Data-driven Decision Support for Component Flow Turnaround Time Reduction in Aircraft Maintenance

Case Study at KLM Engineering & Maintenance

T.F. Spaan

Data-driven Decision Support for Component Flow Turnaround Time Reduction in Aircraft Maintenance

Case Study at KLM Engineering & Maintenance

by

T.F. Spaan

to obtain the degree of Master of Science
in Aerospace Engineering
at Delft University of Technology
to be defended publicly on Friday August 31, 2018 at 10:00 AM.

Student number:	4205472
Project duration:	December 15, 2017 – August 31, 2018
Supervisors:	Wim Verhagen (TU Delft)
	Scout Herremans (KLM)
	Jan Willem van Woerdekom (KLM)

TU Delft

Executive Summary

Aircraft maintenance is inherently unpredictable due to differences in failure rates between components. This often leads to high variation in incoming component flow, which in case of insufficient and inflexible capacity results in backlogs of work, longer turnaround time, and the need for high stock levels. There are two main ways in which a production environment can deal with high variance in demand to obtain the desired service level. The first method accepts the highly variable demand and aims to control the available resources such as capacity, while the second aims to control the number of incoming components under fixed conditions. The aim of this research project is to develop two approaches: *"Determine the effect of service level optimisation in the presence of highly variable demand from an aircraft maintenance supply chain perspective, by means of a capacity optimisation technique and a multi-criteria decision-making method".*

Academically, the focus is on the use of technique and the combination of several research areas and tactical and operational methods. In order to apply both approaches and test the effect in a real-life situation, a case study at KLM Engineering & Maintenance is carried out. The scope is limited to a repair shop EWF, which is a key business entity in the supply chain that repairs high business value components and is having problems in terms of performance and turnaround time. For KLM, the aim is to determine the possible savings (quantitatively and qualitatively) that can be obtained by implementing either of the two models. The novelty of this research project is found in the combination of the research areas: *service level optimisation* in the presence of *highly variable demand* from an *aircraft maintenance supply chain.* Specifically by developing approaches that eliminate assumptions made on demand characteristics, and translating service level optimisation to TAT reduction (and eventually stock reduction) rather than focusing on inventory optimisation.

Phase I aims at determining the required and optimal capacity in terms of net manpower per day in the presence of highly variable inflow, thus creating sufficient flexibility in the process to handle the incoming flow. To do this, a greedy algorithm is used. The algorithm simulates the path that each component follows in steps of 1 day for the chosen time interval. The model is initiated at 0 and the capacity is increased until the service level is sufficient. The greedy algorithm is preferred over other exact techniques such as linear programming, due to its simplicity, transparency, high computational speed, and flexibility regarding generalisation. The main risk of a greedy algorithm (obtaining a local optimum rather than a global) is avoided by initiating the model from 0 and increasing capacity by 1 until the desired performance is obtained. The lowest capacity at which this occurs is automatically the best.

Results from the case study show that there is a direct relationship between available capacity in terms of manpower and service level, which is dependent on the distribution of regular/disrupted flow, productivity, demand scenario, chosen priority procedure, and personnel scheduling. The implementation of disrupted flow assumes a certain percentage of incoming components will be overdue, regardless of the available capacity. For the case study based on historical data from 2017 this results in a maximum performance of 83%, far below the desired 95%. The higher the productivity of technicians, the lower the waste in the repair process, implying lower required capacity. Regarding the demand scenario, the most important parameters are the variance in incoming inflow and the average number of incoming components per day. Increasing (or decreasing) demand by steps of 10%, results in an almost linear relationship between required capacity and incoming demand for the case study at KLM. Perhaps surprisingly, the effect of high peaks on the required capacity is limited, suggesting the process itself contains a certain degree of flexibility. The priority procedure in many cases is a strategical choice, but has shown to impact the performance significantly. While FIFO is the most logical procedure, another possibility is 'first-from-buffer', which provides the illusion of a significantly higher performance, but yields a major increase in average days overdue (in case of insufficient capacity). Finally, working multiple shifts, weekends, and/or nights, results in a different relationship between available capacity and service level, as it effectively yields a lower RPT or longer contracted TAT. In the case study the scenario of working weekends is tested, resulting in an average increase in performance of 10%, regardless of the current capacity.

For the case study at hand, the required capacity is computed based on historical inflow data from 2016-2017,

taking account only regular flow, and the new industry standard of 14 calendar days as contracted TAT. This yields a required capacity of 19 net fte per day: an increase of 6 fte from the current 13. The effect of highly variable flow can be seen in the difference in required capacity between the actual situation (19 fte), and a simulated situation in which the variance is 0 (15 fte). Implementing a capacity of 19 in the shop yields on average 4 days TAT reduction for the IDG, and 2 days for the VFSG. By means of a component-specific user function, this is translated to a reduction in required units in stock of 2 for the IDG, and 1 for the VFSG; yielding a total saving of approximately USD 1.2 million.

Phase II focuses on controlling the demand by assisting in in- or outsourcing decision-making, for which the weighted sum method (WSM) is used. The WSM is simple, transparent, fast, and allows for generalisation of the model for other applications. The main weakness of the WSM is the inability to incorporate different units, which is eliminated by normalisation of scales. The three considered criteria are (effect on) service level, expected days overdue, and direct cost. For each incoming component, the model advises one of two alternatives: repair the component in-house or outsource it to an outside vendor. To limit subjectivity in determining the level of importance of each of the criteria, 1000 combinations of weights are used to visualise the effect on the best alternative.

Several parameters were tested to analyse the possible impact on the best alternative: chosen time interval, available capacity, peak size, and component-repair combination. Using data from the KLM case study, it can be concluded that the available capacity has the largest impact on the decision to in- or outsource. The reason for this is that available capacity over a certain time interval affects the entire shop state: buffer-size, work in stock, shop performance, and shop TAT. For a capacity of 19, the decision support system advises to repair all incoming components in-house, up to a maximum peak of 24 incoming components in one day. This supports the conclusion from Phase I that using a capacity of 19 fte creates sufficient flexibility, and thus robustness, in the process to maintain the desired service level of >95%. For a capacity of 10, or even 15, the advice is to outsource the first (and following) incoming components, in order to restore a stable and sufficient shop service level. The component-repair combination might impact the decision to in- or outsource a component, only in case of insufficient capacity and high repair cost. Maximum peak size to be handled in-house is also highly dependent on the available capacity; for a capacity of 15 or lower, the maximum peak size is 1. On the other hand, for a capacity of 18 the maximum peak size is 14, and for a capacity of 19 the maximum peak size is 24.

When testing the MCDM model on a historical inflow scenario, and comparing the results to the actual situation, the first conclusion is that the model advises to outsource more than five times as many components as the amount actually outsourced: 10 vs. 52. This results in a USD 1,200,000 increase in repair cost, but also an increase in service level of 30% on average: from a 60% average with high variation, to a stable 90%. Making the link to the supply chain by means of TAT and required stock, this translates to an *increase* in TAT of 3 days. This is explained by the significantly longer TAT in case of outsourcing, which is five time higher in the modelled scenario. However, if KLM follows up on contract agreements with outside vendors (maximum TAT of 28 (current) or 15 (new)), the shop TAT is reduced by 1 and 4 days, respectively. This yields a saving of 1 VFSG in required stock for the analysis between January and March 2017.

Besides the quantitative benefits discussed above, there is one major qualitative benefit for both approaches: the shop service level is increased significantly, and perhaps more importantly, stable. This yields higher reliability of the shop,which extrapolates to the entire supply chain. For the MCDM, an additional benefit is that by quickly making in- or outsourcing decisions (preferably when the component is still at the customer), valuable transportation time can be saved by outsourcing directly from the customer.

Main limitations of the models include the inability for flexible capacity and priority procedures for the greedy algorithm, and the use of only three criteria for the weighted sum method. Regarding the case study, major limitations are the focus on regular flow and the use of historical distributions which are likely to change in the future. Besides that, only the top two components are taken into account for supply chain analysis, and the focus is on reducing the required units in stock. In reality there are other parameters with significant impact on the supply chain, such as Customer Service Level and cost related to additional lease-in of components. Finally, the case study for Phase II only takes into account a limited time interval of 3 months.

The model is validated by means of expert validation as well as a sensitivity analysis in which three scenarios are developed with varying weights. All measured parameters remain within 10% upper- and lower-limit of the modelled scenario, from which it can be concluded that the model is robust.

While a large degree of specialisation is introduced by introduction and analysis of the case study, the methodology for both phases allow for generalisation. The main reason for this is that both approaches discussed in Phase I and II are straightforward, transparent, and relatively simple. Besides that, they consist of 'building blocks', which can be adapted and/or expanded depending on the application. Specialisation is introduced mainly by use of application-specific inputs and assumptions, which therefore allows for significant flexibility in application. Overall, both approaches are successful in increasing SL and reducing TAT in the presence of highly variable flow, focusing on applications in an operational business entity in an aircraft maintenance supply chain that have problems in terms of service level, flexibility, and/or available capacity, with a process that allows for outsourcing demand, and for which a time step of 1 day is sufficiently accurate.

While this research project introduces two approaches to service level optimisation in the presence of highly variable demand from an aircraft maintenance supply chain perspective, there are opportunities for further research. The main focus of additional research should be on increasing complexity to obtain a higher level of accuracy and possibilities for generalisation of application. Besides that, a combination of both approaches is to be researched to analyse a possible optimal combination of the two approaches, yielding a strategical and operational model.

Preface

This report contains the result of the nine-month MSc Thesis Project performed as part of a joint MSc thesis between Delft University of Technology and KLM Royal Dutch Airlines. This report is the final deliverable for the course AE5310 - MSc Thesis as part of the MSc Aerospace Engineering specialising in Air Transport and Operations.

The aim of this research project is to contribute to the research gap identified between service level optimisation in the presence of highly variable demand from an aircraft maintenance supply chain perspective. This is done by developing two approaches. The first method accepts the highly variable demand and aims to control the available resources such as capacity, while the second aims to control the number of incoming components under fixed conditions. The link to the supply chain is made by translating the unit turnaround time reduction to supply chain turnaround time reduction and effect on the required stock levels. Pre-existing knowledge on- or experience in aircraft maintenance processes is advised when reading this report.

I would like to express my gratitude and appreciation to my supervisors Wim Verhagen (TU Delft), Scout Herremans (KLM), and Jan Willem van Woerdekom (KLM). While Wim provided me with academic insights regarding the scientific content, Scout and Jan Willem presented me with an incredible amount of information regarding the practical applications and processes. All guided me in the right direction and helped shape not only my research, but also provided me with invaluable knowledge for the rest of my career - and personal life. I would also like to thank the entire CS2.0 team for their input, support and all the laughter during my time at KLM - it would not have been the same without you. Finally, I would like to thank my family and friends for their unconditional love and support throughout these past six years as a student.

