraft registration shown in Figure 7.10, more extra checks are required for the first eight aircraft registrations and less for the other registrations. The model provides a suggestion of when to allocate additional maintenance checks and to which aircraft registration, such that the workload may be allocated most efficiently.

Finally, it is important to note that all of the results and observations shown and discussed are heavily dependent on the inputs and therefore subject to change. The workload of modifications is constantly subject to the addition of new modifications. Identically, the required capacity of the OOP tasks and all other routine tasks for that matter, are subject to changes of the AMP; inclusion of new tasks, declaring former tasks obsolete or other alterations of tasks. Furthermore, a more accurate prediction of the workload of non routine tasks will impact the available capacity of the various maintenance checks. The sensitivity of the model to such changes will be discussed more elaborately in Chapter 9. The inevitability of such changes also speaks in favor of an automated optimization model such as the present model. The constant changes to maintenance tasks and checks require a fully automated model which can cope with such changes and optimally and efficiently recompute a new task schedule for the full planning horizon.

7.5. Summary on Anomaly

As mentioned in the introduction of this Chapter, an anomaly was detected during the analysis of the results. This anomaly was subsequently corrected for and the model was rerun to produce the results presented in this Chapter. However, the anomaly revealed a potential weakness in the present model formulation. Hence, this Section will briefly summarize the anomaly, its effect on the optimality of the model results and the detected potential weakness in the model formulation.

Description of the Anomaly

During the original analysis of the results, it was observed that the objective function value was dominated almost completely by the excess penalty. The initial run of these results had an objective function value of 150.18 and an excess penalty of 150.17 million. Moreover, the excess penalty in turn was predominated by one single excess penalty of 150.00 million. Further analysis revealed that this excess penalty corresponded to a single extra check which had more than EH3 excess hours. More specifically, it was found that a single modification task of YMOD hours has been scheduled to this extra check, which had 'only' ANORM hours of available resource capacity. For all other aircraft registrations the modification has been scheduled to the C01-checks of the respective aircraft. However, for a single aircraft, the due date of the task happened to be prior to the check date of the C01-check. In order not to violate this due date, the task was scheduled to an extra maintenance check as this would provide the largest, yet insufficient, available resource capacity.

Effect of the Anomaly on the Model Results

As a result of the anomaly, the objective function had an order of magnitude of approximately 150 million. As such, when CPLEX ran the initial optimization run it found an integer solution which was within the permissible 0.01% from the best feasible solution within 586 seconds and thus CPLEX was stopped. However, provided the magnitude of the objective function value this is equivalent to an absolute error of 15,000. This is significantly larger compared with the absolute difference of 4,000 observed for the results presented in this Chapter. Moreover, if one would exclude the singular penalty from the objective function value, the remaining objective function value would be in the order of magnitude of 100 to 200 thousand (similar to what has been observed in this Chapter). This would make the gap of 15,000 equivalent to approximately 10 percent, instead of 0.01 percent.

Provided the much larger absolute difference between the best feasible and best integer solutions, the anomaly had a significant impact on the optimality of the solution. This was also observed from other model results, some of which have been summarized in Table 7.5. The 'revised' model results refer to the model results which have been presented in this Chapter, thus after the anomaly was removed.

Table 7.5: Comparison of various model output parameters for original results and revised model results

	Original Results	**Revised Results**
number of extra checks allocated	OSBC	TSBC
percentage of checks with excess	OPET	PET
total excess hours	OTEH	TEH
max. excess hours	OMET	MET
avg. excess hours	OAVEX	AVEX
total shortage hours	OTSH	TSH
max. shortage hours	OMDC	MDC
avg. shortage hours	OAVSH	AVSH
summer checks with excess (%)	OPESUM	PESUM
A-/extra checks with excess (%)	OPEASB	PEASB
total interval days remaining	OTID	TID
average interval days remaining	OAVID	AVID

As can be seen from Table 7.5, there are substantial differences between the two sets of model results. It is apparent that the revised results show less overall excess hours and fewer unused interval days. Moreover, the revised results show much lower excess hours scheduled to the A-checks and extra checks and to the maintenance checks in the summer season. In general the revised results show a more optimal task schedule compared with the original results, which is consistent with the observed lower optimality gap.

The anomaly was removed by simply forwarding the original C01-check date of the aircraft in question by three weeks, such that it would allow for the modification with YMOD labor hours to be allocated to the C01-check. The results presented in this Chapter confirm that this simple change had the desired effect. The modification in question is now allocated to the C01 check, removing the disproportionally large excess penalty and subsequently resulting in a more optimized model result.

Potential Weakness in Model Formulation

Although the anomaly was resolved fairly simply, it has highlighted a potential weaknesses in the current model formulation. As explained, the allocation of a single task incurred an excess penalty which was so disproportional to all other penalties that it reduced the optimality of the overall model results. The substantial excess penalty was imposed because the excess hours exceeded EH3 hours. In the current model setup, the post-slope of the piecewise linear function for the excess penalty is located at EH3 hours. This post-slope is equal to PSX million per hour. Subsequently, the excess hours of the anomaly incurred an excess penalty in the order of 150 million.

As has been observed, overall the piecewise linear relation between the excess hours and penalty appears to work very well, as most excess hours are kept below EH1 hours, at which point the first step in the function is located. However, as has also been observed, in exceptional circumstances, the excess hours may be such that excess hours beyond the 'upper limit' can occur. As singularities such as the anomaly discussed in this Section cannot entirely be avoided, it is recommended that the post-slope of the piecewise linear function be reformulated such that it in case of anomalies the allocated excess penalty is proportional to the other penalties and objective function value. This adjustment will ensure that even in case of anomalies, the optimality of the solution is not compromised.

This Chapter has introduced the results, showing the outcome of the optimization runs, as well as highlighting and discussing the decision making process of the model, its performance and its applicability to the MRO. The following two Chapters will present and analyze the validation of the model results and the sensitivity analysis of the model, thereby assessing the accuracy and robustness of the model respectively.

8

Model Validation

This Chapter discusses the validation of the model results. The validation results provide an important contribution to the research questions formulated in the Introduction of the report. It is important to note that the model and its individual modules have been verified previously, during the development of the model. An elaborate explanation of the verification process will not be provided in this report. Rather, it suffices to state that a satisfactory verification has been performed.

The validation has been performed through a thorough comparison of the task schedule resulting from the model with the actual check data from the MRO's check detail reports (CDR) for a selection of maintenance checks. The performance of the model has been assessed based on accuracy and efficiency. Section 8.1 discusses the validation of task occurrence as it compares whether the tasks scheduled in practice also occur in the model. Subsequently, Section 8.2 presents the validation of task hours, comparing the scheduled labor hours used by the model with practice for all task types. A comparison of the accuracy and efficiency of task scheduling decisions made by the model compared with best practices is discussed in Section 8.3. Finally, a discussion on the validation results is provided in Section 8.4.

The following maintenance checks have been used for the validation:

- A09 check of Aircraft 5 which was scheduled for July 6th, 2018
- A10 check of Aircraft 7 which was scheduled for August 24, 2018
- A02 check of Aircraft 12 which was scheduled for September 19th, 2018
- C01 check of Aircraft 2 which was scheduled for October 7th, 2018

The selection of the sample set of checks has been based on the diversity of letter checks, as well as the diversity of scheduled maintenance tasks; provided the OOP and modifications which were allocated to these checks in practice. Unfortunately, no cabin checks were scheduled by the model prior to the validation by the author; thus these could not be included in the validation. Nevertheless, the variety of maintenance checks and tasks is considered sufficient to validate the model.

8.1. Accuracy of Task Occurrence

This Section presents and discusses the validation of task occurrence. A matching task occurrence implies that the tasks which were scheduled to the sample set of maintenance checks have also been scheduled by the model, and vice versa. It is important to note that this comparison does not require the tasks to be scheduled in the same maintenance check. This will be discussed in Section 8.3. Thus, if there is a mismatching task occurrence, the task either does not occur in the model while it was scheduled in practice, or vice versa. The task occurrence validation results are shown in Table 8.1.

For the block tasks, a comparison of both the MRIs and JICs has been conducted. For the OOP tasks the comparison was restricted to MRIs only. A comparison of the EOs was used for the comparison of task occurrence of modification tasks. As explained in Section 6.2, it has been assumed in the model formulation that an OOP task is a 'stand-alone' task consisting of all its JICs, access panel hours and even its own non routine hours.

Table 8.1: Overview of the percentages of mismatches for the MRIs, JICs or EOs

	Block		OOP	Modification
	MRI	JIC	MRI	EO
A02 - Aircraft 12	3.6	2.7	50.0	100.0
A09 - Aircraft 7	3.2	1.9	30.0	85.7
A10 - Aircraft 5	2.3	2.5	0.0	100.0
C01 - Aircraft 2	-	-	-	19.0
Total Percentage	3.0	2.4	17.9	55.0

In practice, a maintenance planner would remove any JICs from the OOP task if the same JIC already occurs in the workpackage of the check to which the OOP task is allocated. As this is omitted in the model, the OOP tasks will occasionally produce 'duplicate JICs' in the workpackage. The resulting error is conservative and has been accepted. Thus, for this reason the comparison of task occurrence for OOP tasks has been restricted to the MRIs only in the validation.

As can be seen for the block tasks, the matching task occurrence of both MRIs and JICs is relatively high, with only 3.0 and 2.4 percent of the MRIs and JICs mismatching respectively. The single reason for these deviations is a revision of the AMP resulting in the addition or discontinuation of some block tasks. Central Engineering regularly conducts changes to the AMP as needs be. Updates are made accordingly in the MRO's MIS such that the maintenance planners also prepare workpackages with the latest AMP revisions. The model inputs for the block tasks dated back to before the AMP revision and as such did not include these updates. All mismatches in JICs are a direct result of the mismatches in MRIs.

The total percentage of mismatches of OOP MRIs is higher compared with the observed mismatches in block tasks. However, in absolute sense the number of mismatches was still small. Some of these mismatches were also the result of the AMP revision, as explained for the block tasks. Furthermore, for one of the mismatches the interval of the OOP task was not located in the proper designated area in the AMP. Because of this, the model discarded this task as it has been assumed that tasks without an interval specification are not active. As for the AMP revision, this is considered an input error which can easily be rectified. The last cause for mismatches results from the assumption that OOP tasks with an interval of 90 days may be excluded from the model. This assumption is based on the fact that such tasks are in practice often executed at line maintenance, also referred to as M-check tasks. However, analysis revealed that the model excluded one OOP task unjustly based on this assumption, as it appeared to be a regular A-check task. The assumption is considered too ambiguous and has to be reviewed.

At 55%, the EO have a relatively high mismatch. As can be seen from Table 8.1, nearly all EOs of the A-checks mismatch in their occurrence. There are two primary reasons for these mismatches, the first of which is found in the definition of an EO. As explained in Chapter 3, Central Engineering generally creates an EO for any non-recurrent task, among which are modifications but also various other tasks. Some of the mismatches are caused by the fact that while the model only considers EOs originating from modifications, in the data from the actual checks all EOs have been taken into consideration, including EOs other than modifications. Out of all the EOs which have been compared in the validation, almost a third did not originate from modifications but 'other' non recurring tasks.
The second cause for mismatches results again from an input error. A secondary project at the MRO in which the author was involved required a frequent update of the same input file for modifications used by the present model. Consequently, the author mistakenly used the updated file rather than the original input file for the model run. The continuous scheduling and execution of modifications implied that some of the modifications were no longer in the updated input file and therefore also not scheduled by the model. Half of the mismatches may be attributes to this. Together with the aforementioned error in EO definition, this equates to all mismatches. Disregarding these would imply a 100% match in occurrence for the modifications.

Based on the observations discussed in this Section it may be concluded that the accuracy of the model in terms of the occurrence of tasks is high. Most mismatches are due to input errors which can easily be corrected for and which are not attributable to the model itself.

8.2. Accuracy of Task Hours

The accuracy of task hours entails the comparison of the labor hours of tasks as scheduled by the model compared with the actual scheduled hours of the tasks according to the check detail reports. A comparison with the actual spent hours of the tasks according to the check detail report has not been included. The primary objective of the validation is to compare the hours used for planning, rather than to provide a comparison of the planned and spent hours. Subsection 8.2.1 discusses the comparison of block task hours, followed by a comparison of the OOP task hours in Subsection 8.2.2. Subsequently, the comparisons of the modification task hours and other miscellaneous task hours are discussed in Subsections 8.2.3 and 8.2.4.

8.2.1. Comparison of Block Task Hours

Table 8.2 provides an overview of the comparison for the block hours. From the top down, the following parameters are included in Table 8.2: the total hours according to the check data; the total hours as scheduled by the model; the difference of the former two as a percentage; the percentage of tasks with a non-zero absolute deviation indicated as 'qnty. dev'; the maximum absolute deviation of the model hours compared with the check hours; and the average of all tasks with a non-zero absolute deviation. It is important to note that only the hours of those tasks for which the occurrences matched have been included in the comparison.