Tessa
Amsterdam, August 2018

Contents

Executive Summary iii

List of Figures xi

List of Tables xiii

Abbreviations xv

1 Introduction 1

2 Theoretical Background and State-of-the-art 5
- 2.1 Aircraft Maintenance Supply Chain and the Impact of Demand. . . . 5
- 2.2 Phase I: Optimisation of Shop Capacity 7
  - 2.2.1 Exact Solution Methods 7
  - 2.2.2 Heuristics 8
- 2.3 Phase II: Decision Support System 9
- 2.4 Novelty & Research Contribution 12

3 Methodology Phase I: Optimisation of Shop Capacity 15
- 3.1 Choosing a Solution Technique 15
- 3.2 Approach 15
- 3.3 Assumptions 19
- 3.4 Data Gathering & Reliability. 20
- 3.5 Uncertainty Scenarios to Analyse the Relationship between Capacity and Shop-Performance . . 20
- 3.6 Required Shop Capacity and the Effect on the Supply Chain 21
- 3.7 Strengths, Weaknesses & Limitations of Approach 22

4 Implementation Phase I: Optimisation of Shop Capacity 23
- 4.1 Initialisation 23
- 4.2 Greedy Algorithm 23
- 4.3 Verification of Greedy Algorithm 26

5 Case Study 27
- 5.1 General Shop Information 27
- 5.2 Analysis of Inflow 28
- 5.3 Scope of Case Study. 30

6 Results & Discussion Phase I: Optimisation of Shop Capacity 31
- 6.1 Results of Uncertainty Scenarios 31
- 6.2 Link to Supply Chain 36
- 6.3 Conclusions Phase I. 38
  - 6.3.1 Conclusions: Uncertainty Scenarios 38
  - 6.3.2 Conclusions: Required Capacity and Effect on Supply Chain. 38

7 Methodology Phase II: Decision Support System 41
- 7.1 Choosing a Multi-Criteria Decision-Making Method 41
- 7.2 Approach 42
  - 7.2.1 Determination of Direct Cost 42
  - 7.2.2 Determination of the Effect on the In-Shop Service Level 42
  - 7.2.3 Determination of Expected Days Overdue 43

7.3 Set-up of Scales & Grading of Alternatives. 43
7.4 Assumptions 44
7.5 Data Gathering & Reliability. 45
7.6 Solution Technique for Analysing the Effect of Parameters on Best Alternative 45
7.7 Approach for Analysing Result of Using MCDM Method on Historical Inflow Scenario 46
7.8 Strengths, Weaknesses & Limitations of Approach 47

8 Implementation Phase II: Decision Support System 49
8.1 Initialisation 49
8.2 Operational Model 49
8.3 Multi-Criteria Decision-Making: Weighted Sum Method 50
8.4 Verification of Multi-Criteria Decision-Making Method 51

9 Results & Discussion Phase II: Decision Support System 53
9.1 Effect of Varying Individual Parameters on Decision-Making 53
9.1.1 Effect of Time Interval/Inflow Scenario 53
9.1.2 Effect of Capacity 54
9.1.3 Effect of Component-Repair Combination. 55
9.1.4 Effect of Inflow Peak Size. 55
9.2 Effect of MCDM on Shop and Supply Chain Based on Historical Inflow Scenario 56
9.3 Conclusions Phase II 59
9.3.1 Conclusions: Effect of Parameters on Decision to In- or Outsource 59
9.3.2 Conclusions: Case Study and Effect on Supply Chain 59

10 Validation 61
10.1 Expert Validation 61
10.2 Sensitivity Analysis 61
10.3 Limitations of Validation Strategy. 64

11 Comparison Application and Results Phase I and II, & Possibilities for Generalisation 65
11.1 Similarities and Differences in Application Phase I & Phase II. 65
11.2 Similarities and Differences in Case Study Results Phase I & Phase II 66
11.3 Possibilities for Generalisation 67

12 Conclusion & Recommendations 69
12.1 Conclusions. 69
12.2 Recommendations 71

Bibliography 73

A Appendix: Visualisation Greedy Algorithm 77

B Appendix: Demand Forecasting 79
B.1 Theoretical Background Demand Forecasting Techniques 79
B.2 Results Demand Forecasting Shop EWF. 82

C Appendix: Expert Validation Scenarios 83

List of Figures

2.1 Categorisation of Demand Patterns [7] . . . 6
2.2 AHP Structure [28] . . . 11

3.1 Schematic Visualisation of Abstract Locations Used in Greedy Algorithm . . . 16
3.2 Visualisation of Key Simulation Actions Greedy Algorithm: Day 4 FIFO . . . 17
3.3 Visualisation of Key Simulation Actions Greedy Algorithm: Day 6 FIFO . . . 17
3.4 Flowchart for Greedy Algorithm Determining Required Shop Capacity . . . 18

5.1 Visualisation of Shop Process in Shop EWF . . . 27
5.2 Daily Inflow Shop EWF 2017 . . . 28
5.3 Total Number of Incoming Components per Type 2017 . . . 28
5.4 Weekly Inflow vs. Work in Stock & Weekly Inflow vs. SL 2017 Shop EWF . . . 29

6.1 Capacity vs. Performance Shop EWF Based on March 2017 Inflow: Regular vs. Disrupted Flow . 31
6.2 Capacity vs. Performance Shop EWF Based on Inflow from Several Months 2017: Regular Flow . 32
6.3 Capacity vs. Performance Shop EWF Based on Jan-Mar 2017 Inflow: Regular vs. Disrupted Flow 32
6.4 Required Capacity per Month for SL > 95%, 2016 and 2017 Inflow Data . . . 33
6.5 Correlation between Required Capacity and Key Characteristics of Inflow Data 2017 . . . 34
6.6 Required Capacity for Multiple Growth Scenarios (Baseline Jan-Mar 2017) . . . 34
6.7 Effect of Working Weekends for Jan-Mar 2017 Inflow . . . 35
6.8 Effect of Different Priority Procedures on Performance, Inflow Jan-Mar 2017 . . . 35
6.9 Relationship between end-to-end TAT and Required number of IDGs and VFSGs . . . 37

7.1 Flowchart Weighted Sum Method . . . 42

8.1 Verification Best Alternative Varying Weights: 10 data points . . . 52
8.2 Verification Best Alternative Varying Weights: 100 data points . . . 52
8.3 Verification Best Alternative Varying Weights: 1000 data points . . . 52
8.4 Verification Best Alternative Varying Weights: w1 vs. w2 . . . 52
8.5 Verification Best Alternative Varying Weights: w1 vs. w3 . . . 52
8.6 Verification Best Alternative Varying Weights: w2 vs. w3 . . . 52

9.1 Best Alternative for Varying Weights (Capacity = 19) - Time Intervals: 1, 3, 6, and 12 Months . . . 53
9.2 Best Alternative for Varying Weights (Capacity = 15) - Time Intervals: 1, 3, 6, and 12 Months . . . 54
9.3 Best Alternative for Varying Weights: Capacity = 10 . . . 54
9.4 Best Alternative for Varying Weights: Capacity = 15 . . . 54
9.5 Best Alternative for Varying Weights: Capacity = 19 . . . 54
9.6 Best Alternative for Varying Weights - Effect of Component/Repair: Capacity = 19 . . . 55
9.7 Best Alternative for Varying Weights - Effect of Component/Repair: Capacity = 15 . . . 55
9.8 Effect of Incoming Component on In- or Outsourcing Decision for Multiple Capacity Scenarios 56
9.9 Number of Outsourced Components (cumulative) for Actual Situation and Simulation over Time 57
9.10 In-Shop Service Level over Time for Actual Situation and Simulation . . . 57

10.1 Cumulative Number of Outsourced Components and SL for Various Validation Scenarios . . . 63

A.1 Visualisation of Greedy Algorithm . . . 77
A.2 Visualisation of Key Simulation Actions Greedy Algorithm: Day 6 First-from-buffer . . . 78

B.1 Mean Average Deviation for Varying n: MA, WMA . . . 82
B.2 Mean Average Deviation for Varying Alpha: SES, CRO, SBA, SY, and TSB . . . 82

List of Tables

4.1 Verification Scenarios for the Greedy Optimisation Model . . . 26

5.1 Historical Repair Process Time Shop EWF 2017 . . . 29
5.2 Turnaround Time 2017 for Varying Scenarios Shop EWF . . . 29
5.3 Key Numbers and Percentages Regarding In- and Outflow of Components Shop EWF . . . 30

6.1 Average Inflow, Maximum Inflow, CV^2, and ADI Values for Varying Inflow Scenarios . . . 32
6.2 Required Capacity for Varying Time Intervals . . . 36
6.3 Savings in Required Stock per Component Type in TAT and USD . . . 37

7.1 Baseline Scenario 1 and 2 Decision Support System . . . 46

8.1 Verification of Total Score Multi-Criteria Decision-Making Model . . . 52

9.1 Distribution of Outsourced Components in Shop EWF: January-March 2017 . . . 57
9.2 Difference in Repair Cost: Actual vs. Modelled Scenario . . . 57
9.3 Average Shop TAT for Several Scenarios - Actual vs. Modelled . . . 58
9.4 Average Shop TAT in Case of Varying Outsourcing Contracted TAT . . . 58

10.1 Results of Expert Validation . . . 61
10.2 Validation Scenarios with Varying Weights . . . 62
10.3 Repair Cost and Difference from Baseline for Validation Scenarios . . . 63
10.4 Average in-shop TAT Including and Excluding Outsourced Components for Validation Scenarios 63

C.1 Scenario 1 . . . 83
C.2 Scenario 2 . . . 83
C.3 Scenario 3 . . . 83
C.4 Scenario 4 . . . 83
C.5 Scenario 5 . . . 84
C.6 Scenario 6 . . . 84
C.7 Scenario 7 . . . 84
C.8 Scenario 8 . . . 84
C.9 Scenario 9 . . . 84
C.10 Scenario 10 . . . 84

List of Abbreviations

ADI	Average Inter-Demand Arrival Interval
AHP	Analytic Hierarchy Process
ANP	Analytic Network Process
AOG	Aircraft On Ground
AUP	Auxiliary Power Unit
BUG	Back-Up Generator
CIRO	Customer Interface Repair Officer
CSL	Customer Service Level
CV	Coefficient of Variance
DP	Dynamic Programming
DSS	Decision Support System
E&M	Engineering & Maintenance
EUR	Euro
FH	Flight Hours
FIFO	First-In-First-Out
GP	Goal Programming
IDG	Integrated Drive Generator
IP	Integer Programming
KLM	Koninklijke Luchtvaart Maatschappij
KPI	Key Performance Indicator
LP	Linear Programming
MAUT	Multi Attribute Utility Theory
MAVT	Multi Attribute Value Theory
MCDM	Multi Criteria Decision Making
MRO	Maintenance Repair Overhaul
MSc	Master of Science
MSE	Mean Squared Error
MTBR	Mean Time Between Removal
OEM	Original Equipment Manufacturer
QPA	Quantity Per Aircraft
RPT	Repair Process Time
SBA	Syntetos-Boylan Approximation
SES	Single Exponential Smoothing
SL	Service Level
T&M	Time & Material
TAT	Turn Around Time
TU	Technische Universiteit
USD	US Dollar
VFSG	Variable Frequency Starter Generator
WIO	Work in Overflow
WIP	Work in Process
WIS	Work in Stock
WPM	Weighted Product Method
WSM	Weighted Sum Method

1
Introduction

For many airlines, aircraft maintenance is considered a necessary evil due to the lack of income while the aircraft is grounded. For others, aircraft maintenance is part of their core business which they use to increase their profit, expand their network, and build their brand. Either way, aircraft maintenance plays a significant role within the aviation industry, its main purpose being the assurance of serviceability and prolonging the aircraft's operational life. The level of importance is proportionate to the level of complexity of the aircraft maintenance industry, which is characterised as heavily regulated, time-driven, and cost-intensive [47]. Aircraft maintenance consists of a vast amount of -and variation in- repair tasks. Each task requires a unique combination of skilled personnel, equipment, and material, implying the need for not only a large organisation, but also an enormous capital investment [33].

Another key characteristic of aircraft maintenance is that it is inherently unpredictable due to differences in failure rates between components [6]. This, in turn often leads to lumpy or intermittent demand of components, which is notoriously difficult to forecast and plan [8][68]. Due to the high level of uncertainty in forecasting demand combined with the time-sensitive industry, stock levels of aircraft spare parts are required to be relatively high. Maintaining a large inventory of spare parts is very costly, both due to the required storage space, as well as the high cost of aircraft components [37]. Insufficient stock can be caused by 1) unexpected and unpredictable external factors, 2) poor planning and investment tactics, and 3) high turnaround time (TAT) in the supply chain, suggesting the components are not serviceable on-time. While causes 1 and 2 are fairly subjective and difficult to quantify, reducing TAT in any part of the supply chain has a more direct impact on the stock levels and/or customer service level.

Reduction of turnaround time can be achieved in multiple manners throughout the supply chain, for which a key requirement is meeting the service level. Obtaining and maintaining a high and stable service level (or performance) is challenging in the aircraft maintenance industry, due to the highly unpredictable and variable demand. Generally, there are two main ways in which a production environment can deal with high variance in demand to obtain the desired service level. The first method accepts the highly variable demand and aims to control the available resources such as capacity, while the second aims to control the number of incoming components under fixed conditions:

1. Creating sufficient space in the process to handle high inflow peaks
2. Anticipating on- and managing of incoming flow

The first option is to create sufficient space in the process, which can be done in multiple ways such as buffering, work planning, or personnel scheduling. For this research project however, the scope is limited to the optimisation of required capacity in the presence of highly variable inflow. Regarding the management of inflow, the focus is on assisting in in- or outsourcing decision-making in case the current state in the repair shop has reached critical limits and service level is in danger. Another manner of anticipating on inflow is to predict demand, however previous research has shown that to be a very complex task in the presence of high variation [25][46][24].

Combining the above-mentioned focus areas, the aim of this research project can be summarized as follows: *"determine the effect of service level optimisation (both in terms of capacity optimisation and the use of a multi-criteria decision-making method) in the presence of highly variable demand from an aircraft maintenance supply chain perspective"*. In order to discuss both methods, the research as well as the report are divided into two phases: Phase I covers the capacity optimisation, while Phase II discusses the decision-making method.

A key requirement for this research project is the applicability in the aircraft maintenance industry. In order to test the applicability and obtain a better understanding of implementation and results, a case study is carried out in collaboration with KLM Engineering & Maintenance (E&M), Division Component Services. The problem described above is one that KLM E&M, as one of the world's leading Maintenance Repair and Overhaul (MRO) organisations, is very familiar with. KLM E&M maintains a large pool of spare parts for its own fleet as well as its external partners and customers, worth EUR 275 million [19]. In order to minimise stock levels while maximising service level to the customer, KLM's Component Services (CS) is undergoing a complete redesign. The main goal of this redesign is to become market leader in Component Availability for the Boeing 737 and 787, while lowering the turnaround time (TAT) and increasing the service level (SL) to 95%. Within this project, reducing the TAT is crucial, as it will contribute to an increase in availability and service level, while simultaneously leading to lower stock levels and thus reduced cost [37].

Regarding the scope of the case study, the focus is on optimising the performance of one unit in the supply chain: the repair shops. To limit the scope further, shop Hydraulics 1 - workcentre EWF is considered. The main reason for this is that the inflow of shop EWF is highly variable, which in combination with the long Repair Process Time of the components, has caused significant problems regarding the service level and a large backlog of work. Besides that, shop EWF contains several high business value components, among which the Variable Frequency Starter Generator: a key component in the Boeing 787 with expected growth given the expansion of the contracted 787 fleet. Reducing the buffer-, transport- and waiting-time for components within this shop can lead to lower cost for investment in stock.

As this project is a collaboration between two parties, both have their own set of requirements and desires, resulting in different research objectives and -questions. While TU Delft aims at scientific contribution, KLM E&M is looking for a practical solution to their problem.

Scientific Research Objective & Questions:

"Determine the effect of service level optimisation in the presence of highly variable demand from an aircraft maintenance supply chain perspective by means of a capacity optimisation technique and multi-criteria decision-making method".

1. Which solution technique is best to obtain the required shop capacity, taking into account the practical application requirements?
2. What parameters affect the relationship between capacity and service level, and to what extent?
3. What is the effect of highly variable inflow on the required shop capacity?
4. Which multi-criteria decision-making method is best suited to assist in in- or outsourcing decision-making, taking into account practical application requirements?
5. What parameters affect the decision to in- or outsource, and to what extent?
6. What validation techniques are best suited to validate the multi-criteria decision-making method and its results, given the operational nature of the application?

Practical Project Objectives & Questions:

Phase I: *"Contribute to a reduction in shop TAT by determining the optimal capacity in terms of manpower in order to obtain a sufficient and stable shop performance."*

Phase II: *"Contribute to a reduction in transport- and waiting-time of the component flow by assisting in in- or outsourcing decision-making using a data-driven decision support model combining tactical and operational approaches"*

1. In what manner does the available capacity influence the shop performance?
2. What is the required capacity in the shop given the actual historical inflow data?
3. What savings - qualitatively and quantitatively - can be obtained by implementing the optimal shop capacity?
4. Which parameters have the largest impact on the decision to outsource a component or repair a component in-house?
5. What savings - qualitatively and quantitatively - can be obtained by implementing a multi-criteria decision support system?

While there is plenty of literature available on each of the individual research areas, many papers focus on inventory models and inventory optimisation, rather than service level optimisation in the supply chain. Novelty of the research project can be found in the combination of the research areas: *service level optimisation* in the presence of *highly variable demand* from an *aircraft maintenance supply chain* perspective. While highly variable demand is thoroughly researched, in the majority of studies some type of demand distribution is assumed [27][8]. This research project aims to develop two approaches that are applicable regardless of the demand type and thus eliminate assumptions made in existing studies on demand characteristics. Current literature contains a wide variety of complex linear programming approaches to the capacity optimisation problem [53][42][18], of which the models and thus application are highly specific. This project focuses on the development of two approaches that are more generally applicable. Regarding the MCDM, the integration of an operational model in a tactical decision support model increases the novelty of this research project. Instead of focusing on inventory optimisation, this project aims to translate service level optimisation in one part of the supply chain to TAT reduction in the entire supply chain, which eventually is linked to stock reduction by means of a user function. Additional contribution is added by the research and implementation of two service level optimisation methods applied to a case study at a large MRO organisation.

As mentioned in this chapter, this research project is divided into two phases. The reason for this division is the difference in objective and approach. The report begins with Chapter 2, in which the theoretical framework and state-of-the-art of both phases of the research areas are presented. After this, Phase I is initiated with the methodology in Chapter 3, including model description, assumptions and solution techniques. This is followed by the implementation of this approach in Chapter 4, containing initialisation and pseudo-codes for the model. Chapter 5 covers the description and context of the case study used for this project. Using the information from the case study, results for Phase I are presented and discussed in Chapter 6, in which a distinction is made between results from uncertainty scenarios and the link to the supply chain. After discussing the results from Phase I, Phase II is initiated with the methodology in Chapter 7. This is followed by the implementation and results based on the previously discussed case study in Chapters 8 and 9. Chapter 10 covers the validation of the model. Finally, Chapter 11 discusses the comparison of Phase I and Phase II, as well as possibilities for generalisation. The report ends with conclusions of the findings and recommendations for further research in Chapter 12.

2

Theoretical Background and State-of-the-art

This chapter covers the theoretical background of research areas introduced in Chapter 1. The chapter starts with an introduction on the aircraft maintenance supply chain the effect of demand characteristics in Section 2.1. As discussed before, Phase I of the project focuses on the optimisation of capacity, for which several solution techniques are discussed in Section 2.2. This is followed by the theoretical background in multi-criteria decision-making methods in Section 2.3, which will be used in Phase II of the project to assist in in- or outsourcing decision-making. Finally, Section 2.4 discusses the research gap and contribution of this research project, as well as project novelty.

2.1. Aircraft Maintenance Supply Chain and the Impact of Demand

A supply chain can be defined as: *"a system of organisations, people, activities, information, and resources involved in moving a product or service from supplier to customer"* [21]. For an aircraft maintenance supply chain, the products can be specified as aircraft components or subparts. The activities are focused on repair and maintenance of the components, suppliers are the repair stations, and customers are contracted airlines. In order to optimise the supply chain processes, Tzafestas and Kapsiotis [55] identify three options:

1. Manufacturing facility optimisation: minimise cost from manufacturing only
2. Global supply chain optimisation: assumes direct and cooperative relationship between all stages of the supply chain
3. Decentralised optimisation: individual optimisation of each of the supply chain entities

Option 1 is not applicable given the specialisation into manufacturing. Option 2 introduces many inter-dependencies between the different business entities in the supply chain, which has a risk of lower accuracy. For that reason the focus is on option 3: decentralised optimisation with the aim of linking the results to the supply chain. Many papers focus on inventory optimisation to minimise cost and maximise service level, but another interesting business entity in the aircraft maintenance supply chain are the repair shops. Repair shops can be seen as a specific type of production environment, handling highly complex components, often with highly variable demand patterns. In many cases this combination leads to a variable output flow, which affects the variability of the entire supply chain, resulting in either a lower customer service level or a higher required stock as a buffer.

There are two main (quantitative) performance measures categories in supply chain management that are used to measure the efficiency and/or effectiveness of the system [11]: cost and (customer) service level. Cost in the aircraft maintenance supply chain consist of both direct as well as indirect cost. Direct cost include cost for transport, storage, repair, and personnel. Indirect cost are caused by buy-in and lease-in of additional stock due to disruptions in the supply chain, such as longer turnaround time or reduced service level. Customer service level is defined as the percentage of components delivered within contract agreements. However, service level (or performance) can also be used for individual entities within the supply chain, for example in the repair shop. As briefly mentioned above, both performance metrics are inter-related; a lower service level often results in higher cost (either direct or indirect), and vice versa. Other key parameters that affect the SL and cost are the available resources, productivity, and demand patterns.