Table 8.2: Comparison of model and check detail report block task hours

Parameter	**A02**	**A09**	**A10**
Total CDR hours	ABH1	ABH2	ABH3
Total model hours	MBH1	MBH2	MBH3
Total hours diff. (%)	-7.0	-0.7	-2.8
Qnty. dev. (%)	8.1	6.4	5.8
Max dev. (hrs)	8.25	2.23	2.25
Avg. dev. (hrs)	1.50	0.83	0.79

The total hours scheduled by the model are lower than the actual check hours. As expected, the mismatch in MRI occurrences discussed in the previous Section contributes at least in part to the observed mismatch in scheduled hours. Furthermore, analysis has shown that there are some JICs for which the actual scheduled hours according to the check detail reports are exactly half of the hours scheduled by the model. Inquiry with maintenance planners and maintenance program engineers has revealed that the correct labor hours are contained in the MIS, but somehow the hours are not reported correctly in the check detail reports. For these JICs two maintenance workers are required, but it would appear that the hours reported in the check detail reports only account for one.

Furthermore, there is a single MRI which occurs in each A-check of which the actual scheduled hours deviate consistently from the scheduled hours of the model and the scheduled hours as prescribed in the MIS. This particular task is used to account for the labor time required by the lead engineers to sign off of all maintenance tasks as stipulated by regulations. Oftentimes, any additional time spent on overhead concerning the management of the check is also listed under this task. For both the A02 check of Aircraft 12 and the A10 check of Aircraft 5, the maximum deviation in block hours is caused by the difference in scheduled hours for this task. In case of Aircraft 12's check, 8.5 hours were scheduled by the maintenance planner, compared to only 0.25 hours scheduled by the model. The MRO's MIS prescribes only 15 minutes, or 0.25 hours, for the task. It remains unclear why more hours were allocated for this single task to these checks in practice.

In general, the differences in block hours may be considered relatively small. Moreover, the observed errors may all be attributed to reporting errors, rather than errors in the model design.

8.2.2. Comparison of OOP Task Hours

Table 8.3 provides an overview of the comparison for the OOP MRI hours. Provided the relatively small number of OOP tasks scheduled to a single check, the comparison is reported for all OOP tasks collectively, rather than per check.

The XMOH extra hours scheduled by the model equate to an additional 67% compared with the actual scheduled hours. In addition to the large deviation in total hours, the percentage of OOP tasks of which the model hours and check detail report hours mismatch (see 'Mismatches' in Table 8.3) indicates a high frequency of different hours. The last row of Table 8.3 is the relative average deviation, which indicates the average absolute deviation (which includes tasks without deviation) compared with the average hours of all OOP tasks. The average hours of an OOP task have been based on the hours reported by the CDR, not of the model. The relative average deviation is an indication of how substantial the average absolute deviation is, as also reported in Table 8.3. At 67% the relative absolute deviation of the OOP tasks is high. The hours of the OOP tasks scheduled by the model are both frequently and substantially above the hours scheduled in practice. There are four primary reasons for the observed deviations in the scheduled hours of the OOP tasks.

Table 8.3: Comparison of model and check detail report OOP task hours

Parameter	Value
Total hours CDR	AOH
Total hours model	MOH
Total hours diff (%)	67.3
Mismatch (%)	65.2
Max. dev. (hrs)	4.23
Avg. dev. (hrs)	1.17
Rel. avg. dev. (%)	67.3

One of the reasons for the deviations in the OOP hours is the occurrence of duplicate JICs. As explained in Section 8.1, the model does not correct for duplicate JICs resulting from the allocation of OOP tasks to a check with one or more of the same JICs already scheduled as part of the check's block tasks. This approach is considered a conservative simplification for OOP tasks. For example, in the A10 check of Aircraft 5 the two tasks *OOP-1* and *OOP-2* are scheduled both by the model and in the actual check. Both MRIs consists of 7 JICs in total, four of which they have in common, equivalent to 2.63 hours. Where the maintenance planner takes into consideration that duplicate JICs are only executed once in the check, the model does not. Thus, yielding a 2.63 hour increase of the model's scheduled hours.

The second primary reason for differences in OOP hours concerns the associated access panel hours. Until recently maintenance planners did not account for the access panel hours in preparing the workpackage of a check. Out of the three A-checks in the validation, two do not include any access hours. More on this will be discussed in Section 8.2.4. Consequently, this results in a higher scheduled workload for the model, which does include access panel hours for all checks. Additionally, the aforementioned assumption that each OOP task consists of all its JICs results not only in duplicate JICs, but also in duplicate access panel hours.
For example, the *OOP-3* and *OOP-4* tasks which are scheduled to the A09 check of Aircraft 7 both consist of 7 JICs, equivalent to a total of 6.69 hours. However, the model schedules an additional 1.34 hours for access hours for both MRIs (a total of 2.68 hours extra) which is not accounted for in practice. However, the access panels of these two MRIs are actually identical. Thus, in practice these access hours should have only been included once (1.34 hours extra). Further still, it may well be that some of these access panels will match with the access panels for the block tasks of the check, in which case there would be further duplicate hours. Thus, the model correctly includes access panel hours which have not been included in practice, but it does not correct for duplicate access panel hours, resulting in more scheduled hours.

As has been explained in the model formulation in Section 6.2, the required labor hours of each OOP task consist of the sum of its JIC hours, access panel hours and its estimated non routine hours as predicted by the Non Routine Predictor (NRP). The inclusion of non routine hours as part of the required labor hours of each OOP also results in mismatches with the OOP hours reported in the check detail reports. The latter reports the non routine hours of the check separately from the scheduled hours of the task which only consist of the JIC hours. However, these mismatches are considered reporting differences, rather than modeling errors.

The last reason for the differences in OOP hours may be attributed to actual differences of some of the JIC hours. As with the block tasks, some of the JICs in the checks have exactly half of the hours required by the model. As explained, it would appear that this is a reporting error in the generation of the check detail reports. The matter is currently under investigation. However, as for the block task hours, the model uses the correct hours as reported in the MIS.

Thus, the differences in OOP hours result primarily from assumptions in the model design and from reporting errors. Whereas the latter are not to be attributed to the model, it has been shown that the model assumptions do yield small discrepancies in the model hours and its outcome.

8.2.3. Comparison of Modification Hours

Table 8.4 gives an overview of the comparison of modification hours. It is important to mention again that the hours of the EOs which have been compared correspond only to those EOs which occurred in both the model and in practice. Thus, only the hours of modifications have been compared. As with the OOP tasks, the hours reported in Table 8.4 are the combined hours of all modifications from the various checks.

As can be seen from Table 8.4, there is a substantial difference of approximately XMD hours between the total scheduled hours of the model and actual scheduled hours. The model schedules over PERMODS percent less modification hours compared with the actual scheduled hours. Moreover, for MISMODS% of all modifications the hours mismatch. Thus, both the frequency and size of the deviations is substantial. The largest reported deviation is MMAX hours and the average absolute deviation (including those modifications of which the hours match) is AVGMODS hours.

Table 8.4: Comparison of model and check data record modification hours

Parameter	**Value**
Total hours CDR	AEOH
Total hours model	MEOH
Total hours diff (%)	-PERMODS
Mismatches (%)	MISMODS
Max. dev. (hrs)	MMAX
Avg. dev. (hrs)	AVGMODS

If one removes the three modifications of which the hours have the largest differences, the difference in total scheduled hours would be reduced to SMPER percent, where the model schedules RMEOH hours compared with RAEOH hours in practice. Thus, the effect of these three outliers on the observed level of accuracy of the modification hours is substantial.

A logical, possible explanation for the observed differences would be the assumed modification hours. As explained in Chapter 6, the labor hours have been assumed for those modifications of which the hours were not yet specified in a Data Exchange Platform. However, analysis of all modifications included in the comparison reveals that for just two modifications the hours have been assumed. Both modifications have a preference for A-checks and as such in both cases the model assumed HAMOD labor hours required. In both cases this assumption appeared conservative, as the hours scheduled in practice were MOD1 and MOD2 respectively. For all other modifications the hours as reported in the Data Exchange Platform have been used by the model. Thus, the assumed model hours do not explain the observed differences.

Thus, the observed difference represents a difference between two of the MRO's systems, the MIS and the Data Exchange Platform, rather than a difference induced by the model design. As explained in Chapter 3, Central Engineering creates the EO which includes an estimate of the required labor hours as provided by the OAM. The direct support personnel in the maintenance hangars subsequently reviews these hours and adjusts these hours as deemed necessary. Based on the observed difference it would appear that these corrections are made in the MIS but not in the Data Exchange Platform. Provided that in practice the MIS is used for scheduling and not the Data Exchange Platform, it would appear that the hours are only corrected for in the MIS. Although in practice the MIS is used for the actual allocation of modifications to workpackages, it is still worthwhile for the correct hours to be reported in the Data Exchange Platform for analyses on the modification workload.
The difference in hours may either be resolved by adjusting the model design or through stricter observance

of reporting during the EO execute process. The model design could be extended with a feedback loop which corrects the Data Exchange Platform hours with the MIS hours if differences are observed. More importantly however, it is recommended that the MRO takes the necessary precautions to avoid differences between these two systems.

8.2.4. Comparison of Non Routine, P-Package and Access Panel Hours

This Section contains some brief remarks on the comparison of non routine hours, access panel hours and p-package hours. A comparison of the hours of each of these categories for the A-checks is shown in Table 8.5.

Table 8.5: Comparison of model and check data record (CDR) of non routine hours, p-package hours and panel hours for the A-checks considered in the validation

	A02 - Aircraft 12	A09 - Aircraft 7	A10 - Aircraft 5
NR CDR Hours (hrs)	ANR1	ANR2	ANR3
NR Model Hours (hrs)	MNR1	MNR2	MNR3
P-package hours CDR (hrs)	APP1	APP2	APP3
P-package hours Model (hrs)	PNORM	PNORM	PNORM
Panel Hours CDR (hrs)	APH	-	-
Panel Hours Model (hrs)	MPH1	MPH2	MPH3

Non Routine Hours

As explained in Chapter 6, the non routine hours of the blocks and of the OOP tasks are estimated by the MRO's non routine prediction algorithm (NRP). The non routine hours reported by the CDR encompass all non routine hours of the maintenance check. The non routine hours of the model only refer to the non routine hours of the routine block tasks. As aforementioned, the non routine hours for the OOP tasks are included in the model as part of the labor hours of the OOP tasks. The error between the non routine hours of the block tasks and the combined non routine hours of the block and OOP tasks is considered very small, provided that only a few OOP tasks are scheduled per maintenance check. Thus, the comparison made in Table 8.5, is still considered a representative comparison.

The non routine hours of the A02 and A09 checks are relatively comparable. However, there is a substantial difference between the actual and model non routine hours of the A10 check. Further analysis of the A10 check revealed that during the operational check of the emergency lighting system a fault or failure was detected which required ELNR hours to be corrected. The exact details of the nature of the fault have not been researched.

Although the non routine hours scheduled by the model are fairly similar for two of the maintenance checks as they occurred in practice, the current NR prediction algorithm is still only a trial version and as such still limited in its accuracy. As explained in Chapter 6, this is in part due to the inaccuracies in the historical records on which the non routine estimates are based, as well as due to limitations in the NRP's design.

P-Package Hours

As explained in Chapter 6 it has been assumed, in correspondence with current best practices at the MRO, that PNORM hours are reserved for line maintenance related work in the letter checks, also referred to as the P-package. For the three A-checks included in the validation study, the actual workload of the P-package tasks performed is substantially lower. However, it is known to the author that the opposite may also occur. The difficulty of the P-package is that it is often finalized mere days before the actual maintenance check takes place. The norm of PNORM hours is used as a guideline. In practice the work scheduled may be higher or lower, and the rest of the workload, including modifications and OOP tasks is adjusted accordingly. Thus, a relatively heavy workload of P-package tasks, may necessitate that some of the other tasks are removed from the workpackage. The model does not provide such flexibility and assumes the PNORM hours to be fully utilized. Nevertheless, from a planning perspective, it is correct to assume that all PNORM hours will be required when allocating the rest of the workpackage.

Access Panel Hours

As aforementioned, until recently the hours required for access panels were not included by maintenance planners in the compilation of the workpackages. Of the A-checks considered in the validation, only for the A02 check of Aircraft 12 the access panel hours were included. As can be seen from Table 8.5, the access panel hours considered by the model for the block tasks, closely match with the APH hours observed for the A02 check. The exact cause for the observed difference could not be determined, but the difference is considered sufficiently small. Based on this, it may be assumed that the model provides an accurate prediction of the access panel hours, although comparison of future checks and their respective access panel hours would provide further validation.