Demand, specifically demand characteristics are an important parameter with a large impact on the supply chain and its performance. Historically, four categories of demand have been identified [6][7][25], based on modified Williams [61] criteria:

- *Smooth demand*: little to no variation in interval and quantity

- *Intermittent demand*: highly sporadic demand, but little variation in quantity
- *Erratic demand*: large variation in quantity, but little variation in interval
- *Lumpy demand*: large variation in both interval as well as quantity

The characteristics of each above-mentioned category are based on two parameters: the average inter-demand interval (ADI) and the coefficient of variation (CV^2). ADI determines the average number of time periods between two successive demands, whereas CV^2 is the standard deviation of the demand divided by the average demand. In order to classify the demand as smooth, intermittent, erratic, or lumpy, boundaries are to be set. While both Williams [61] as well as Eaves [9] set specific cut-off values, according to Syntetos, Boylan, and Croston [26] these cut-off values have been chosen based on the particular empirical situation analysed. Syntetos, Boylan, and Croston have shown the boundaries to be at an ADI value of 1.32, and a CV^2 value of 0.49 [26]. These values were the result of a numerical analysis, where the Mean Squared Error (MSE) was compared for several forecasting techniques (Single Exponential Smoothing (SES), Croston's method, and the Syntetos-Boylan Approximation). The result can be seen in Figure 2.1.

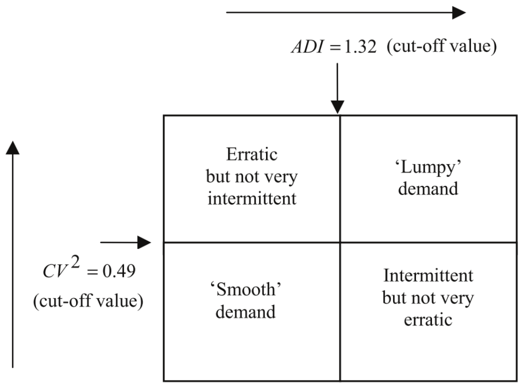

Figure 2.1: Categorisation of Demand Patterns [7]

In case of smooth demand, the level of uncertainty is minimised, yielding a fairly stable inflow of work. However, when considering erratic or lumpy inflow, the variance in incoming work is large. This results in a relatively high level of uncertainty, which can significantly impact the performance of business entities within the supply chain. In many instances, capacity is not very flexible, especially in high-complexity environments such as aircraft maintenance.

As supported by Baghalian et al. [2], demand uncertainty can cause disruptions in the operations in the supply chain and thus result in the requirement for additional inventory and/or the decrease in service level. Gupta and Maranas [5] identify two strategies that businesses can position themselves as in order to deal with demand uncertainty:

1. **Adapter:** the company does not try to control the demand, but rather focuses on increasing flexibility in the operation
2. **Shaper:** the company aims to control the demand in order to limit the risk by means of contract limitations or outsourcing agreements.

Increasing flexibility in the operations can be done in multiple manners such as buffering, work planning, and personnel scheduling. For this project the focus is on personnel scheduling, more specifically the optimisation of shop capacity. Solution techniques for this method are discussed in Section 2.2. Controlling the incoming demand can be done by either contract limitations on number of incoming components in a certain time period, or by outsourcing components to maintain a stable in-shop performance. As mentioned in Chapter 1, the focus is on the latter, of which several solution techniques and theoretical background is discussed in Section 2.3.

2.2. Phase I: Optimisation of Shop Capacity

Theoretically, determining the optimal shop capacity can be classified as a production scheduling problem, in which the optimal capacity is subject to a number of boundary conditions and constraints. In 1954, Edie [32] was one of the first to discuss the personnel scheduling problem, which over time has increased in complexity due to more flexible working conditions for employees. Beliën and Forcé [22], and Cardoen et al. [10] classified problems and solutions methods based on the following characteristics:

1. **Personnel characteristics**
   Regarding personnel characteristics, one key difference can be found in full-time or part-time contracts. Another parameter of importance is the skill required to perform the task at hand, in which can be distinguished between homo- and heterogeneous skill-sets. A final classification can be made in case a crew (or team) is to be scheduled for a task, rather than an individual.

2. **Constraints, performance measures, and flexibility**
   This characteristic starts with the required coverage to cover the workload, which according to Van den Bergh et al. [23] is present in more than 75% of all literature. The main objective is to determine the minimum number of employees to cover the demand. Another well-covered topic is that of flexibility, on which Topaloglu and Ozkarahan [54] state that organisations use multiple shift times, -lengths, and process sequences to create additional flexibility in their schedules.

3. **Solution method and uncertainty**
   The vast majority of reviewed papers use some type of mathematical programming, such as linear, integer, dynamic, or goal programming. Another well-covered solution method are heuristics, especially constructive and meta-heuristics. One important characteristic that influences the choice for a solution method is uncertainty. As stated by Van den Bergh et al. [23], the three main categories of uncertainty are: uncertainty of demand, uncertainty of arrival, and uncertainty of capacity.

4. **Area of application**
   According to Van den Bergh et al. [23], there are six key areas in which the scheduling problem are prominent: services, transportation, general, manufacturing, retail, and military.

Given the project context and research objectives from Chapter 1, the scope is limited to individual task scheduling in a production/manufacturing environment. The key objective is to determine the required manpower to meet the workload, suggesting additional complexities such as flexibility in shifts, contract-types, and process sequences are not the main priority, but could be added to the model in the future to enhance accuracy. One important complexity given the problem at hand is the inclusion of uncertainty of demand and arrival: aircraft maintenance is unpredictable, and so is the required repair time. Several solution techniques are discussed briefly below in order to provide an overview of the different techniques and build a foundation which can be used to choose the best suited technique for the application at hand.

2.2.1. Exact Solution Methods

Linear Programming (LP). Linear Programming problems are characterised by the need to make a decision. This decision is to be made by either minimising or maximising the objective function, which is subject to a set of linear constraints, as shown by Eq. 2.1.

$$\begin{aligned} \text{maximise:} \quad & \mathbf{c}^T\mathbf{x} \\ \text{subject to:} \quad & A\mathbf{x} \leq \mathbf{b} \\ & \mathbf{x} \geq \mathbf{0} \end{aligned} \tag{2.1}$$

Linear Programming is a powerful technique that is applicable to a large variety of problems. It is a fairly simple technique that is adaptable and flexible in use and analysing the problem. Cormen et al. [58] states that a weakness of LP is that the linearity assumption in many cases is unrealistic, often resulting in large discrepancies between model results and the real-life situation. Besides that, LP is a static technique, suggesting that it does not deal with change in variables well.

Integer Programming (IP). Integer Programming is a subset of linear programming, in which the solution is limited to integers. The main difference with LP is in the application; some applications require discrete solutions (e.g. the number of men required to cover the workload), while other decisions can be continuous (e.g. minimisation of cost).

Dynamic Programming (DP). Dynamic programming is a specific type of linear programming, in which the problem is broken down into simpler sub-problems in order to reduce complexity. These sub-problems are solved recursively, after which their solutions are combined in order to solve the original problem, as stated by Cormen et al. [58]. Rather than having one objective function, the decision is broken down into a sequence of decisions over time, characterised by value functions. These value functions are used to determine the optimal values required to obtain the optimal solution. While DP breaks down the problem in smaller, simpler problems, it decreases computation speed significantly.

Goal Programming (GP). Goal Programming is a solution method in which the alternative is chosen with the shortest distance (deviation) from a pre-defined goal or target [56] [4] [12]. The main strength of GP is that it can handle large-scale problems, and translates well to real-world situations. Besides that, GP can be used for both linear as well as non-linear problems, discrete and continuous variables. Weaknesses include the inability to determine weights and the relative high complexity. Also, a goal or target must be defined, which is not always the case in real-world situations.

2.2.2. Heuristics

As stated in Section 2.2.1, exact techniques often result in long computation time, given the fact that they require all possible solutions to be analysed in order to find the optimum. Heuristics attempt to yield a good solution, but not necessarily the optimum [44]. This is done by using accessible strategies derived from previous experience to problems with similar characteristics in order to lower the computational workload an complexity. In real-life problems, decisions often have to be made in a reasonable time interval, something that is not always possible when using LP methods. Besides computation time, Martí et al. [44] provides four other reasons for using heuristics:

1. There is no known method to solve the problem optimally
2. The exact method to solve the problem cannot be used on available hardware
3. The heuristics method provides a higher degree of flexibility compared to the exact method
4. The heuristic method is part of an exact procedure for problem solving

While there are many categories of heuristic methods, there are three main characteristics that a heuristic model should contain, according to Martí et al. [44]. Firstly, the solution should be obtainable with reasonable computational effort and time. Besides that, the solution should be near optimal. Finally, the probability of obtaining a bad solution should be low. Five categories are briefly outlined below.

- *Decomposition Methods:* break down the problem into sub-problems.
- *Inductive Methods:* generalize simpler cases to be used in the overall problem.
- *Reduction Methods:* introduce pre-defined properties as boundaries to the problem, as these properties are proven to be largely fulfilled by good solution options.
- *Constructive Methods:* construct a solution from scratch, usually step-by-step, and based on the best choice at each iteration.
- *Local Search Methods:* initiated with a feasible solution to the problem and tries to improve this with every iteration.

Greedy Algorithm. As defined by Cormen et al. [58], a greedy algorithm makes a locally optimal choice at every step, with the intent of finding the global optimum. Given this definition, the greedy algorithm can be categorised as each of the five above-mentioned classifications, dependent on the application. For some applications, the greedy algorithm could lead to the global optimum, meaning it can be categorised as an exact solution technique.

The main strength of the greedy algorithm is its computational speed. Moreover, the greedy algorithm has a relatively low level of complexity, making it easy to use and understand. Rather than the model being a black-box, the user is able to back-track why certain decisions are made. Besides that, the greedy algorithm can easily be adapted and expanded. A major weakness of the greedy algorithm is that it will not always lead to a global optimum. The level of risk regarding the difference between a local and global optimum is dependent of the type of application.

2.3. Phase II: Decision Support System

The term decision support system (DSS) was first used in literature in 1969 by Ferguson and Jones [48] and has since led to the more updated definition by Marakas in 2003 [20]: "DSS is an interactive computer-based system or sub-system to help decision-makers use communications technologies, data, documents, knowledge and/or models to identify and solve problems, complete decision process tasks, and make decisions." Of the five main categories identified by Power [14] (communications-, data-, document-, knowledge-, and model-driven), model-driven DSS are mainly used in the aviation industry, as stated by Zhang et al. [41]. To be more precise, a specific type of model-driven DSS: multi-criteria decision-making (MCDM) tools are used in the aviation industry with the objective of determining the optimal solution to a certain problem in the presence of multiple criteria. Belton and Stewart [63] identified three main sub-categories of MCDMs:

1. *Value measurement models*: "numerical scores are constructed in order to represent the degree to which one decision option may be preferred to another. Such scores are developed initially for each individual criterion, and are then synthesized in order to effect aggregation into higher-level preference models"

2. *Goal, aspiration, or reference models*: "desirable or satisfactory levels of achievement are established for each criterion. The process then seeks to discover options which are closest to achieving these desirable goals or aspirations";

3. *Outranking models*: "alternative courses of action are compared pairwise, initially in terms of each criterion in order to identify the extent to which a preference for one over the other can be asserted. In aggregating such preference information across all relevant criteria, the model seeks to establish the strength of evidence favouring selection of one alternative over another"

In the context of this research project, the focus is on the development and implementation of a value measurement model, the choice of which will be elaborated on in Chapter 3. The most relevant value measurement techniques are briefly highlighted below.