8.3. Accuracy and Efficiency of Task Scheduling

The following Section will discuss the task scheduling accuracy and efficiency. The task scheduling efficiency refers to the use of task interval and the available resources of the maintenance checks. To determine the task scheduling efficiency, the unused or lost interval days of all tasks allocated by the model and in practice have been compared. The task scheduling accuracy assesses the feasibility of the task schedule of the model. The feasibility was determined by assessing whether the allocation of tasks to checks satisfies all task or check requirements.

Task Scheduling Efficiency

The lost interval days of all tasks has been defined as the number of days between the due date of the respective task and its allocated execution date (check date). An overview of the total, maximum, minimum and average number of lost interval days is shown in Table 8.6. To improve the accuracy of this comparison, the analysis was conducted by considering all OOP and modification tasks for a one year period, instead of only those scheduled to the four checks included in the rest of the validation. The actual check data was based on the year 2017, whereas the model data was based on the year 2019. C-checks were excluded in this second analysis. The summed values of the lost interval days of all modification and OOP tasks over a one year period are listed in the second half of the Table.

Table 8.6: Comparisons of the lost interval days of the OOP and modification tasks scheduled by the maintenance planner and by the model respectively for a selection of maintenance checks

	Total	**Max**	**Min**	**Avg.**
Planner (days)	TAID	MAXAID	MINAID	AAVID
Model (days)	TMID	MAXMID	MINMID	MAVID
Diff (days)	TDID	MAXDID	MINDID	AVDID
Diff (%)	-15.3	-76.9	-100	-29.5

As can be seen from Table 8.6, the lost interval days are lower for the tasks scheduled by the model. The total, maximum and average number of lost interval days are 15.3%, 76.9% and 29.5% lower for the model compared to the MRO's best practices. Thus, with the current setup of the model, it would appear that the model achieves a higher task scheduling efficiency compared to the current standard at the MRO.
Contrary to the operational maintenance planner, the model considers a much longer planning horizon; is able to allocate extra maintenance checks wherever these are most needed; and is able to evaluate vast numbers of different scenarios to find the 'optimal' outcome. Subsequently, the model yields a much more efficient schedule in terms of utilizing task interval.

In addition to the interval efficiency, the efficiency with which the available resource capacity is used has also been analyzed. The resource efficiency has been assessed by comparing the total scheduled hours of a selection of A-checks with the capacity norm of the A-checks. Provided the relatively small sample group included in the validation, a larger sample group was used for this assessment as well. For a selection of 46 A-checks, the actual scheduled hours of the routine block tasks, OOP tasks, modification tasks and the actual spent hours of the non routine tasks have been summed based on the check detail reports. The same has been done for all TAC A-checks in the four year planning horizon of the model. The total hours of each check have subsequently been compared with the A-check norm of ANORM hours. An overview of the deviation from the norm for the model and for the actual values is shown in Table 8.7.

Table 8.7: Overview of the average, maximum and minimum deviation from the A-check norm of a selection of A-checks as scheduled by the model and in practice

	Model	**Practice**
average	-ADM	+ADMP
max.	+MAXDM	+MAXDMP
min	-MINDM	-MINDMP

The negative values indicate an *excess* of scheduled hours compared to the norm, whereas the positive indicate a *shortage* of scheduled hours below the norm. As can be seen from Table 8.7, on average the model schedules ADM hours above the norm, whilst in practice the scheduled hours are on average ADMP hours below the norm. Moreover, from the maxima and minima it is observed that the deviation of the scheduled hours of the model is much smaller compared with the deviation which has been observed from the scheduled hours of the CDR data. It is important to note that two of the 46 actual checks have been excluded from the results in Table 8.7 because these A-checks had been combined with a cabin-check. Of the TAC A-checks of the model, NEA checks had more than ANORM hours scheduled, equivalent to approximately EXPERA percent. Of the 44 actual A-checks, AAAN checks had more than ANORM hours scheduled, equivalent to PAAAN percent.

It is difficult to draw a conclusion on the resource efficiency from this analysis. The comparison is somewhat biased since the results of the 46 actual A-checks did not account for P-package hours, while these were accounted for in the hours of the model. Thus, the hours of the actual checks will most likely be slightly higher when also taking into consideration the P-packages. Excluding the P-packages from the model results would also not yield a fair comparison, since in practice the scheduled hours are adjusted according to the size of the P-package. Nevertheless, the analysis does indicate that in general the model has a higher utilization of the available resource capacity of its A-checks, compared with current best practices at the MRO. However, whether a maximum resource efficiency is always desirable is another matter. The higher efficiency also implies a heavier check with less margin for error in case tasks require more capacity than initially expected, which is quite likely with the occurrence of non routines.

Task Scheduling Accuracy

Although analysis shows that the model schedules tasks more efficiently than current practices, in practice there are many other considerations that constrain the scheduling flexibility, including dependencies between tasks; limited shelf life of materials; limited warranty from OEMs; extensive turn around times; power-off conditions; network restrictions; airworthiness implications; passenger convenience based priority; etc. It is important to beware that none of these other considerations are included in the present model. Hence, in addition to the scheduling efficiency, the scheduling accuracy should also be considered. The task scheduling accuracy implies the extent to which the model's proposed task schedule complies with these other scheduling considerations. Provided the great diversity and quantity of task conditions which may affect scheduling, an exhaustive analysis is difficult to conduct and beyond the intend of this model's validation. However, a simple analysis based on the more prominent task conditions revealed that for at least half of the tasks conditions apply which may potentially impact the task scheduling or which will certainly affect the task scheduling. Some examples from the sample set of maintenance checks included in the validation are elaborated upon:

- A set of five OOP tasks refer to so-called freeplay checks of the various flight control surfaces. During these tasks, the play in the flight controls has to be tested. According to the Aircraft Maintenance Manual (AMM) stable conditions are necessary for this procedure. As such, it advises to avoid any disturbances to the airframe during the execution of these tests, including noticeable winds and personnel walking through the aircraft's cabine. Provided these restrictions, the maintenance planners have scheduled these OOP tasks in a completely separate maintenance check. However, as these restrictions are unknown to the model, it has scheduled these tasks to be executed during the A10 maintenance check of Aircraft 5. Given the large number of routine and non routine maintenance tasks located in the aircraft cabin during any A-check, this task allocation could not be realized in practice as these tasks would severely interfere with the routine A-check tasks.

- One of the modification tasks concerns a modification to the air conditioning packs temperature control systems. The turn around time or duration of the modification is estimated at GTMOD hours. The maintenance planners have taken this into consideration and scheduled this task in the C01 check of Aircraft 2. Provided that the model does not include any ground time restrictions and provided that the required capacity is HMLH hours, the model has allocated this task to SB30, an extra check with A-check capacity (ANORM hours) and ground time (GTA hours). Although the capacity of an extra check would be sufficient to meet the required capacity of this task, the ground time is far less than the required GTMOD hours. Thus, it would be impossible to execute this task in an extra check as proposed by the model.
- One of the modification tasks concerns an improvement to the nitrogen generation system. The task requires the aircraft to be in a power-off condition for nearly POH hours. The maintenance planners have allocated this task to the C01 check for Aircraft 2. The model has allocated the task to the Cabin-01 check, which would greatly improve the scheduling efficiency. However, provided the power-off time, it is likely that this task would severely interfere with the required cabin maintenance tasks. Therefore, it is likely that this combination would not be feasible in practice.
- The execution of four OOP tasks related to the left engine has been combined by maintenance planners in the A09 check of Aircraft 7. Provided that these four tasks share much of the same access panels, by so doing the maintenance planners maximize efficiency. The model has allocated these four tasks to 3 different maintenance checks, to maximize the scheduling efficiency. Although this is feasible, it may be more desirable to combine these tasks. Combining these tasks will likely result in a shorter overall ground time of the aircraft as the same access panels only have to be accessed once, rather than three different times. Thus, although no hard restrictions apply as in the previous examples, other considerations may also implicate the task scheduling. Moreover, the ETOPS significance of some of these OOP tasks may further restrict the checks to which these tasks are allocated, depending on the routine block tasks in those checks.

As aforementioned, a basic analysis revealed that at least 50% of all modification and OOP tasks considered in this validation step have conditions which will or are likely to implicate the model's proposed schedule. Thus, the task scheduling accuracy of the present model is still severely limited. It should be noted however, that these additional scheduling considerations were purposely left out of the scope of the current research.

In addition to the accuracy of task scheduling, the model is also limited in its check scheduling. As explained in Chapter 6, the dates of the letter checks have been based on the due dates of their respective block tasks. As a result of this approximate method, some of the aircraft have the same check dates. For example, Aircraft 1 has been scheduled to come in for an extra check on the 15th of April 2019. On the same day, the A05 check for Aircraft 11 has been scheduled. Together, the checks require over CH hours of resource capacity which is undesirable. In practice, either of these checks would be executed earlier. Similar to the task scheduling, the check dates have been approximated with a maximum interval usage in mind and have not been optimized with respect to one another. However, as explained in Chapter 6, the optimization of maintenance check allocation to aircraft registrations was also considered out of the scope of the present model.

8.4. Discussion on Validation Results

The following Section briefly summarizes the main findings of the validation and subsequently draws a conclusion on the accuracy and efficiency of the model.

Table 8.8 provides a summarized overview of the validation outcome for the various performance indicators and task categories discussed in this Chapter. The performance is classed either as 'good', 'satisfactory' and 'poor'. The primary causes of lesser performance are also listed as 'Primary Errors', where a distinction is made between, out of scope; input errors; reporting errors and model design errors.
The outcome of the validation as indicated by the 'Performance' column in Table 8.8 shows a very mixed performance of the model. However, when considering the primary causes of lesser performance, the majority

Table 8.8: Summary of validation outcomes

Performance Indicator	Category	Performance	Primary Errors
Accuracy of Occurrence	Block tasks	Good	Inputs
	OOP tasks	Satisfactory	Inputs
	Modifications	Poor	Inputs & Scope
Accuracy of Hours	Block tasks	Good	Inputs & Reporting
	OOP tasks	Poor	Model design & Reporting
	Modifications	Poor	Inputs
	Non routine	Satisfactory	Inputs
	P-package	Good	Model design
	Access panels	Good	Reporting
Efficiency of Scheduling	Interval	Good	-
	Resources	Good	-
Accuracy of Scheduling	-	Poor	Scope

may be attributed to either input and reporting errors. These include the AMP revision which affected the accuracy of occurrence and hours for the block and OOP tasks, as well as the faulty hours for the modifications in the Data Exchange Platform. Furthermore, some of the observed non-performance was due to the limited scope of the present research. The exclusion of EOs other than modifications in the model resulted in an observed difference of EO occurrence. More importantly, the task scheduling accuracy is severely limited as a result of the exclusion significant scheduling conditions and considerations such as the available groundtime, power-off conditions, task dependencies, etc., which has resulted in unfeasible task allocations. However, non of these errors are attributable to the model itself. Faulty inputs, reporting errors and limited scope may have resulted in an unfavorable outcome, but this has not been caused by poor performance on the part of the present model. Some assumptions of the model have attributed to a different outcome compared with best practices. For example, the design of the P-packages which is fixed at PNORM hours, rather than dependent on the actual line maintenance workload. Furthermore, the assumption that an OOP task consists of all of its JICs and access panels has resulted in doubled scheduled hours.

Overall, the validation shows good performance of the present model in its current scope. Moreover, the validation has shown the need for accurate and frequent updating of all inputs required by the model. Furthermore, the validation highlights the need to expand the scope of the present model to include some of the other significant scheduling conditions, including task dependencies, groundtime requirements and others. These extensions will yield an even more valuable model for the MRO capable of actually improving the efficiency of its maintenance program and thereby contributing to the desired fleet availability.

9

Sensitivity Analysis

This Chapter will discuss the sensitivity analysis which has been performed on the model. First, the setup of the sensitivity analysis will be discussed in Section 9.1. The outcomes of the sensitivity analysis for the various parameters of interest will be discussed in Section 9.2. Lastly, Section 9.3 will provide a discussion on the observations from the sensitivity analysis and an assessment of the model's robustness.

9.1. Setup of the Sensitivity Analysis

The parameters of interest and their respective values for the base run and sensitivity runs are shown in Table 9.1. The parameters of interest have been selected based on their influence on the model penalties, and therefore on the objective function value and decision making of the model. The sensitivity analysis reveals what effects and with what magnitude these parameters of interest affect the task planning.

For *each* parameter of interest two sensitivity runs have been performed. The general approach for the sensitivity analysis has been such that for the parameter of interest a lower and higher value relative to the base value have been evaluated. *Please note that during the run of any given parameter of interest, the values of all other parameters of interest are set to their base run values.*

Of the ten parameters of interest, the first 6 in Table 9.1 are chosen for their direct effect on one or multiple penalties in the objective function value, and as such on the decision making of the model. A detailed explanation of the relation between the various parameters of interest and the model penalties is provided in Chapters 4 and 5. The last four parameters are chosen from a scheduling perspective, as these parameters are likely to have an immediate effect on the model's task schedule. The results from the sensitivity analysis will be discussed in the next Section.