Weighted Sum Method (WSM). The WSM was introduced in 1963 by Zadeh [31] and has been prominent in literature ever since. It is based on weights that are determined by the decision-maker. Each criterion is given a (non-negative) weight, and each alternative is ranked by computing the weighted sum of the criteria.

$$A_i = \sum_{j=1}^{n} w_j a_{ij} \tag{2.2}$$

Here A_i is WSM score for alternative i, n is the number of alternatives, w_j is the relative weight attributed to criterion C_j, and a_{ij} is the performance value of alternative i evaluated in terms of criterion C_j. The main advantage of the WSM is that it is very straightforward and simple to understand and implement. Mela et al. [29] and Triantaphyllou [16] even argue that WSM should become a standard for evaluation MCDM methods due to its high performance in single-criteria problems. A major disadvantage of the weighted sum method is its limitation in application. One limitation of this method is that it requires all data to be expressed in the same units. Another difficulty is that in case criteria dimensions vary in the order of several magnitudes, normalisation is required. Besides that, Marler and Arora [52] state that is application is limited to optimisation problems with a maximum of two objective functions. Finally, they state that even though WSM is easy to use, it only provides a linear approximation of the preference function.

Weighted Product Method (WPM). The WPM is similar to the WSM, in the sense that weights w_j and performance values $a_i j$ of certain criteria C_j are determined by the decision-maker. In this case however, instead of ranking the alternatives based on the weighted sum, the alternatives are ranked based on the weighted product. There are two manners in which this can be done. In the first, pair-wise comparison is used, as can be seen in Eq. 2.3. Here, the relative performance value is computed for alternative A_K relative to alternative A_L. In case the result is greater than 1, A_K performs better, and in case it is smaller than 1, A_L performs better. The second method, of which the equation can be seen in Eq. 2.4, computes the total performance value of alternative A_K (instead of the relative one) when all criteria are considered simultaneously.

$$P(\frac{A_K}{A_L}) = \prod_{j=1}^{m} (\frac{a_{Kj}}{a_{Lj}})^{w_j} \tag{2.3}$$

$$P(A_K) = \prod_{j=1}^{m} (a_{Kj})^{w_j} \tag{2.4}$$

As discussed before, WSM is limited to single-dimension problems. A major advantage of WPM is its dimensionlessness; due to the relative nature, any units of measure are eliminated. This allows the WPM to be used in both single- as well as multi-dimensional MCDM problems. Besides that, Atmojo [50] states that WPM is more efficient compared to other MCDM methods, which can be largely attributed to its fast computation speed and simplicity. As stated by Mela et al. [29], a downside of the WPM is that it does not work in case the criteria values are zero or negative. This is a minor problem, they argue, as this almost never occurs, but it is something to keep in mind.

Multi-Attribute Value Theory (MAVT). MAVT is a decision support tool for which the foundation was laid during the late 1960s and early 1970s with works by Fishburn [39] and Keeney [49]. MAVT is a tool that assists decision-makers in assigning value to a finite and discrete set of alternatives that are to be evaluated in the case of conflicting objectives. The following steps are to be taken in a MAVT process:

1. Define alternatives
2. Select and define criteria
3. Assess score for each alternative in terms of each criterion
4. Rank the alternatives by applying a value function v

While steps 1-3 are common for most multi-criteria decision models, step 4 is specific for MAVT. As stated by Von Winterfeldt and Edwards [13], the value function v represents the preferences of the decision-maker. Value in this context is defined as a measure of preference under certainty. A characteristic of MAVT is that it is compensatory, meaning that it aggregates the performance across all evaluated criteria.

Let $c_1, ..., c_n$ be a set of attributes associated with the outcome of the problem, and $a_1, ..., a_m$ be the performance value for alternative j. The value function is then equal to $v(c_1(a_j), ..., c_n(a_j))$, and the alternative with the maximum v is the best alternative. If $m = 1$, one can simply choose the alternative with the best score. However, in the case that $m > 1$, the value of each of the attributes is to be taken into account in the overall value. This can be done in two manners: additive and multiplicative. In the additive model, the value function v is split up into multiple functions v_i. The additive model is the simplest form, however it should be noted that this can only be used if there is additive independence between multiple attributes. The multiplicative model is more efficient in its computing speed and allows for interaction between the attributes. A major disadvantage of this model is its complexity.

MAVT has several advantages, starting with its ability to accommodate both quantitative as well as qualitative data. Besides that it aids in structuring and understanding a problem, as its user(s) are required to compose both the value function representing their preference, as well as objectives, criteria, and alternatives. Moreover, the additive form of MAVT is simple to use, and robust, as stated by Keeney and Raiffa [49], and often used in practice. However, they also state that limiting the model to to its most simple (additive) form means it becomes unrealistic. This is due to the lack of uncertainty and the assumption of independence of preferences. Another disadvantage of MAVT is its assumption of complete compensability of criteria, meaning all criteria are expressed in the same unit of measure. For that reason, MAVT is often referred to as a weak form of decision making. As mentioned before, a major limitation of MAVT is that it can only accommodate problems involving a finite and discrete set of alternatives, which is many real-life applications is not the case. Regarding future work, as stated by Herwijnen [34], no applications of non-additive MAVT are found in literature, which is necessary to demonstrate its use and application.

Multi-Attribute Utility Theory (MAUT). MAUT is a more complex form of the multi-attribute value theory discussed previously. Compared to MAVT, Loken [15] states: *"it is a more rigorous methodology for how to incorporate risk preferences and uncertainty into multi criteria decision support methods"*. MAUT is an expected utility theory that aids decision-makers in assigning utility to certain alternatives, considering the decision-maker's preferences. This is done by evaluation of the outcomes taking into account multiple attributes, and finally combining these individual assignments to obtain overall utility measures, as described by Velasquez [35]. Utility in this context is defined as a measure of preference under uncertainty.In order to determine the utility, a utility function must be set up. As Mateo [28] states: a utility function is a representation that quantifies the preferences of the decision-maker by assigning a numerical index to varying levels of satisfaction of a certain criterion. This is done in similar fashion as for MAVT, but instead of using a v to indicate value, u is used to indicate utility.

One of the main advantages of MAUT is that it can take into account uncertainty. According to Velasquez [35], many applications of MAUT rely on this strength of dealing with uncertainty, which is often seen in the fields of economics, finance, and agriculture. Herwijnen on the other hand argues that MAUT is very difficult to apply and no real-world applications are known. Another advantage of MUAT is that it allows the consideration of both qualitative as well as quantitative criteria, and is able to include preference of the decision-maker into the model. This leads to a major disadvantage of MAUT, namely that it requires a lot of input data for it to be successful, in order to accurately record preferences. Moreover, preferences are difficult to apply precisely, and often lead to a relatively large level of subjectiveness.

Analytic Hierarchy Process (AHP). The AHP was developed by Saaty in 1980 [59], and has been widely used in multiple criteria decision-making ever since [35]. The objective of AHP is to identify the preferred alternative, together with determining a ranking of the alternatives when considering all criteria simultaneously. This is done by breaking down an unstructured problem into several components, of which hierarchical levels are arranged, as shown in Figure 2.2.

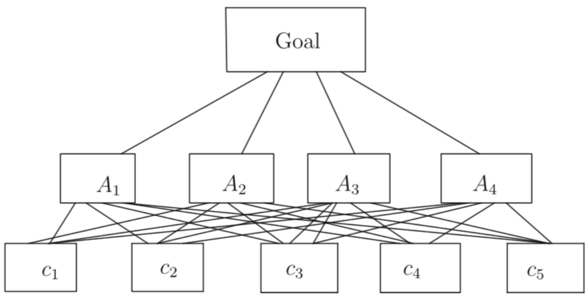

Figure 2.2: AHP Structure [28]

Here, the goal is the objective, A_i, are the alternatives, and c_n are the criteria on which the alternatives are graded. A characteristic of AHP is the use of pair-wise comparisons, both to determine weights as well as comparing alternatives. These pair-wise comparisons are represented by the assignment of numerical values by the decision-maker. The next step is to synthesise these judgments in Eigen vectors in order to determine priority of variables.

AHP is one of the most well-known MCDM techniques, which implies significant advantages over other methods. One of these advantages, as stated by Saaty [60] is its ease of use. Besides that, it is scalable and due to its hierarchical structure, it is easily adjustable to accommodate decision-making problems. While AHP requires sufficient data to be able to properly perform pair-wise comparison, this is significantly less than the amount of data required for MAUT. However, AHP also has some disadvantages, starting with problems of interdependence between alternatives and criteria. Besides that, Konidari and Mavrakis [40] argue that it can be subject to inconsistencies in judgment and performs poorly in identifying weaknesses and strengths due to its pair-wise comparison. Another major criticism has been the fact that AHP is subject to rank reversal, meaning that alternatives being added later in the process often leads to reversal in the final rankings. A limitation of AHP is that it can only be used in linear decision problems [28], and that it is unable to deal with uncertainty.

Regardless of the disadvantages of the method, AHP has seen many applications in areas such as resource management, corporate policy, strategy and risk assessment [35]. According to Subramanian et al. [38], AHP has been used in supply chain management, more specifically for supplier selection, outsourcing and stock management, but its use is limited for other applications within the aviation industry.

Analytic Network Process (ANP). The ANP is a more generalised form of the AHP, similar to the relationship between MAVT and MAUT. Instead of using hierarchies, such as those in AHP, ANP structures objectives, alternatives, and criteria as networks, as stated by Wang in 2012 [57]. The use of networks instead of hierarchies allows for the prioritisation and clustering of elements. This means that ANP is able to handle dependence within a cluster (inner dependence), as well as dependence between different clusters (outer dependence), as stated by Yang et al. [69].

A major advantage of ANP is that it is nonlinear, opposed to AHP which is a linear method. Besides that, ANP outperforms AHP in terms of handling interdependence and can thus be used in more complex decision-making problems, including various intangible criteria, as stated by Tsai et al. [64]. Applications of ANP are mostly found in project selection, optimal scheduling planning, and supply chain management [69].

2.4. Novelty & Research Contribution

Using the information from this chapter discussed in previous section, the main keywords of the research project are: *service level optimisation*, *capacity optimisation*, *multi-criteria decision-making*, *highly variable demand*, and *aircraft maintenance supply chain*. The first (and main) novelty of this research project is found in the combination of the above-mentioned research areas. While the individual areas are thoroughly researched and covered in literature, there is a research gap when it comes to the combination of service level optimisation in the presence of highly variable demand from an aircraft maintenance supply chain perspective.

Highly variable demand in the aircraft maintenance industry has been researched thoroughly [24][27][57]. A major assumption made in these studies is that the demand is exclusively lumpy, intermittent, erratic, or smooth. In reality, it often occurs that demand (and thus demand characteristics) vary over time. Another assumption that is often made in the presence of highly variable demand is the use of certain demand distributions, most commonly the Poisson distribution [30][8]. While in many cases this is the best fit when considering stochastic modelling, it often results in significant discrepancies compared to the actual situation and thus does not handle highly variable inflow well. This research project aims to develop two approaches that are applicable regardless of the demand type and thus eliminate assumptions made in existing studies on demand characteristics.