To analyze the effect of the various parameters of interest on the task scheduling outcome, several key performance indicators (KPIs) are evaluated for each sensitivity run. The following KPIs have been used: shortage-; excess-; extra check-; and interval- quantities and penalties; the objective function value; and the run times. The shortage and excess hour penalties of the last optimization run are used for their respective KPIs. Thus, for an optimization consisting of 19 consecutive optimization runs, the shortage and excess penalties of the 19th run are used. These values correspond to the total penalties, provided that with each consecutive optimization run the excess and shortage penalties increase cumulatively. Moreover, the objective function value corresponds to the sum of the overall interval and extra check penalties, and the sum of the shortage and excess penalties of the last optimization run. As such, the objective function value reported as KPI represents the true total objective function value of the sum of all penalties of all consecutive optimization runs. Due to the sequential nature of the iterative solution technique, these formulations are necessary to see the true effect of a parameter of interest over all optimization runs, rather than a single run.

Table 9.1: Overview of the parameters of interest for the sensitivity analysis with their respective values for the sensitivity runs and base run

Runs	Parameter of Interest	Base Value	SA Value 1	SA Value 2
Runs 1 & 2	Labor cost per hour	CLAB	CLAB_LO	CLAB_HI
Runs 3 & 4	Check factor	1.25	5	10
Runs 5 & 6	Summer factor	1.25	5	10
Runs 7 & 8	Cost per extra check	CEXTRA	CEXTRA_LO	CEXTRA_HI
Runs 9 & 10	Cost of aircraft unavail. per hour	LOSTP	LOSTP_LO	LOSTP_HI
Runs 11 & 12	Cost factor for shortage hours	1	CLAB	2 · CLAB
Runs 13 & 14	Available extra checks per month	SBN	0	SBN_LO
Runs 15 & 16	Last optimization run to allocate extra checks	0	1	2
Runs 17 & 18	Norm of all checks	-	-10 hrs	+10 hrs
Runs 19 & 20	Number of tasks	-	-10%	+10%

9.2. Results of the Sensitivity Analysis

This Section will discuss the results of the sensitivity analysis for each parameter of interest. As explained, the effect of these parameters on the task allocation is assessed by a set of KPIs.

9.2.1. Labor Cost per Hour

The result for the sensitivity analysis of the labor cost per hour are shown in Table 9.2. The base run value for the labor cost per hour is CLAB. A lower and a higher value of CLAB_LO and CLAB_HI respectively have been used for the sensitivity analysis.

Table 9.2: Overview of KPIs for sensitivity run on the labor cost per hour

KPI	Diff. Base to SA 1 (%)	Diff. Base to SA 2 (%)
extra checks	-6.5	0.0
extra checks penalty	-48.1	44.4
total excess hours	15.3	-0.4
total excess penalty (1e5)	-3.9	49.3
total shortage hours	-3.8	-0.6
total shortage penalty	-3.8	-0.6
total interval days lost	3.1	-0.4
total interval penalty (1e5)	-41.6	46.8
obj. function value (1e5)	-24.7	46.7
run time (s)	-7.5	-7.6

As can be seen from Table 9.2, changing the labor cost per hour has big effect on the extra check, excess and interval penalties. This is analogous as all of these penalties are scaled proportionally with the labor cost per hour. As observed, a reduction of the labor costs also yields a reduction of these penalties, and vice versa. The task and extra check allocation do not appear to be affected significantly by either an increase or reduction of the labor costs, as signified by the excess hours and shortage hours and the number of extra checks. These parameters all show close resembles with the base run values, especially those of SA Run 2. The combination of large changes in the penalties without any apparent change of scheduling behavior may be attributed to the fact that the labor cost per hour scales all penalties equally, except for the shortage penalty. Thus, as all penalties are scaled equally, the weights of the penalties does not change relatively. The shortage penalty is relatively too small to induce a scheduling difference.

9.2.2. Check Factor

The results for the sensitivity analysis runs for the check factor are shown in Table 9.3. The check factor was introduced to discourage excess hours on A-checks and extra checks by amplifying the excess penalties for these check types. Provided that the initial results appeared to have little effect on the task scheduling, the values of the two sensitivity runs have been set to four and eight times the base run value. The base run value for the check factor was set to 1.25.

Table 9.3: Overview of KPIs for sensitivity run on the check factor

KPI	Diff. Base to SA 1 (%)	Diff. Base to SA 2 (%)
extra checks	-4.3	0.0
extra checks penalty	-4.3	0.0
total excess hours	15.8	25.7
total excess penalty (1e5)	288.8	613.6
total shortage hours	-6.6	-4.0
total shortage penalty	-6.6	-4.0
total interval days lost	9.4	15.9
total interval penalty (1e5)	16.2	28.7
obj. function value (1e5)	135.2	285.8
run time (s)	-8.0	-7.5

As shown by Table 9.3, the increase of the check factor has a significant effect on the excess penalty. As explained in Chapter 5 the check factor scales the slopes and steps of the piecewise linear functions which correlate the excess hours with the excess penalties. Thus, with a higher check factor a higher excess penalty is allocated for the same excess hours. However, the results in Table 9.3 show a positive correlation between the excess hours and the check factor. A further analysis of the excess hours per check type has been performed to assess the effectiveness of this parameter. An overview of the average, maximum and total excess hours per check type for the base and sensitivity runs is shown in Table 9.4.

Table 9.4: Average, maximum and total excess hours per check type for the base run and two sensitivity runs with 1.25, 5 and 10 as values for the check factor respectively

Parameter	Check	BasE	SA 1	SA 2
average	A	BRAVA	S1AVA	S2AVA
	SB	BRAVSB	S1AVSB	S2AVSB
	B	BRAVB	S1AVB	S2AVB
	C	BRAVC	S1AVC	S2AVC
maximum	A	MXBRA	S1MXA	S2MXA
	SB	MXBRSB	S1MSB	S2MSB
	B	MXBRB	S1MXB	S2MXB
	C	MXBRC	S1MXC	S2MXC
total	A	TBRA	TS1A	TS2A
	SB	TBRSB	TS1SB	TS2SB
	B	TBRB	TS1B	TS2B
	C	TBRC	TS1C	TS2C

As becomes clear from the additional results shown in Table 9.4, the check factor induces a reduction of the average and total excess hours for the A-checks. Furthermore, Table 9.4 shows a clear increase of the average, maximum and total excess hours for the Cabin (B-check) and C-checks with an increase of the check factor. Thus, an increase of the check factor yields a re-allocation of the maintenance tasks away from the A- and extra checks and to the Cabin and C-checks. This re-allocation of tasks is further affirmed by the higher observed interval days remaining in Table 9.3. Based on the large quantity of excess hours still scheduled to the A- and extra checks, it may also be concluded that the effectivity of the check factor is limited. It is recommended to conduct more sensitivity runs to determine the check factor value which will yield a more desirable outcome.

9.2.3. Summer Factor

Similar to the check factor, the summer factor increases the excess penalty incurred for all maintenance checks with excess hours during the summer season. The summer factor was set to 1.25 on the base run. Just as for the check factor, in the two sensitivity runs for the summer factor the values were set to 5 an 10 respectively. A comparison of the base and sensitivity run outcomes is shown in Table 9.5.

Table 9.5: Overview of KPIs for sensitivity run on the summer factor

KPI	Diff. Base to SA 1 (%)	Diff. Base to SA 2 (%)
extra checks	-15.2	-23.9
extra checks penalty	-15.2	-23.9
total excess hours	47.0	54.1
total excess penalty (1e5)	214.1	60,104
total shortage hours	-14.6	-21.9
total shortage penalty	-14.6	-21.9
total interval days lost	27.3	31.0
total interval penalty (1e5)	36.3	43.7
obj. function value (1e5)	109.5	26,795
run time (s)	-5.6	-8.1

As can be seen from Table 9.5, a strong positive correlation is observed between the excess penalty and the summer factor. Although this is largely because the summer factor is designed to amplify the excess penalty, Table 9.5 also shows an increase in the excess hours with an increase of the summer factor. A further analysis reveals that the summer factor affects the allocation of excess hours, as is shown in Table 9.6.

Table 9.6: Comparison of average, maximum, total excess hours, percentages of summer checks with excess and percentages of summer checks with more than XEH hours of excess for the sensitivity runs and base run

Parameter	Base Run	SA 1	SA 2
Percentage of summer checks with excess	BP	S1P	S2P
Average excess	BA	S1A	S2A
Total excess	BT	S1T	S2T
Maximum excess	BM	S1M	S2M
Percentage of checks with >XEH hrs excess	BPX	S1PX	S2PX
Percentage of extra checks in summer	BPS	S1PS	S2PS

The results shown in Table 9.6 show a reduction of the total and average excess hours allocated to maintenance checks in the summer season for an increasing summer factor. Moreover, although the percentage of checks with excess remains relatively constant, the percentage of maintenance checks in the summer with more than XEH hours of excess clearly decreases for an increasing summer factor. Thus, the summer factor appears to be effective in redistributing tasks such that fewer maintenance checks during the summer season have large excess hours. This re-allocation is affirmed by the positive correlation between the interval days remaining and the summer factor, as well as the negative correlation between the shortage hours and the summer factor, as observed in Table 9.5.

Table 9.5 also shows the negative correlation between the number of extra checks and the summer factor; overall fewer extra checks are allocated with an increasing summer factor. As explained in Chapter 5, the summer factor increases the extra check penalties of all extra checks allocated to the summer season. As further shown in Table 9.6, the percentage of extra checks in the summer season decreases for an increasing summer factor. However, provided that XSB of the SBN available extra checks have check dates in the summer season, the higher summer factors not only yield a reduction of the number of extra checks in the summer season, but an overall reduction of the allocation of extra checks.

9.2.4. Cost per Extra Check

The cost per extra check refers to the penalty imposed when the model allocates an extra maintenance check. The results of the sensitivity runs conducted for the cost per extra check are shown in Table 9.7. One sensitivity run was conducted at half of the base run cost, and the second run evaluates double the base costs.

Table 9.7: Overview of KPIs for sensitivity runs on the cost per extra check

KPI	Diff. Base to SA 1 (%)	Diff. Base to SA 2 (%)
extra checks	13.0	-34.8
extra checks penalty	-43.5	30.4
total excess hours	-13.7	48.6
total excess penalty (1e5)	1.7	81.2
total shortage hours	11.2	-25.9
total shortage penalty	11.2	-25.9
total interval days lost	-6.5	9.5
total interval penalty (1e5)	-2.9	9.5
obj. function value (1e5)	-4.5	42.7
run time (s)	-8.1	-8.0

The variation of the extra check cost has a clear impact on the number of extra checks allocated by the model. A reduction of the cost yields a modest increase, but an increase of the cost appears to result in quite a considerable reduction. The changes in the number of extra checks allocated result in a different allocation of maintenance tasks over the available resource capacity. With fewer extra checks available due to an increased cost, there is a substantial increase of the number of checks with excess hours, as well as a noticeable reduction of the number of checks with shortage hours. Additionally, more tasks are executed further ahead of their respective due dates in order to avoid large excess penalties, as can be seen from the higher number of interval days remaining. The opposite is observed in case the cost per extra check are reduced. Thus, the cost per extra check also have a profound effect on the planning.

9.2.5. Cost of Aircraft Unavailability per Hour

The aircraft unavailability cost per hour represent the height of the steps in the piecewise linear functions for the excess penalties. The base value is set at LOSTP. An elaborate explanation on how this value was defined can be found in Chapter 5. The effect of this parameter on the model outcome has been assessed by evaluating half and double the base value. The results of these sensitivity runs are shown in Table 9.8.

Table 9.8: Overview of KPIs for sensitivity runs on the cost for aircraft unavailability

KPI	Diff. Base to SA 1 (%)	Diff. Base to SA 2 (%)
extra checks	-4.3	0.0
extra checks penalty	-4.3	0.0
total excess hours	1.6	0.0
total excess penalty (1e5)	-0.9	17.8
total shortage hours	-5.6	0.0
total shortage penalty	-5.6	0.0
total interval days lost	2.9	-2.0
total interval penalty (1e5)	3.2	1.2
obj. function value (1e5)	0.4	8.4
run time (s)	-7.0	-7.8

As can be seen from Table 9.8 for all KPIs the differences between the sensitivity runs and the base run are relatively small. Particularly, SA Run 2 appears to be nearly identical to the base run, with exception of the total excess penalty. Further analysis showed that SA Run 2 has four checks which exceed EH1 excess hours and are therefore penalized extra by the aircraft unavailability penalties. The base run has only ABC checks which are penalized with the base value for aircraft unavailability penalty. Furthermore, the ABCD checks in SA Run 2 are all scheduled during the summer season, thus the additional seasonality factor is also applied to the aircraft unavailability penalty. For the base run, only two checks are scheduled in the summer season.

What remains to be answered is why SA Run 2 has a larger number of checks which are penalized with the aircraft unavailability penalty. Theoretically, the opposite would be expected, since a higher penalty should discourage scheduling behavior that yields more penalty. The difference may be attributed to the non-integral approach of model's consecutive optimization runs. A detailed explanation of the effect of the iterative solution technique on the task scheduling has been discussed in Chapter 7. As is explained there, this approach can cause the model to make scheduling decisions during the initial runs, which will induce large excess hours in subsequent optimization runs and which cannot be corrected for anymore in the later optimization runs. The same has occurred in SA run 2; a relatively small scheduling difference in SA run 2 compared with the base run has resulted in the incurrence of an additional aircraft unavailability penalty.

Apart from highlighting one of the model limitations, it may be concluded from the sensitivity runs that the aircraft unavailability penalty has little effect on the scheduling outcome of the model.

9.2.6. Cost Factor for Shortage Hours

The cost factor for shortage hours is set to unity for the base run. Consequently, the shortage hour are penalized least heavily. However, too much unused available resource capacity is also undesired. In the sensitivity runs, the cost factor for the shortage hours has been set to equal and double the CLAB penalty for excess hours. Thus, the results of the sensitivity runs show the effects on task scheduling when shortage and excess are weighted equally and when shortage is weighted more heavily than excess. An overview of the results is shown in Table 9.9.

Table 9.9: Overview of KPIs for sensitivity runs on the cost for shortage hours

KPI	Diff. Base to SA 1 (%)	Diff. Base to SA 2 (%)
extra checks	-54.3	-60.9
extra checks penalty	-54.3	-60.9
total excess hours	100.2	121.4
total excess penalty (1e5)	184.0	258.9
total shortage hours	-31.9	-33.2
total shortage penalty	2964.7	5912.7
total interval days lost	13.6	29.7
total interval penalty (1e5)	17.7	17.2
obj. function value (1e5)	144.4	236.7
run time (s)	-7.9	-3.9

As expected, the higher penalty for shortage hours yields a reduction of the total shortage hours scheduled. The correlation between the excess and shortage penalties is also reflected in the total excess and shortage hours. The excess and shortage hours are nearly identical for equally weighted penalties in SA Run 1. In SA Run 2, where the shortage hours are weighted more heavily than the excess hours, the results show less shortage than excess hours. The non linear segments of the excess penalty imply that the excess penalties are proportionally still bigger than the shortage hour penalties. The redistribution of tasks to find the optimal balance between shortage and excess is further affirmed by larger number of interval days remaining and corresponding penalties.

The increased shortage hour penalty also has a strong effect on the number of extra checks allocated. As has been observed for the base run results in Chapter 7, the available resource capacity is often not used completely. As the shortage hours are now penalized more heavily than the excess hours, the total penalty incurred by an extra check apparently makes the extra checks a less viable option.

9.2.7. Available Extra Checks per Month

It has been assumed that there are SBN extra maintenance checks available each month for the airline's wide-body fleet. A detailed explanation on this assumption is provided in Chapter 5. In the future at least some of these slots will be used for aircraft that will be added to the airline's current wide-body fleet. As such this number is likely to reduce in the future. Thus, the sensitivity runs assume a pessimistic scenario in which only

SBN_LO and even no extra checks at all will be available. The latter scenario provides a schedule in which only the letter checks with their current norms are available for scheduling of OOP tasks and modifications. The results of the sensitivity runs are shown in Table 9.10.

Table 9.10: Overview of KPIs for sensitivity runs on the number of extra checks available per month

KPI	Diff. Base to SA 1 (%)	Diff. Base to SA 2 (%)
extra checks	-100.0	-28.3
extra checks penalty	-100.0	-28.3
total excess hours	277.8	45.5
total excess penalty (1e5)	6,605,670	74.8
total shortage hours	-33.6	-19.6
total shortage penalty	-33.6	-19.6
total interval days lost	44.5	21.0
total interval penalty (1e5)	31.0	16.0
obj. function value (1e5)	2,943,077	37.1
run time (s)	-99.6	-6.1

As can be seen from Table 9.10, not having any extra checks available has a profound impact on the results. The excess hours are substantially larger, yielding a significant increase of the excess penalties and subsequently of the objective function value. Of the TLC letter checks, YEX1 checks have excess hours, equivalent to PYEX1%. Moreover, YEX2 checks have more than XEH hours of excess; YEX3 checks have more than EH1 excess hours; and YEX4 have more than EH3 hours; equivalent to PYEX2%, PYEX3% and PYEX4% of all maintenance checks respectively. The last group with more than EH3 excess hours is responsible for the excessive penalties. As explained in Chapter 5, above EH3 excess hours the excess penalty scales with the slope of the steps in the piecewise linear function (PWLF), which are equivalent to PSX million per hour. With the maximum excess hours at UDXYZ hours, this results in the large excess penalties observed. The significantly lower run time of SA Run 1 is the result of CPLEX reaching the desired optimality gap of 0.01%. The large objective function value of SA Run 1 yields a larger absolute optimality gap, which results in an earlier cut-off of the optimization. A similar case has been described in Chapter 7, where an anomaly resulted in significant excess hours which subsequently compromised the optimality of the model's solution. As has been recommended there, this sensitivity run affirms again the need for a reformulation of the excess penalty's PWLF for exceptionally high excess hours. The results of SA Run 1 clearly show that with the current norms, the letter checks alone do not suffice to accommodate the workload on the wide-body fleet and consequently additional measures are required, such as extra maintenance checks.

A comparison of SA Run 1 and SA Run 2 shows the profound impact of the extra check availability. With just SBN_LO extra checks available per month, the excess hours are still 45% higher than the base run, but significantly lower than the case without any available extra checks.

9.2.8. Last Optimization Run to Allocate Extra Checks

In the base run it has been assumed that extra checks may only be allocated in the initial optimization run, to avoid under-utilization of the extra checks. However, it has been observed that the workload of the tasks allocated during the subsequent optimization runs often yields excess hours. Thus, a sensitivity run has been conducted in which the model is permitted to also allocate extra maintenance checks during the second and third optimization runs respectively. It is important to note that for these consecutive runs only those extra checks were available which remained after the allocation in the previous runs. The results of the sensitivity runs are shown in Table 9.11.

As may be expected, by increasing the number of optimization runs during which extra checks may be allocated, a larger number of extra checks are allocated by the model. Moreover, with more extra checks the excess hours are reduced and the shortage hours increased. There is simply more available capacity to allocate maintenance tasks to and inherently also more unused available capacity.

Table 9.11: Overview of KPIs for sensitivity runs on the last optimization run to allocate extra checks

KPI	Diff. Base to SA 1 (%)	Diff. Base to SA 2 (%)
extra checks	6.5	30.4
extra checks penalty	6.5	30.4
total excess hours	-5.2	-31.5
total excess penalty (1e5)	7.9	-21.5
total shortage hours	5.4	26.2
total shortage penalty	5.4	26.2
total interval days lost	-2.4	-4.1
total interval penalty (1e5)	2.0	-1.7
obj. function value (1e5)	5.2	-6.8
run time (s)	-4.9	-8.0

The lower objective function value of SA Run 2 would suggest that the solution is more favorable than the outcome of the base run. Provided that the extra check penalty is equivalent to the penalty for XEH excess hours, further analysis was conducted to see if the number of checks with less than XEH hours of excess has indeed increased with the larger number of extra checks allocated. Moreover, the total unused available capacity has been investigated. An overview of these results is shown in Table 9.12.

Table 9.12: Comparison of unused capacity of extra maintenance checks and excess hours for all maintenance checks

Parameter	Base	SA 1	SA 2
Percentage of all checks with <XEH hrs excess	BRX	S1X	S2X
Total unused capacity of extra checks	TBRSH	TS1SH	TS2SH
Avg. unused capacity of extra checks	AVBRSH	AVS1SH	AVS2SH
Max. unused capacity of extra checks	MXBRSH	MXS1SH	MXS2SH

The results in Table 9.12 affirm that the number of maintenance checks with less than XEH excess hours increases when more extra checks are allocated. However, the results also show an increase for the unused available capacity of the extra checks. Although the maximum unused capacity is constant, the total and average unused available capacity of the extra maintenance checks increase. Thus, although the allocation of extra maintenance checks at later optimization runs alleviates the excess hours of maintenance checks, it also implies less efficient use of the available resource capacity of these extra checks. Further analysis would be required by maintenance planners to assess if the reduced pressure on the capacity of maintenance checks is worth the additional downtime incurred by the extra maintenance checks that would be required. An alternative option would be to increase the norm of the maintenance checks. The sensitivity analysis conducted on the norm of the maintenance checks will be discussed in the next section.

9.2.9. Norm of Maintenance Checks

As already noted in the previous Section, an alternative to scheduling extra maintenance checks would be to increase the norm of the maintenance checks; to provide more available capacity per maintenance check without increasing the downtime of the aircraft. To study the effect of the norms a sensitivity analysis was conducted in which the norms of all maintenance checks were increased and decreased by ten hours respectively. The results of the sensitivity runs are shown in Table 9.13.

As may be expected, an increase in the norm (SA run 2) results in fewer extra maintenance checks being allocated. Both measures fulfill the same purpose of providing more available capacity. With more capacity provided by the increased norm, the need for extra checks decreases. The opposite is also observed (SA Run 1). Similarly, with a higher norm per maintenance check, the shortage hours increase whilst the excess hours decrease. Interestingly, the excess penalty also increases. This may be attributed to the non-integrality of the optimization runs and the fact that extra checks are only allocated on the initial optimization run. On the initial optimization run, the UDSB allocated extra checks and higher norms will be sufficient for the required capacity. However, on the subsequent optimization runs, this capacity no longer suffices, resulting in excess hours. More specifically, compared to the base run which had XYZ maintenance checks with more than EH1

Table 9.13: Overview of KPIs for sensitivity runs on the norms of the maintenance checks

KPI	Diff. Base to SA 1 (%)	Diff. Base to SA 2 (%)
extra checks	8.7	-26.1
extra checks penalty	8.7	-26.1
total excess hours	58.9	-12.2
total excess penalty (1e5)	111.0	10.2
total shortage hours	-23.4	15.3
total shortage penalty	-23.4	15.3
total interval days lost	14.9	-7.8
total interval penalty (1e5)	6.8	-4.7
obj. function value (1e5)	52.8	0.2
run time (s)	-7.2	-7.3

hours of excess, SA Run 2 has WXYZ checks with more than EH1 hours of excess. The latter results in an increase of the excess penalty, despite the net excess hours being lower. Thus, it is apparent that an increase of 10 hours is insufficient to compensate for the NUMSB extra checks which have been dismissed. Lastly, the results show that tasks are allocated more closely to their due dates with higher norms as indicated by the fewer lost interval days, and vice versa. Again, this is expected as there is simply more or less capacity available within the vicinity of the due date of the task.

9.2.10. Number of Maintenance Tasks

The number of tasks determines the workload of the aircraft and fleet collectively and is subject to constant change. A change of the workload may come from new modification tasks or newly defined or discontinued OOP tasks, etc. To assess the effect of such workload changes to the scheduling outcome, two sensitivity runs (SA Run 1 and SA Run 2) have been conducted in which 10% fewer and 10% extra tasks need to be scheduled in comparison with the base run. A random number generator has been used to randomly select tasks from the base run workload to either be removed or duplicated in the workload. The random number generator was applied to the workload of each aircraft individually. This ensures that the workload is increased or decreased proportionally; since the workload of the older wide-body aircraft is larger compared with the workload of the more recent aircraft. An overview of the overall workload alterations is shown in Table 9.14. The results of the sensitivity runs are shown in Table 9.15.

Table 9.14: Overview of quantity and hours of modification and OOPs tasks removed (SA Run 1) and added (SA Run 2) as selected per random number generation

Parameter	SA Run 1 - Removed	SA Run 2 - Added
Qnty. modifications	QM1	QM2
Qnty. OOPs	QO1	QO2
Avg. hours modifications	MAV1	MAV2
Avg. hours OOPs	OAV1	OAV2
Total hours modifications	MT1	MT2
Total hours OOPs	OT1	OT2

As can be seen from the distribution of the quantity of tasks and their respective hours in Table 9.14, the number of tasks added and removed is very similar between the two sensitivity runs. There are slightly more OOP tasks added or removed in comparison with the modifications. On the other hand, the hours of the modifications are larger compared with the OOP tasks.