Service level optimisation can be applied to almost all problems that in some way measure performance. In terms of capacity optimisation, many studies focus on exact methods such as linear programming [23]. While this has shown to incorporate multiple parameters [53][42][18], the models and thus application are highly specific. Significantly less information is found on the use of simpler techniques such as the greedy algorithm and the use of it in heuristic solution methods. This research aims to contribute to the body of research performed in this area and develop a model that is more widely applicable.

Similarly to serivce level optimisation, multi-criteria decision-making models can be applied in practically any field for a wide variety of problems [20][35]. Walker et al. [65] identified a research gap in terms of inclusion of uncertainty in decision support models, which this research project aims to contribute to by limiting subjectivity in weighing criteria. Moreover, the integration of an operational model in a tactical decision support model increases the novelty of this research project as well.

When considering the link between service level optimisation and the (aircraft maintenance) supply chain, the majority of studies focus on inventory optimisation, which combines both research areas directly [66][36]. These models often handle highly variable inflow poorly due to the assumptions made on demand distribution. Besides that, they focus primarily on cost minimisation in one part of the supply chain: the warehouse or storage location. This research project aims to translate service level optimisation in one part of the supply chain to TAT reduction in the entire supply chain, which eventually is linked to stock reduction by means of a user function.

Finally, application in the aircraft maintenance industry adds to the novelty, especially in combination with the case study performed in an operational environment.

3
Methodology Phase I: Optimisation of Shop Capacity

The aim of Phase I is to determine the required and optimal shop capacity in terms of manpower per day based on historical inflow data in the presence of highly variable demand. Moreover, analysis of multiple scenarios can provide insights in the effect of capacity on shop performance, and vice versa. Besides that, a link is to be made to the supply chain, which is used to analyse the results in terms of key parameters in the supply chain. The chapter is outlined as follows. Section 3.1 covers the choice of solution technique for this application, followed by the approach taken in Section 3.2. Section 3.3 covers the assumptions and their possible impact on the solutions, which is followed by the data gathering and reliability of data in Section 3.4. The uncertainty scenarios and approach to yield the desired results are discussed in Sections 3.5 and 3.6. Finally, the strengths, weaknesses and limitations of the chosen approach are discussed in Section 3.7.

3.1. Choosing a Solution Technique

Phase I is concerned with developing a model that determines the capacity required to obtain a service level (also called shop performance) of at least 95% based on highly variable demand. Key criteria in choosing a solution technique for this application are:

- Computational speed
- Simplicity
- Transparency
- Flexibility

All four above-mentioned criteria are of importance for the implementation and use of the model. As stated in Chapter 2, the scope is limited to value measurement models due to the simplicity compared to both goal- and outranking models. Besides that, goal models require the existence of an ultimate solution, which in this case is not present, or unknown. Given the theoretical background on scheduling problems in Chapter 2, the choice is made to use a greedy algorithm, specifically a constructive greedy algorithm. One of the main reasons for this choice is the computational speed; the method provides fast results. Besides that, the greedy algorithm is relatively simple. Generally, more complex methods require more (and highly specific) assumptions, implying the accuracy for a specific application is high, but reducing possibilities for generalisation. Simpler methods are more easily transferable to other applications. Another key criterion on which the greedy algorithm performs well is transparency and flexibility. The main limitation for a greedy algorithm is the risk of yielding a local optimum instead of a global. For this application the optimum is the lowest possible capacity required to meet the demand (inflow). Since the approach is to start from a capacity of 1 and increase with every iteration, the risk of obtaining a local optimum rather than a global is minimised. A linear programming technique was also considered, however given the many inter-dependencies between variables and the reduced level of transparency and simplicity, a greedy algorithm is preferred. Also, as stated in Chapter 2, the majority of studies focus on implementation of linear programming techniques, while the use of a greedy algorithm is not as well covered by literature.

3.2. Approach

The approach to solving the above-mentioned problem is as follows: starting with a capacity of 1 and a certain inflow scenario, the model simulates the path each component follows in steps of 1 day. The time step of 1 day is chosen based on the available information and shop processes; the TAT is measured in days, incoming components is measured per day, and work is distributed and divided per day. Higher accuracy would therefore be unnecessary and require a significant change for both the shop as well as the supply chain processes.

Each component that enters the shop has a certain workscope, which corresponds to an expected Repair Process Time (RPT). The model assumes three abstract locations: repair, buffer, and overflow buffer. When a component is in repair, a mechanic is working on the component. The total amount of components in repair cannot exceed the available capacity. When a component is in the buffer, it is waiting for available capacity in repair. There is no theoretical limit on buffer-size, however with long Repair Process Times and limited capacity, the size of the buffer will most likely affect the shop performance. Finally, when a component has exceeded its maximum buffer time (maximum buffer time = contracted TAT - expected RPT), the component will be transferred to the overflow buffer. Depending on the priority procedure of choice, components can either first be taken from the buffer or the overflow buffer in case of available capacity in repair. The model runs until there are no components left in repair, and the shop performance is computed using Equation 3.1:

$$P = (1 - \frac{n_{overdue}}{n_{total}} * 100\%) \tag{3.1}$$

Where,
$n_{overdue}$ = number of components in the overflow buffer
n_{total} = total number of components in chosen time interval

If the shop performance is lower than 95%, the capacity will be increased by 1, after which the entire process is repeated. This continues until the performance is 95% or higher, resulting in the required capacity for the chosen inflow scenario. Figure 3.4 provides a visual representation of the approach in a flowchart. In order to provide a better overview of key simulation actions performed by the greedy algorithm, a visualisation of two days is presented in Figures 3.2 and 3.3. The full scenario for FIFO is shown in Appendix A, together with a visualisation of a different priority procedure: first-from-buffer. The example has the following characteristics:

- Inflow = [1 2 3 2 0]
- Capacity = 2
- Repair time = 3 days
- Contracted TAT = 5 days
- Maximum buffer time = Contracted TAT - Repair time = 2 days

Figure 3.1 shows a schematic representation of the abstract locations used in the greedy algorithm, which are used in Figures 3.2 until A.2 and Appendix A. Components entering on a given day enter in the *inbound*. In case there is no available capacity, the incoming components go directly to the *buffer*. The location from which the components are taken in case of available capacity depends on the priority procedure. In case of FIFO, If the *overflow buffer* is empty (Figure 3.2, step 2), components enter the repair directly from buffer. However, if there are components in the overflow buffer, they are given priority since they entered the shop first (Figure 3.3, step 2). In case of 'first-from-buffer', components are first taken from the buffer, regardless of the time spent in overflow, which can be seen in Figure A.2 in the Appendix. It should be noted that when a component enters one of the abstract locations its count restarts to 0.

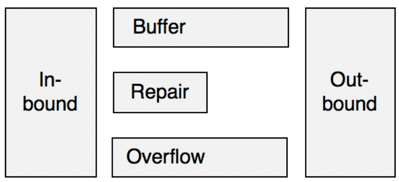

Figure 3.1: Schematic Visualisation of Abstract Locations Used in Greedy Algorithm