As can be seen from Table 9.15, the change in workload affects the number of extra checks allocated as there is simply more or less resource required. NSBX1 fewer extra checks are allocated for a smaller workload, compared to NSBX2 more extra for a larger workload in SA run 2. This is at least in part due to the fact that the increase of the workload was slightly smaller than the reduction of the workload.
As expected, the workload also affects the excess and shortage hours; a larger workload yields more excess

Table 9.15: Overview of KPIs for sensitivity runs on the number of maintenance tasks

KPI	Diff. Base to SA 1 (%)	Diff. Base to SA 2 (%)
extra checks	-15.2	2.2
extra checks penalty	-15.2	2.2
total excess hours	-15.4	27.6
total excess penalty (1e5)	-5.2	56.9
total shortage hours	2.8	-14.4
total shortage penalty	2.8	-14.4
total interval days lost	-10.7	15.0
total interval penalty (1e5)	-8.5	13.5
obj. function value (1e5)	-7.5	31.2
run time (s)	-8.0	-8.1

and less shortage, and vice versa. Moreover, as the workload of tasks increases, so do the lost interval days as tasks are scheduled further from their respective due dates, and vice versa.

Thus, the fluctuations in the workload indeed affect the task scheduling of the model. Although based on the results from the sensitivity runs which have been conducted, the changes to the planning outcome are moderate. It would appear that the changes in the KPIs are somewhat proportional to the percentages by which the tasks have been increased and decreased.

9.3. Discussion on the Sensitivity Analysis

The sensitivity analysis performed in this study has been limited due to time constraints on the study. As a result, the analyses from the current study only provide an indication of the effect of a parameter of interest on the KPIs, but it does not necessarily show the rate of change. It is therefore recommended that a further analysis be conducted which is more extensive than the one in the current study. However, several relevant observations and conclusions may be drawn on the basis of the results of the analysis performed in this study.

The sensitivity runs have shown a strong correlation between the model parameters and the model outcomes. With the exception of the labor cost per hour and the aircraft unavailability cost, all other parameters of interest which have been have a noticeable impact on the KPIs and subsequently also on the actual task and check allocation. The study affirms the importance of the relation between the weights of the different penalties, and the profound impact which changing these relations has on the scheduling outcome.
Based on the observations in this analysis, it may be concluded that the present model is highly susceptible to changes of the input conditions or assumptions. However, this sensitivity may be expected when considering the formulation of the model and the various interdependencies between the different variables and parameters. More important however is what the implications of the model's sensitivity are for its usability by the MRO?
As has been explained in Chapter 3, the MRO's Tmin planning policy prescribes that the workpackage of an A-check is finalized two weeks prior to the maintenance check. For the C-checks the workpackage has to be finalized six weeks prior to the C-check [3]. The remaining time period is used for the preparations of the maintenance check (material, equipment orders) based on the confirmed workload. However, as aforementioned, in practice the workload is constantly changing. Thus, in practice the model would need to efficiently process any workload changes, but it would also need to provide a stable outcome for the upcoming maintenance checks in accordance with the Tmin policy. Currently, the high sensitivity of the model would likely yield a slightly different scheduling outcome with each input change and model run. To make the model more usable for the MRO, the model would need to be extended. A possible solution would be to include 'earmarking' features which fix the allocation a task to a check according to the input of a maintenance planner. This would prevent the model from reallocating the tasks assigned to upcoming check, whilst still allowing it to optimize the task schedule of all other tasks and maintenance checks.

The sensitivity analysis also reaffirms the need for additional resources given the current workload of the fleet and the norms defined for the letter checks. As has already been shown in Chapter 7, the letter checks by themselves currently do not provide the available resource capacity sufficient to meet the demand of the maintenance task workload. The two measures studied in the sensitivity analysis, allocating extra maintenance checks and increasing the norms of the maintenance checks, both have proven to be effective and necessary measures to reduce the excess hours on the maintenance checks. A further study is recommended into both measures using the model to determine the desired balance between the two measures given the workload on the fleet. This trade-off study should take into consideration the actual costs of extra downtime on the aircraft induced by the extra checks, the cost of extra maintenance personnel required when increasing the norms and any dependencies between the various tasks, currently not included in the model.

In addition to the observed sensitivity of the model to change, the sensitivity study for the number of available extra checks highlighted once again the need to change the formulation of the post-slope of the piecewise linear function that control the excess penalty. It is recommended to formulate this penalty more proportional to the other penalties such that the optimality of the model's solution is maintained even in case of large excess hours. Furthermore, both the sensitivity study on the norms of the maintenance checks, as well as the study on the cost of aircraft unavailability showed the need for an integrated modeling approach, rather than the current iterative solution technique. It is therefore recommended to consider an alternative approach which allows for accurate scheduling of the recurring OOP tasks in a single optimization run. Such an approach will produce a more optimal task and check schedule.

The sensitivity analysis has stressed the importance of the proper definition of the weights of the model penalties. More specifically, the importance of the correct ratios between the penalty weights as these affect the model's decision-making and therefore the task schedule. However, in practice the decision-making by maintenance planners in producing a maintenance schedule is not specifically quantified but rather driven by experience and 'know-how'. Consequently, in formulating mathematical models similar to the model designed in the present study, one has to capture non-quantified knowledge into quantified data. This process requires human interpretation which makes it inherently susceptible to errors. Thus, the decision making of an operational model, which is designed to automate and optimize operations without human interference, is still subject to human interpretation in its design. It is critical therefore to the usability of the model that a thorough study be conducted to determine and further validate these weights. This may be achieved by conducting a sensitivity analysis after formulating the model and fine tuning the weights iteratively, and alternatively or additionally by conducting a thorough study which is aimed at understanding and quantifying human decision making and trade-offs.

10
Conclusion

This report has described the research of an integrated and automated task scheduling optimization model. The primary research question has been defined as: *What is the added value of the proposed integrated and automated task scheduling optimization model applied to a tactical planning horizon for all base maintenance of the wide-body fleet of a large European airline?* This Chapter aims to answer that question by drawing conclusions based on the observations which have been discussed elaborately in the report, according to the sub-questions formulated in the Introduction. A final conclusion will be drawn in answer of the main research question. A brief review of the academic relevancy of the present research will be provided first.

Academic Relevance of the Research

The present research contributes to the applied task scheduling and resource allocation literature. Contrary to existing literature on aircraft maintenance task scheduling it includes all base maintenance tasks types: routine, non routine and modifications. Moreover, the emphasis of the present model on resource capacity efficiency has resulted in a very different task scheduling model compared with the conventional task scheduling models found in literature. The soft resource capacity constraints and corresponding penalties have been based on model formulations found in conventional task scheduling literature, but have been adjusted to meet the needs of the MRO.

What insights does the task scheduling model provide in terms of the required and available resource capacity?

The model produces an optimized, integral overview of all significant base maintenance tasks and checks over a tactical planning horizon for the airline's wide-body fleet. The schedule is task specific and includes due dates and capacity constraints. It provides a suggestion on how to divide the complete workload of each aircraft over the various letter checks according to their available capacity and the due dates of the tasks. Up until now such an integral, task specific schedule for a tactical planning horizon did not exist within the MRO.

The schedule presented in this research has shown that with the current capacity norms for letter checks (2018) and the known workload, there will be a shortage of available capacity within the period from June 2018 till August 2022. Thus, additional measures are required by the MRO to avoid maintenance induced disruptions. In turn, the model may be used to assess the effectivity of any proposed measures by the MRO planning policy makers, such as an increase of the capacity norm for the letter checks or the allocation of extra maintenance checks.

Furthermore, compared with the current non-task specific tactical maintenance planning, the model's schedule gives a much more accurate indication of the actual required resource capacity. Its schedule further provides maintenance planners and policy makers with suggestions on where and when to schedule extra resource capacity over the full four year planning period to accommodate the complete maintenance workload. This will help the MRO reduce the potential negative impact of poor planning on the airline's fleet availability.

Additionally, the model's schedule is considered particularly valuable for the MRO to schedule modification tasks. Unlike other maintenance tasks, these are not specifically bound to any check type; generally have due dates quite far into the future; and have so far provided a larger workload than expected. It is anticipated that the model's suggested schedule will help the MRO to implement modifications more quickly, yielding operational and financial benefits for the airline.

Lastly, the insights provided by the model may also be valuable in discussions between the MRO's planning policy makers and the operator. The task specific schedule implies that the MRO can provide the operator with a more firm commitment of when certain modifications will be completed on their aircraft. Moreover, it may assist the MRO in negotiations with the operator on accepting future modification workloads.

How accurate is the planning provided by the task scheduling model, in terms of the accuracy of occurrence of maintenance tasks?
As shown by the model validation in Chapter 8, the accuracy of occurrence of maintenance tasks is excellent provided the model inputs are up-to-date. All routine maintenance tasks for the A- and B-checks as contained in the AMP are included in the model, as well as all active modification tasks.
However, the model does not include EOs other than modifications. Furthermore, the routine block tasks of the C-checks are not included in the present model. Moreover, the model does not include routine tasks smaller than A-check type tasks, e.g. line maintenance tasks, engine and APU changes and deferred defects.

How accurate is the planning provided by the task scheduling model, in terms of the accuracy of the scheduled labor hours of the maintenance tasks?
The accuracy of the scheduled labor hours of all routine tasks is good and consistent with the hours in the MRO's MIS, provided that the inputs are kept up to date. The assumption that the labor hours of each OOP task are based on all its JICs and access panels will sometimes result in duplicate JIC and access panel hours. However, these errors are considered relatively small and conservative.
An error was observed for the scheduled labor hours of the modifications. Analysis revealed that this error is the result of a discrepancy between the modification information contained in the Data Exchange Platform and the MIS. Currently no feedback loop exists between the two systems which has allowed this discrepancy to go unnoticed.

The MRO's trial-version of the non routine predictor (NRP) was used to estimate the workload of non routine tasks for all routine tasks scheduled by the model. However, validation revealed that in its present form the NRP does not provide a satisfactory prediction of the non routine workload. Much of this inaccuracy is caused by the corrupted input data on which the NRP predictions are based. Without a proper correction of these input records the current NRP cannot provide an accurate enough estimate of the non routine workload.

How accurate is the planning provided by the task scheduling model, in terms of the accuracy of task scheduling?
The task scheduling accuracy of the present model is good when based on the scope of the present research, but insufficient when considering a broader, more desirable scope. Given the limited scope of the present research, the model produces a task schedule which satisfies all due date constraints and which optimizes for the most efficient use of available resource capacity. However, in addition to the due date and resource capacity constraints, there are numerous additional maintenance task and check related conditions which affect the feasibility of a task schedule, including: elapse times; skills; power-off conditions; ETOPS regulations; material availability; task dependencies; etc. These conditions are ignored in the present model, resulting in practically unfeasible task allocations. Hence, although the model performs well within its limited scope, its capabilities need to be extended for it to produce an optimized and completely feasible task schedule.

Moreover, the current model's simplified approach to estimate the dates of the maintenance checks will sometimes yield unfeasible schedules where, for example, two different aircraft are required to have an A-check on the same day. An optimization of the maintenance check allocation was considered outside of the scope of the present research. However, the addition of such an optimization would further enhance the scheduling accuracy and operational consistency of the model. An extension of the MRO's current maintenance check optimization algorithms could provide the model with the desired optimized maintenance check allocation.

How efficient is the planning provided by the task scheduling model, in terms of the efficiency of using the available task interval and resource capacity?

Analysis has shown that the interval efficiency of the model is superior to the current best practices at the MRO. The analysis showed that on average the task allocation of the model yields on average PAI% fewer lost interval days compared with the MRO's current standards. Similarly, analysis has shown that the efficiency with which the model utilizes the available resource capacity is superior to current best practices at the MRO. Results show that on average the model schedules slightly above the available norm, maximizing the usage of available resources. However, it should be noted that the current schedule of the model does not consider the task dependencies and conditions as has been explained for the previous sub-question. Furthermore, the relatively heavier maintenance checks scheduled by the model also imply that there is less margin for error in case of extra required capacity during the check itself.

How is the performance of the task scheduling model, in terms of the optimality of its solution?

The present model achieves an optimality gap of 3.95% on the initial optimization and optimality gaps of 0.01% on the subsequent runs. The optimality gap on the initial run implies that there is still an appreciable deviation from the optimal feasible solution. This optimality gap has been accepted due to time constraints but should ideally be reduced by extending the model's run time.
However, more importantly is the limited optimality caused by the model's current design. The model iterative solution technique implies that the model's solution is inherently sub-optimal. With its current design the model optimizes each run singularly, which yields a sub-optimal overall scheduling result. Essentially, the model is constantly making decisions based on partial information. It has been observed that because of this, some maintenance checks are allocated excessive hours while other checks still have sufficient resource capacity available. Redefining the model formulation to a single, integrated optimization in which all maintenance tasks and checks for the complete planning horizon are considered at the same time will likely produce an optimal task schedule.

How is the performance of the task scheduling model, in terms of its robustness?

The sensitivity analysis of the model and study of the results have shown that the model is sensitive to even relatively small changes to the model inputs. This necessitates some additional features to be build into the model to enhance its practical use as a planning tool for the MRO. One such feature is the ability for the model users to "earmark" tasks to maintenance checks. This will ensure that the tasks of upcoming maintenance checks are not re-allocated by the model, in accordance with the Tmin planning policies of the MRO.
Furthermore, the model is currently not equipped to handle large excess hours. In such cases the allocated excess penalties are disproportional with the rest of the penalties imposed by the model. It has been shown that this subsequently also has a negative effect on the optimality of the model outcome. Thus, the excess penalty of the model must be reformulated to ensure that the model can deal with singularities, such as very large excess hours.

How is the performance of the task scheduling model, in terms of its run times?

The data preparation process of the present model requires approximately 4.5 hours to produce all the input data required for the optimization. Following this data preparatory process the optimization runs will require approximately 2.5 hours to complete. Although the practicality of the model will depend on its exact use and frequency of use, a total run time of approximately 7 hours is considered satisfactory.

Having considered the answers to all previous sub research questions, the main research question remains to be answered.

What is the added value of the integrated and automated maintenance task scheduling optimization model applied to a tactical planning horizon for all base maintenance of the wide-body fleet of a large European airline?

Up until now, the task-specific insights provided by the model did not exist, making the model and its insights a valuable contribution for the MRO. The task specific tactical maintenance planning produced by the model shows the actual relation between the required and available resource capacity for the coming years. Overall, the model's insights will benefit the tactical maintenance planning of the MRO and as such may contribute to cost reductions for both the MRO and the airline.

However, limitations in the model design and scope currently restrict the added value of the present model. The outcome of the model is limited to an initial planning suggestion, rather than a maintenance schedule which can fully be depended upon. Nevertheless, the present research has highlighted the potential of an integrated, automated and fully optimized task schedule for the MRO. It is therefore strongly recommended that the necessary model design improvements and capability extensions be considered and implemented, such that the full potential of an automated, integrated and optimized maintenance planning model may be used to benefit the efficiency of the MRO's maintenance planning and subsequently improve the availability and profitability of the airline's fleet.

11

Recommendations

The previous Chapter has discussed the primary conclusions drawn based on the results and analyses regarding the value of the model to the MRO and the airline. This Chapter will make further recommendations based on these conclusions which, when implemented, will enhance the value of the model even further.

- It is recommended that the current model formulation be extended to also include ground time optimization. The current model is limited to a resource capacity optimization problem only. Although essential, resource capacity by itself is insufficient to produce a feasible task schedule. The addition of task dependencies and task conditions (power-off, ETOPS regulations, work-order costraints, etc.) as ground time constraints have to be taken into consideration. Potentially, the present model could be extended into a bi-objective model which aims to minimize the makespan (turnaround time) of all tasks, whilst also maximizing the resource capacity efficiency.
- It is recommended that the iterative solution technique producing consecutive optimization runs is replaced by an integrated optimization method. Instead of the consecutive optimization runs where the successive runs use the outputs of previous runs as inputs, it is recommended that the model allocates all tasks and checks in a single optimization run. This will ensure a more optimal task and check allocation.
- It is recommended that the excess penalty for large excess hours is reformulated. More specifically, it is recommended that the post-slope of the piecewise linear function which control the excess penalties be redefined to the same order of magnitude as the other penalties of the model. This will ensure that even in the event of exceptionally large excess hours, the model's optimality is not compromised.
- It is recommended that a more accurate prediction algorithm for the non routine workload be developed. Although the algorithm itself was outside of the scope of the present study, its accuracy affects the accuracy of the task schedule of the model.
- It is recommended to include the routine maintenance tasks of the C-checks such that the complete base maintenance workload is included in the model. Moreover, this will yield the actual remaining available resource capacity of the C-checks for other tasks such as the modifications rather than the assumed capacity currently used in the model.
- It is recommended that a check date optimization algorithm be developed and build into or connected to the task scheduling optimization model. Such algorithms are presently already under developed and partially already in use at the MRO. The scope of these existing algorithms is too limited for the requirements of the model. The addition of optimized check dates will ensure a more accurate schedule which is consistent with operational constraints.
- To enhance the practical use of the model in helping maintenance planners allocate maintenance tasks to checks, it is recommended that an earmarking feature is build into the present model (or something similar) which fixes the task schedule of upcoming maintenance checks. This will prevent the model

from re-allocating the tasks of upcoming maintenance checks elsewhere and would make the model conform to the Tmin planning policy at the MRO.

- It is recommended that the model is run without time constraints to determine the convergence rate up until the desired optimality gap of 0.01 percent. Based on the outcome of this run, it is recommended to perform a trade-off study in which the run time is weighted with the the practicality of the model and with the increased optimality of the solution. Depending on the outcome of this study it may be desirable to extend the solver's present run time constraint.
- It is recommended that the weights of the penalties, and particularly the ratios of the respective weights be subjected to a thorough study by additional sensitivity analysis as well as by performing additional studies which may help better quantify the weights. This will ensure that the decisions made by the model more accurately reflect the objectives of the airline and the MRO.
- It is recommended that a feedback loop between the MRO's MIS and the Data Exchange Platform be implemented such that there can be no more discrepancies in information between these two systems.

Bibliography

[1] Ozguven, C., Ozbakir, L. & Yasemin, Y. Mathematical Models for Job-Shop Scheduling Problems with Routing and Process Plan Flexibility. *Applied Mathematical Modelling*, 34(6):1539–1548, 2010.

[2] Brucker, P. & Knust, S. *Complex Scheduling*. Springer-Verlag Berlin Heidelberg, second edition, 2012.

[3] Team Leader Planning and Workload Control C-line, The MRO. T-min Policy PS&FC - CE - Powerpoint Presentation, 5th December 2017. For internal use only.

[4] IATA's Maintenance Cost Task Force. . Airline Maintenance Cost Executive Commentary - An Exclusive Benchmark Analysis (FY2014 data) by IATA's Maintenance Cost Task Force. Online PDF, 2017. https://www.iata.org/whatwedo/workgroups/Documents/MCTF/MCTF-FY2016-Report-Public.pdf, last visited 6th November 2018.

[5] Icon Experience. Dot Matrix Icon. https://www.iconexperience.com/i_collection/icons/?icon=dot_matrix, last visited November 13, 2018.

[6] E. White - The Noun Project. Normal Distribution Icon. https://thenounproject.com/term/normal-distribution/991167/, last visited November 13, 2018.

[7] Pixabay. Bullit List Icon. https://pixabay.com/en/list-bullet-list-icon-html-1294835/, last visited November 13, 2018.

[8] Engineering Department, Division Aircraft Maintenance Programs. *Ground Rules for the set-up of the wide-body aircraft AMP and its translation into task cards*. The MRO, 2015.

[9] Stalnaker, T., Khalid, U., Taylor, A., Alport, G. Airline Economic Analysis - Oliver Wyman. Online PDF, 2018. see page 49, https://www.oliverwyman.com/content/dam/oliver-wyman/v2/publications/2018/January/Airline_Economic_Analysis_AEA_2017-18_web_FF.pdf, last visited 6th November 2018.

[10] Saltoglu, R., Humaira, N., Inalhan, G. Aircraft Scheduled Airframe Maintenance, and Downtime Integrated Cost Model. *Advances in Operations Research*, Volume 2016(2):1–12, 2016.

[11] Saranga, H. Opportunistic Maintenance using Genetic Algorithms. *Journal of Quality in Maintenance Engineering*, 103(1):pp.66–74, 2004. https://doi.org/10.1108/13552510410526884.

[12] Pintelon, L. & Parodi-Herz, A. *Complex System Maintenance Handbook - Chapter 2 - Maintenance: An Evolutionary Perspective*. Springer-Verlag London, first edition, 2008. online version, http://www.springer.com/cda/content/document/cda_downloaddocument/9781848000100-c1.pdf?SGWID=0-0-45-528899-p173750241.

[13] Ahmadi, A. Aircraft Scheduled Maintenance Programme Development - Decision Support Methodologies and Tools. online version, 2010. https://www.diva-portal.org/smash/get/diva2:990771/FULLTEXT01.pdf, Doctoral Thesis at the Division of Operation and Maintenance Engineering, Lulea University of Technology, Lulea, Sweden.

[14] Dinis, D. & Barbosa-Póvoa, A.P. On the Optimization of Aircraft Maintenance Management. In *Operations Research and Big Data: IO2015-XVII Congress of Portuguese Association of Operational Research (APDIO)*, pages 49–57. Springer International Publishing, 2015.

[15] Yan, S., Yang, T. & Chen, H. Airline Short-Term Maintenance Manpower Supply Planning. *Transportation Research Part A: Policy and Practice*, 38(9-10):615–642, 2004.

[16] Rodrigo, W.T.S., Wijethunga, A.P.L.M., Botheju, W.T.L., Wijepala, A.C.S. & Wickramasinghe, K. Optimized Manpower Capacity Planning Model for Aviation Industry. online version, 2016. http://www.kdu.ac.lk/proceedings/irc2016/2016/ENG-021.pdf3.

[17] Gu, J., Zhang, G. & Li, K.W. Efficient Aircraft Spare Parts Inventory Management Under Demand Uncertainty. *Journal of Air Transport Management*, 42:101–109, 2015.

[18] Van den Bergh, J., De Bruecker, P., Belien, J. & Peeters, J. Aircraft Maintenance Operations: State of the Art. online version, 2013. `lirias.kuleuven.be/bitstream/123456789/426843/1/13HRP09.pdf.pdf`, Research Paper, Faculteit Economie en Bedrijfswetenschappen, HU Brussel.

[19] Hillier, F.S. & Lieberman, G.J. *Introduction to Operations Research.* McGraw-Hill, 10th edition, 2015.

[20] Winston, W.L., & Goldberg, J.B. *Operations Research: Applications and Algorithms.* Thomson-Brooks/Cole, 4th edition, 2004.

[21] Fattahi, P., Mehrabad, M.S. & Jolai, F. Mathematical Modeling and Heuristic Approaches to Flexible Job Shop Scheduling Problems. *Journal of Intelligent Manufacturing*, 18(3):331–342, 2007.

[22] Shen, X., Han, Y. & Fu., J. Robustness Measures and Robust Scheduling for Multi-Objective Stochastic Flexible Job Shop Scheduling Problems. *Soft Computing*, 21(21):6531–6554, 2017.

[23] Thornblad, K., Almgren, T., Patriksson, M. & Stromberg, A. Mathematical Optimization of a Flexible Job Shop Problem including Preventive Maintenance and Availability of Fixtures. Proceedings of the 4th World P&OM Conference / 19th International Annual EurOMA Conference, Amsterdam, Netherlands, July 2012 `https://research.chalmers.se/publication/159057`, 2012.

[24] Yang, H., Sun, Q., Saygin, C. & Sun, S. Job Shop Scheduling based on Earliness and Tardiness Penalties with Due Dates and Deadlines: An Enhanced Genetic Algorithm. *The International Journal of Advanced Manufacturing Technology*, 61(5-8):657–666, 2011.

[25] Ahire, S., Greenwood, G., Gupta, A. & Terwilliger, M. Workforce-Constrained Preventive Maintenance Scheduling Using Evolution Strategies. *Decision Sciences*, 31(4):833–859, 2000.

[26] Hartmann, S. & Briskorn, D. A Survey of Variants and Extensions of the Resource-Constrained Project Scheduling Problem. *European Journal of Operational Research*, 207(1):1–14, 2010.

[27] Kolisch, R. Efficient Priority Rules for the Resource-Constrained Project Scheduling Problem. *Journal of Operations Management*, 14(3):179–192, 1996.

[28] Li, S., Jia, Y. & Wang, J. A Discrete-Event Simulation Approach with Multiple-Comparison Procedure for Stochastic Resource-Constrained Project Scheduling. *The International Journal of Advanced Manufacturing Technology*, 63(1-4):65–76, 2012.

[29] Pritsker, A.B., Waiters, L.J. & Wolfe, P.M. Multiproject Scheduling with Limited Resources: A Zero-One Programming Approach. *Management Science*, 16(1):93–108, 1969.

[30] Artigues, C., Demassey, S. & Néron, E. *The Resource-Constrained Project Scheduling Problem, in Resource-Constrained Project Scheduling: Models, Algorithms, Extensions and Applications.* Wiley-ISTE, first edition, 2008.

[31] Bartels, J.H. & Zimmermann, J. Scheduling Tests in Automotive R&D projects. *European Journal of Operational Research*, 193(3):805–819, 2009.

[32] Heimerl, C. & Kolisch R. Scheduling and Staffing Multiple Projects with a Multi-Skilled Workforce. *OR Spectrum*, 32(32):343–368, 2010.

[33] Cavalcante, V.F., Cardonha, C.H. & Herrmann, R.G. A Resource Constrained Project Scheduling Problem with Bounded Multitasking. *IFAC Proceedings Volumes*, 46(24):433–437, 2013.

[34] Afshar-Nadjafi, B., Rahimi, A. & Karimi, H. A Genetic Algorithm for Mode Identity and the Resource Constrained Project Scheduling Problem. *Scientia Iranica*, 20(3):824–831, 2013.