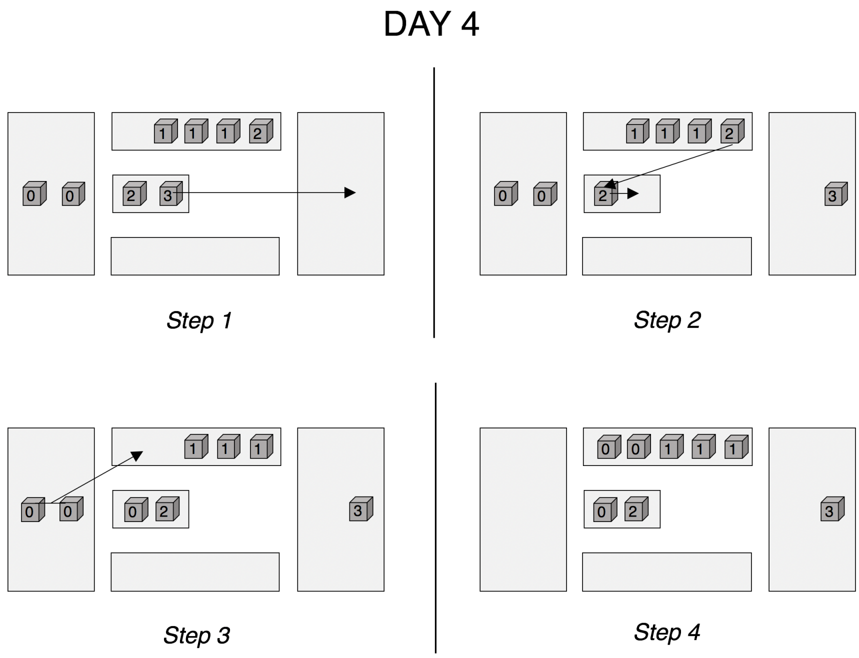

Figure 3.2: Visualisation of Key Simulation Actions Greedy Algorithm: Day 4 FIFO

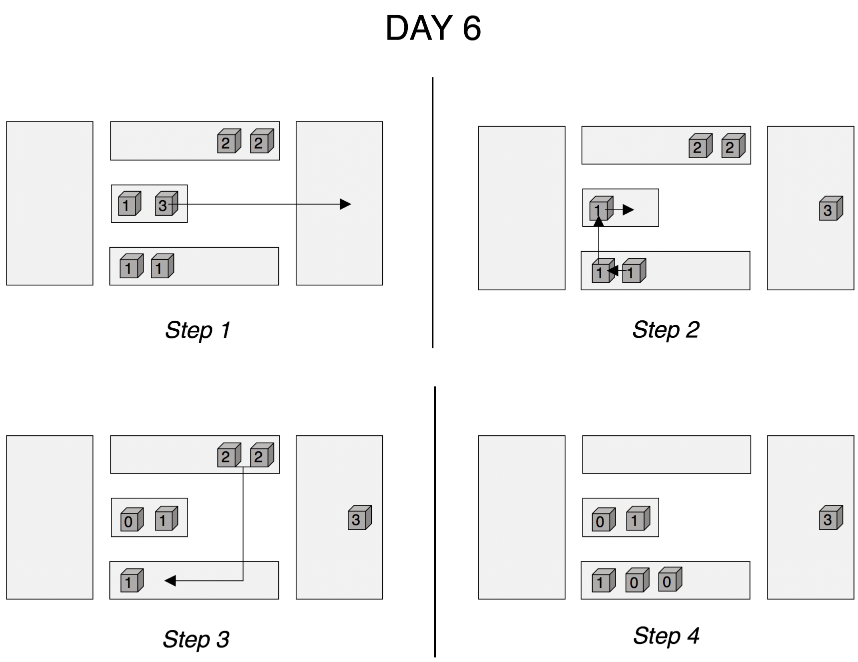

Figure 3.3: Visualisation of Key Simulation Actions Greedy Algorithm: Day 6 FIFO

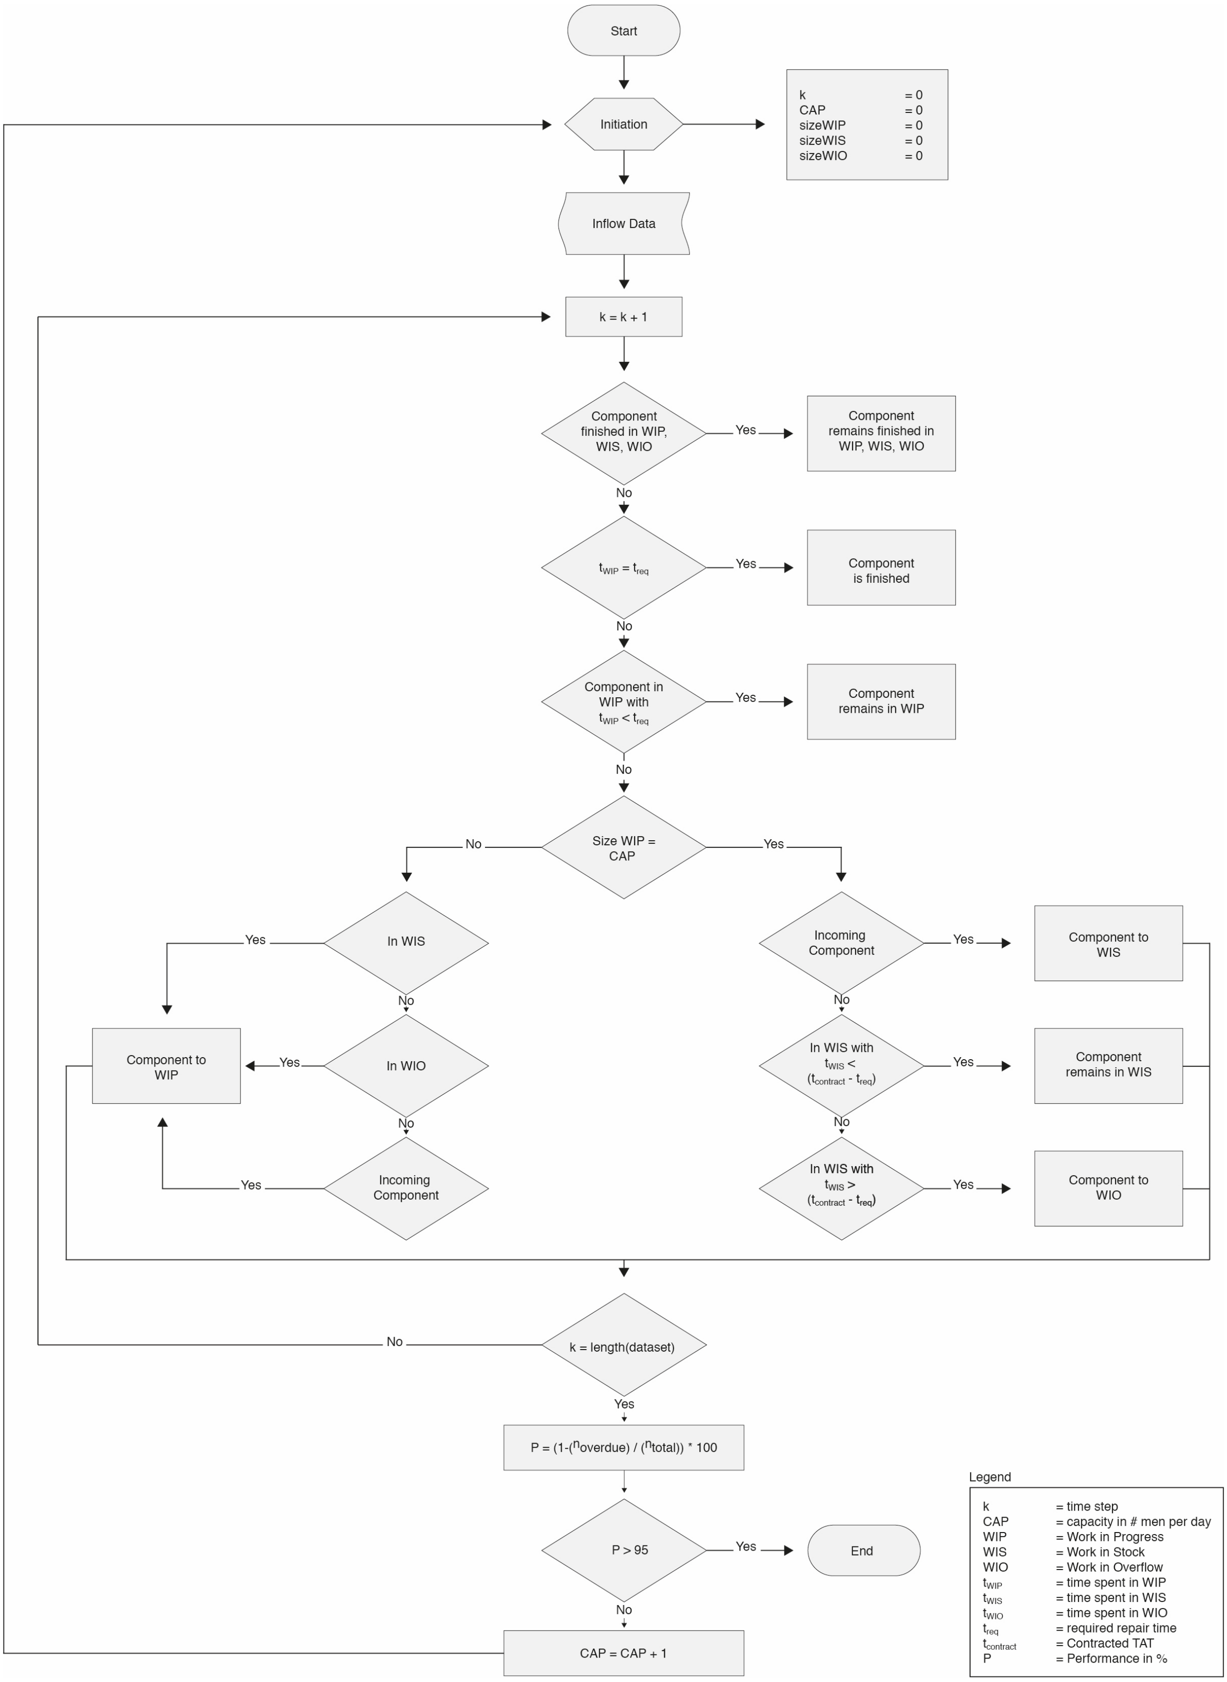

Figure 3.4: Flowchart for Greedy Algorithm Determining Required Shop Capacity

3.3. Assumptions

In order to develop this theoretical model, several assumptions have been made, which are divided into two categories: 1) assumptions that cause little to no deviation from the actual situation in the shop, and 2) assumptions that could, or will, cause discrepancies compared to the actual situation in the shop.

Assumptions with no or minor impact

1. The process considered starts when entering the repair shop and ends when leaving the repair shop
2. One technician works on one component full-time
3. Operations in the shop are shut down during the weekend
4. There is one shift with an effective duration of 6.5 hours
5. The productivity of the technicians in the shop is 80-85%
6. Every component entering the shop follows the same process
7. All technicians have the required skills to be able to repair and handle all components entering the shop
8. The inflow is unpredictable and includes large variation (independent variable)

Assumptions with possible or major impact

9. There is no physical limit on shop capacity
10. There are three main types of repair (minor, major, overhaul), of which the RPT is taken from historical distribution
11. There is one central buffer before the component enters the shop
12. Once a component is in repair, it follows a continuous process
13. Distribution regular/disrupted flow and percentage of disrupted flow resulting in overdue is based on historical data

While assumptions 1 through 8 are relatively straightforward and will result in small to no discrepancies from the actual situation, for assumptions 9 through 13 their effect should be elaborated on further. Starting with assumption 9, which is mostly dependent on the test equipment present in the shop. For the development of the model this limit is not taken into account, but should inflow levels grow significantly, this is an important factor to consider. Regarding the repair process time interval described in assumption 10, the historical data is used to determine the distribution of repair process time for the incoming component based on actual touch-time. While data is available on the expected touch-time for each of the components, the repair process time also includes waste in the process, which is not taken into account in the touch-time (further elaborated on in Section 3.4). For that reason assumption 4 and 5 are used, resulting in 15% waste in the process to be added onto the touch-time. Since the time step chosen for this model is one full day, adding 15% of waste to each touch-time simply resulted in the addition of one day. Implications on the results could be over- or under estimation of the required capacity, depending on the level of productivity. If productivity increases, less capacity is required, and vice versa. The effect of assumptions 11 and 12 in terms of TAT will be minimal, as the sequence is only theoretically altered, however moving the pre-test to after the buffer is often not desired as the pre-test is used to determine the workscope and should thus be performed as quickly as possible. The final assumption, number 13, states the distribution of regular/disrupted flow and its basis in historical data. Again, historical data is used to determine the distribution between regular and disrupted flow, and the percentage of components that end up overdue after being disrupted. This distribution is to be verified by checking it with multiple time intervals. It should be taken into account in the model that this distribution might change over time and should thus be updated with some frequency. Similar to the implications discussed for assumption 10, this can result in both over- or under estimation of the required capacity. Also, if the percentage of overdue from disrupted is decreased, the possibility for meeting the desired service level is increased, and vice versa.

3.4. Data Gathering & Reliability

Based on the approach discussed in this chapter, the following inputs are identified:

Inputs

- Historical inflow data in number of components per day
- Distribution of workscope based on historical data in number of days work
- Distribution of overdue from disrupted flow based on historical data in % of total incoming components
- Contracted TAT in number of working days
- Priority strategy to determine the sequence of component flow

The historical inflow data provides the model with the incoming components per day for a chosen time interval. The distribution in workscope and overdue from disrupted flow are both taken from historical data and are thus dependent on the application. The Repair Process Time (RPT) consists of the touch-time (automatically logged time spent on a component) and the waste in the process. In many cases the waste is not measured, meaning an assumption is made, often based on productivity. The contracted TAT in working days as well as the chosen priority strategy are also dependent on the application, and can be either based on historical, current, or future state. The model outputs are as follows:

Outputs

- Required capacity in number of **net** fte per day
- Shop performance in %

The reliability of the data-sources is important, since unreliable data can cause significant discrepancies between the results obtained from the model and reality. When looking at the inflow data, the reliability is quite high as it is often automatically registered. It is also highly unlikely that a component is not or wrongly entered, as it can otherwise not be handled in sequential steps in the process. The Repair Process Time has a higher degree of unreliability, due to the assumption of adding waste to the (automatically measured, and thus highly accurate) touch-time. Regarding the accuracy of the distribution of RPT used to simulate the workscope of incoming components; the accumulation of assumptions might cause significant discrepancies between simulation and reality. However, these discrepancies can cause skewed distributions to both a larger or smaller required workscope, averaging out the extremes. Regarding the distribution of overdue from disrupted, there will be discrepancies compared to the actual situation, as it is dependent on many factors. The data gathering method is reliable as it obtains data directly from the source.