[35] Wang, X., Dugardin, F. & Yalaoui, F. An Exact Method to Solve a Bi-objective Resource Constraint Project Scheduling Problem. *IFAC-PapersOnLine*, 49(12):1038–1043, 2016.

[36] Klein, R. Project Scheduling with Time-Varying Resource Constraints. *International Journal of Production Research*, 38(16):3937–3952, 2000.

[37] Bertsimas, D., Gupta, S. & Lullic, G. Dynamic Resource Allocation: A Flexible and Tractable Modeling Framework. *European Journal of Operational Research*, 236(1):14–26, 2014.

[38] Safaei, N. Banjevic, D. & Jardine, A.K.S. Bi-Objective Workforce-Constrained Maintenance Scheduling: A Case Study. *Journal of the Operational Research Society*, 62(6):1005–1018, 2010.

[39] Ranjbar, M., Hosseinabadi, S. & Abasian, F. Minimizing Total Weighted Late Work in the Resource-Constrained Project Scheduling Problem. *Applied Mathematical Modelling*, 37(23):9776–9785, 2013.

[40] SOFEMA Aviation Services. The Role of Aviation Engineering Technical Services - Service Bulletins (SB), 2016. `https://sassofia.com/blog/the-role-of-aviation-engineering-technical-services-service-bulletins-sb/`.

[41] A.K. Muchiri. Maintenance Planning Optimisation for the Boeing 737 Next Generation. *Delft University of Technology, Master Thesis*, 2002.

[42] SOFEMA Aviation Services. Considering EASA Compliant Aircraft Certification Principles ISC/MRB/MPD, 2018. `https://sassofia.com/blog/considering-easa-compliant-aircraft-certification-principles-isc-mrb-mpd/`.

[43] Oliver Robelin - Maintenance Review Board Expert - EASA - Certification Directorate/Flight Standards. Maintenance Review Board Process (MRB) and Instructions for Continued Airworthiness - Scheduled Maintenance Instructions, 2010. `https://www.easa.europa.eu/sites/default/files/dfu/ws_prod-g-doc-Events-2010-jan-19-Ref-9.-MRB-process.pdf`.

[44] Kinnison, H.A. *Aviation Maintenance Management.* McGraw-Hill, First edition, 2004.

[45] U.S. Departent of Transportation - Federal Aviation Administration. Advisory Circular - Certification Maintenance Requirements - AC No. 25-19A. online version, 2011. `https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_25-19A.pdf`.

[46] Australian Government Civil Aviation Safety Authority. Airworthiness Bulletin - Certification Maintenance Requirement and Airworthiness Limitation Terminology Explained - AWB 00-022 Issue 1. online version, 2013. `http://en.caa.gov.il/index.php?option=com_docman&task=doc_download&gid=1832&Itemid=`.

[47] U.S. Departent of Transportation - Federal Aviation Administration. Advisory Circular - Extended Operations (ETOPS) and Operations in the North Polar Area - AC No. 135-42. online version, 2008. `https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_135-42.pdf`.

[48] European Aviation Safety Agency (EASA). FAQ n.19494 - Do I have to comply with a 'mandatory' Service Bulletin?, 2017. `https://www.easa.europa.eu/faq/19494`.

[49] The MRO. MRI and JIC information retrieved from the MRO's Maintenance Information System (MIS), 2019. For internal use only.

[50] Kumar, U.D., Crocker, J., Knezevic, J., El-Haram, M. *Reliability, Maintenance and Logistic Support: A Life Cycle Approach.* Springer, 2000.

[51] Papakostas, N., Papachatzakis, P., Xanthakis, V., Mourtzis, D., Chryssolouris, G. An Approach to Operational Aircraft Maintenance Planning. *Decision Support Systems*, 48:pp. 604–612, 2010. `https://doi.org/10.1016/j.dss.2009.11.010`.

[52] European Central Bank. US Dollar - Euro Exchange Rate, 2018. `https://www.ecb.europa.eu/stats/policy_and_exchange_rates/euro_reference_exchange_rates/html/eurofxref-graph-usd.en.html`, visited on October 3rd 2018.

[53] Coolen, D. Maintenance Planning and Scheduling for Aircraft Line Maintenance. 2018. Master Thesis, Aerospace Engineering Department, Delft University of Technology.

[54] IBM Knowledge Center. Solution Status Codes by Number in the CPLEX Callable Library (C API). https://www.ibm.com/support/knowledgecenter/SSSA5P_12.6.3/ilog.odms.cplex.help/refcallablelibrary/macros/Solution_status_codes.html, last visited January 5, 2019.

A

Appendix

Data Preparation Process

As explained briefly in Chapter 6 a data preparation process preceding the optimization model ensures the computation of all relevant input parameters. During this process, the tasks are appointed their applicable aircraft registrations, due dates, check dates, required and available resource capacity, and so on. The purpose of this preparatory process is to ensure that all data required by the optimization model is available in its desired form. The following is an explanation of this process.

A schematic overview of the data preparation process is shown by the flowchart in Figure A.1. Please note that the flowchart only provides a simplified schematic of the actual process. Each of the steps in the data preparation process will be briefly explained:

- Starting with the routine MRIs and JICs inputs, a distinction is made between which MRIs are block tasks and which are OOP tasks.
- The block MRIs are clustered to their respective block
- The aircraft registrations (tails) are allocated to both the OOP and block MRIs. Moreover, the JICs corresponding to each MRI are identified and allocated to the MRIs.
- The due dates of the block MRIs are determined using historical maintenance check records from the MRO's MIS. These records indicate which maintenance checks have yet been executed for each aircraft registration and subsequently indicate which letter check is next to be executed. Having identified the first next check to be executed, the due dates of each MRI of the block corresponding to the letter check is determined using the aircraft utilization rate and the previous execution dates obtained from the historical records. The due dates of the block MRIs are subsequently used to determine the date on which the letter check is to be scheduled. This process is repeated for all consecutive letter checks, until the scheduling date of the letter check falls outside of the specified (four) year planning horizon.
- Contrary to the block MRIs which are computed using the historical records, the due dates of the OOP MRIs are adopted directly from the MIS.
- For each block in the planning horizon the hours of all block MRIs are computed using the JICs of the MRIs. A unique list of JICs is generated from the combined total of all JICs in each block. For this unique list of JICs the combined hours are equivalent to the total hours of each block. Subsequently, the access panel hours of the various JICs are also added to the JIC hours.
- Similarly to the blocks, the hours of each OOP MRI are computed by summing the hours of the JICs that make up the MRI.
- The non routine hours for each block are estimated using the non routine prediction algorithm based on the unique list of JICs from the block. The same process is applied for each OOP MRI.

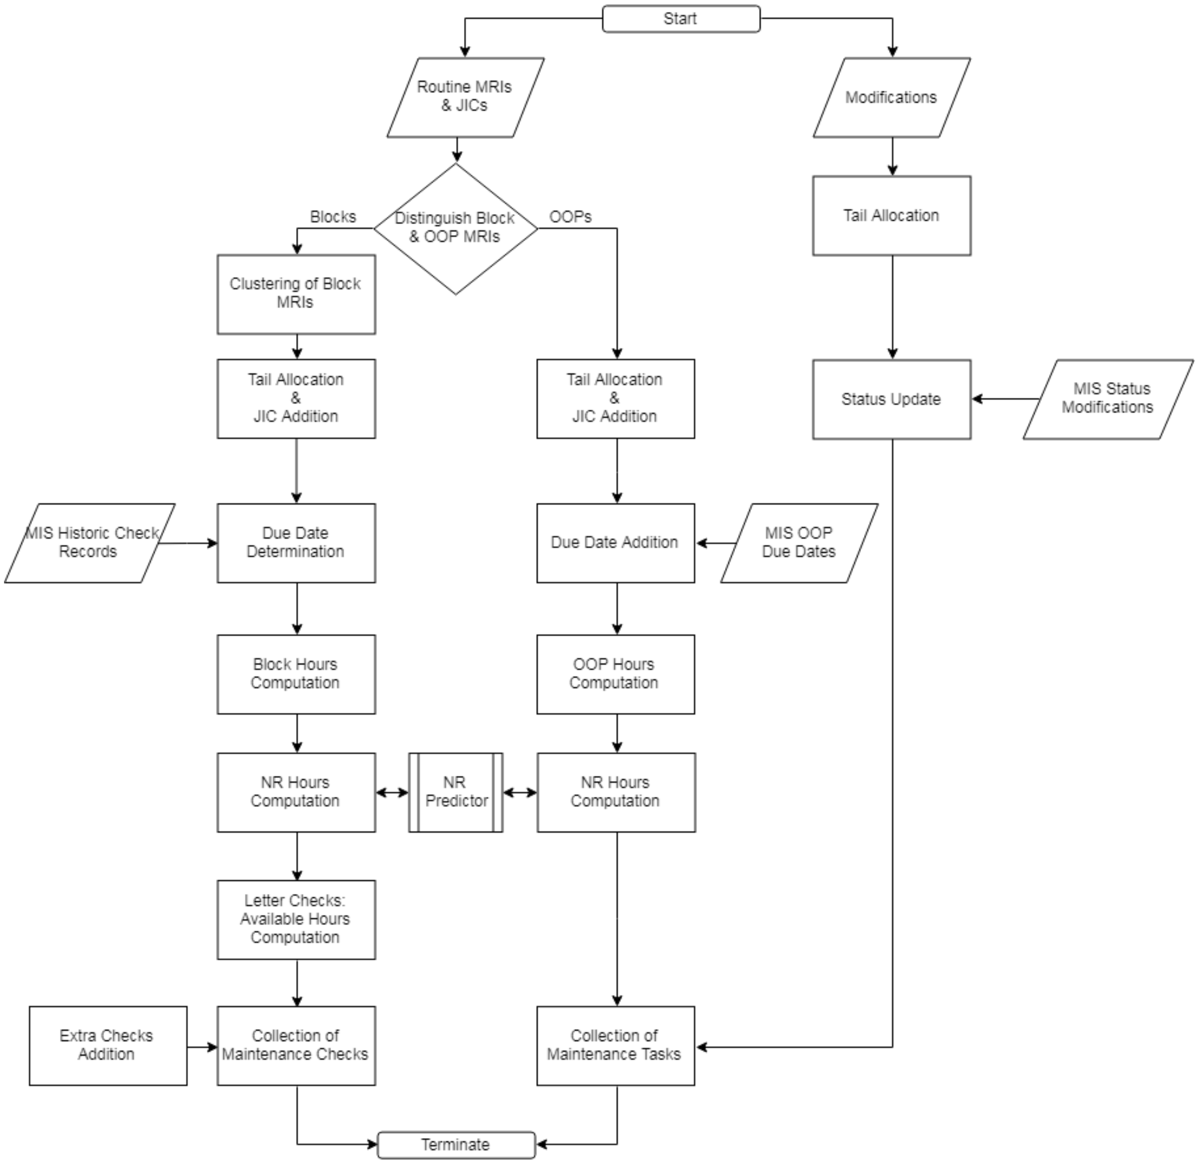

Figure A.1: Schematic overview of the data preparation process to prepare the optimization model inputs

- Having determined the total JIC hours, access panel hours and non routine hours of each block within the planning horizon, the remaining available capacity of each letter check is computed by subtracting the combined total of JIC hours, access panel hours and non routine hours from the predefined capacity norms of each check type. The result is an overview of all letter checks for all aircraft registrations over the planning horizon, with their respective check dates and available capacity.

- The addition of the extra maintenance checks to the overview of letter checks completes the preparation of the maintenance checks for the optimization.

- The modification tasks require fewer preparatory steps. As for the other maintenance tasks, the first step is to designate the various aircraft registrations to each applicable modification, based on the modification input. Secondly, the current status of the modifications is updated using a status report from the MIS. This latter step adjusts the applicability of the modifications according to their current execution status. A number ranging from zero to one is used to indicate the status of the modifications. Each modification with a fully active status receives the number 1, whereas modifications that are completed are assigned a zero. Since modification may consist of various subtasks, if a modification has been only partially completed, its status number is set equal to the fraction representing the portion of the modification yet to be completed. Thus, a modification consisting of two parts of which only one has been completed receives 0.5 as its status number.

- The collection of all modification tasks and OOP tasks with their respective required labor hours, due dates, aircraft registrations and intervals into one combined overview completes the preparation of the maintenance tasks for the optimization.

The running time of the data preparation process is approximately 4 hours. A large portion of this is spent computing the non routine hours for the OOP tasks with the non routine predictor. The outputs of the data preparation process are the combined overviews of all maintenance checks and tasks with their respective due dates, check dates, available and required labor hours, intervals, check types and names, task types and names, and aircraft registrations (except for the extra checks). The outputs of the data preparation process are subsequently used as the primary inputs for the task scheduling optimization model.