3.5. Uncertainty Scenarios to Analyse the Relationship between Capacity and Shop-Performance

As previously discussed, the aim of Phase I is to determine the required capacity in the shop in terms of manpower given a certain inflow scenario. However, another important objective is to analyse the effect of capacity on shop performance in different scenarios. In order to see the impact on the on-time performance or required capacity to obtain the desired service level of > 95%, several uncertainty scenarios are to be tested, in which the following questions are to be answered:

1. Is there a pattern to be observed for different inflow scenarios?
2. What is the required capacity for each month of inflow in 2016/2017?
3. What is the effect on the performance for varying lengths of historical inflow data?
4. What is the effect on the required capacity for varying inflow growth or decline scenarios?
5. What is the effect on the performance/required capacity of working weekends?

6. What is the effect of disrupted flow (and overdue due to disrupted) on the performance/required capacity?
7. Is there any correlation between the maximum- and/or mean inflow and the required capacity?
8. What is the impact of highly variable inflow on the required capacity - compared to inflow with low variance?

In order to answer questions 1 through 3 the model is adapted by using different (time intervals of) inflow data, and analysing the results. Regarding question 4, the actual inflow of a certain time interval is identified as the baseline inflow scenario. This inflow scenario is then reduced or increased by 10% to 50% in steps of 10%, yielding 10 growth or decline scenarios, for each of which the required capacity is computed to maintain a service level of at least 95%. For question 5, the contracted TAT is increased by the number of weekend days in the chosen time interval. When considering disrupted flow, the model is to be adapted so that it can take into account a certain percentage of components that will be overdue, regardless of the capacity. Therefore, the overdue array E is not initiated as zero, but rather consists of a random distribution of *NaN*, the amount equal to the percentage of *too late due to disruption* out of the total incoming components, based on the historical distribution. Question 7 is aimed at providing insight into the ability of the shop to handle peaks, and to find a relationship between mean or maximum inflow and the required capacity. Its answer can be found by plotting the required capacity for several time intervals against the mean and maximum inflow of that time interval. Finally, question 8 focuses on the impact of high variance in demand. This is done by running the model for a scenario with 0 variance and analyse the effect on the required capacity.

3.6. Required Shop Capacity and the Effect on the Supply Chain

As stated in Chapter 2, the two key performance indicators for KLM are: 1) Customer Service Level (CSL), and 2) financial risk (or cost). Increasing the shop capacity results in an increase in shop performance, and an increase in shop performance yields a higher CSL. This in turn results in additional savings by reducing the required stock levels. The relationship between shop capacity, performance, CSL, and required stock level is a complex one that is dependent on many conflicting parameters. This research project does not aim to provide an in-depth analysis on these relationships, rather it uses the dependencies to provide insights on the possible effects of increasing shop capacity in other parts of the supply chain.

The relationship between shop SL and CSL is relatively straightforward if translated to TAT; a decrease in shop TAT with x days results in x days reduction in the total supply chain, assuming other parameters remain unchanged. In order to make the link between shop SL and TAT reduction, the current shop TAT is computed for the top components and compared to the expected TAT when utilising the required capacity in the shop. Financial risk consists of several factors, the most important ones being the required number of units in stock, and the risk of lease-in or buy-in of components. For this research project, the focus is on the first, as there exist a more direct relationship between the TAT and the required number of units in stock. The latter introduces significant uncertainty due to possible Aircraft-On-Ground (AOG) notifications and external factors. In order to determine the relationship between reduction in shop TAT and the required number of units in stock, a user function is to be developed for the top components, which can be used to compute possible savings. Summarising, the following steps are to be taken to obtain an indication of possible effects in the supply chain:

1. Determine required capacity for different time intervals based on historical data
2. Determine relationship between shop SL and shop TAT
3. Compute difference between old and new situation in terms of shop TAT
4. Determine relationship between TAT reduction in the shop and required stock level of top components
5. Compute possible savings

3.7. Strengths, Weaknesses & Limitations of Approach

As with any approach, there are certain strengths and weaknesses. Starting with the strengths, the most obvious one is simplicity. The greedy algorithm is straightforward, simple, and easy to use. Besides that, the number of required inputs is limited and the model handles multiple types of inflow scenarios - not limited to highly variable demand - by simulation the path of each component per day and adapts the required capacity based on that specific inflow scenario. Another strength of this approach is that the model can be adapted and expanded to incorporate higher complexity, for example in case of a different priority procedure in the shop. The algorithm is thus also fairly flexible and a good foundation to solve similar problems. A characteristic of this approach, as for many others, is that the output is as good as the input. This can be seen as both a strength as well as a weakness; a strength since its performance is not limited by the model itself, but also a weakness as it is highly dependent on the quality of input and the user. Another weakness of this approach is that a greedy algorithm chooses the alternative that at a specific moment is the most optimal. This means there is a risk that instead of finding the global optimum, the algorithm outputs a local optimum. For this application however, the risk of obtaining a local optimum rather than a global optimum is very small. The reason being that the algorithm runs until the required performance of at least 95% is obtained. The optimum in this scenario is the lowest capacity for which this will occur, and since the algorithm updates capacity by 1 every iteration, there is no other optimum than the first one found. Another weakness of the model is that it uses historical data and distributions as inputs, which can can discrepancies with the actual situation. Therefore it is important to continuously monitor the inputs and adapt if necessary. A final limitation of the approach is that the link to the supply chain is limited to the effect on the stock levels, excluding lease-in costs, and the effect on the customer service level.

4

Implementation Phase I: Optimisation of Shop Capacity

The aim of this chapter is to cover the implementation of the previously discussed methodology and approach for Phase I of the research project: a greedy algorithm to determine the required capacity in the presence of highly variable demand. Section 4.1 discusses the initialisation of the model, followed by the greedy algorithm in Section 4.2. Finally, the verification of the model is discussed in Section 4.3.

4.1. Initialisation

The initialisation of the greedy algorithm aims to provide the correct inputs into the model. As discussed in Chapter 3, the key inputs are the historical inflow data, contracted TAT, distribution of workscope (also Repair Process Time (RPT) and overdue from disrupted, and priority procedure. The historical inflow is taken from an Excel-file, and the contracted TAT is a fixed value based on industry standards. Both the distribution of workscope and percentage overdue from disrupted are taken from historical data. To implement such a distribution in the model the following steps are taken:

1. Probability of occurrence $p_1, p_1, ..., p_n$ from historical data
2. Total number of incoming components in time interval n_{total}
3. Total expected number of event $1, 2, ..., n = p_1 * n_{total}, p_2 * n_{total}, ..., p_n * n_{total}$
4. Create vector V length n_{total} with correct distribution of events
5. Randomise sequence of events in vector V
6. For each incoming component i, corresponding event is $V(i)$

The RPT distribution is linked to the incoming component once it enters the shop. The distribution of overdue from disrupted is only used if disrupted flow is taken into account, and yields an expected number of components to be overdue regardless of the available capacity. First the total expected number of components overdue from disrupted are to be computed using the historical distribution. These occurrences are then randomly distributed over the total number of incoming components within the chosen time interval. The final input is the priority procedure, which will require adaptation of the greedy algorithm and shop sequence simulation, as shown in Section 4.2.

Finally, several other initial conditions need to be specified. In any case the capacity starts at 0 as the aim is to update the capacity every iteration until the desired performance of at least 95% is obtained. Other parameters are the number of components in repair, buffer, and overflow; each combination of which corresponds to a different initial condition for the model. Also, performance has an initial value of 0, which corresponds to running the model with a capacity of 0.

4.2. Greedy Algorithm

As previously discussed there are two priority procedures to be modeled. Starting with FIFO, the pseudo-algorithm is shown in Algorithm 1, consisting of three procedures. The first is to increase capacity by one in case the performance < 95%. The next step is to simulate the shop process for each component per day, which is shown in lines 5 through 20. After this has been completed the model outputs the capacity and the performance. This process is repeated until the desired performance is met.

Algorithm 1 Greedy Algorithm: Determining Required Shop Capacity - FIFO

```
1: procedure OPTIMISE
2:     while performance < 95% do
3:         procedure INCREASE CAPACITY
4:             capacity = capacity + 1
5:             for time = 1:length(time) do
6:                 for component = 1:sum(inflow) do
7:                     procedure SIMULATION
8:                         if component = completed then
9:                             remain completed
10:                            if component in repair then
11:                                time in repair = time in repair + 1
12:                                if available capacity then
13:                                    component to repair
14:                                else
15:                                    if component in buffer with time in buffer < (contracted TAT - RPT) then
16:                                        time in buffer = time in buffer + 1
17:                                    else if component in buffer with time in buffer = (contracted TAT - RPT) then
18:                                        component to overflow buffer
19:                                    elsecomponent in overflow buffer
20:                                        time in overflow = time in overflow + 1
21:
22:                                procedure RETURN
23:                                    performance = (1 - (components in overflow)/(total components)) * 100%
24:                                    capacity
```

Besides the base model shown in Algorithm 1 in which FIFO is assumed, a second version is developed. Version 2 of the model includes a different priority procedure in which components are first taken from the buffer in order to maximise performance, as shown in pseudo-algorithm 2. The first difference is that a new component is first sent to the buffer (line 14) in order to ensure that new components are not directly sent to repair. The next step is to check the size of the buffer; if 0 this means that the buffer is empty and the component can be send to repair from the overflow buffer. If the size of the buffer is not 0, the first component from the buffer should enter repair. After line 22 both codes are equal.

Algorithm 2 Greedy Algorithm: Determining Required Shop Capacity - BUFFER FIRST

```
procedure OPTIMISATION
    while performance < 95% do
        procedure INCREASE CAPACITY
            capacity = capacity + 1
            for time = 1:length(time) do
                for component = 1:sum(inflow) do
                    procedure SIMULATION
                        if component = completed then
                            remain completed
                            if component in repair then
                                time in repair = time in repair + 1
                                if available capacity then
                                    if new component then
                                        component to buffer
                                    if size buffer = 0 then
                                        component to repair
                                    else
                                        if component in buffer then
                                            component to repair
                                        else component in overflow
                                            time in overflow = time in overflow + 1
                                else
                                    if component in buffer with time in buffer < (contracted TAT - RPT) then
                                        time in buffer = time in buffer + 1
                                    else if component in buffer with time in buffer = (contracted TAT - RPT) then
                                        component to overflow buffer
                                    else component in overflow buffer
                                        time in overflow = time in overflow + 1

                                    procedure RETURN
                                        performance = (1 - (components in overflow)/(total components)) * 100%
                                        capacity
```

4.3. Verification of Greedy Algorithm

In order to verify the model several scenarios are developed specifically to test several parts of the code. Critical parts of the code to be tested are:

- Sequence of for-loops
- If-statements and their sequence
- Computation of buffer-, repair-, and overflow size
- Computation of performance
- Adaptation of capacity when performance is insufficient

The method of verification used for this model is to manually write out scenarios for each time step and compare the results in each of the arrays (buffer, repair, and overflow) to the ones generated by the model. Besides that, the overall performance and is checked and compared with the required capacity for SL > 95%. Table 4.1 shows different scenarios and their purpose.

Table 4.1: Verification Scenarios for the Greedy Optimisation Model

Scenario	Inflow	Required Repair Process Time [days]	Capacity [technicians]
1	[1 2 0 3]	3	2
2	[4 2 0 5 2]	3	3
3	[2 0 0 0 3 1 0 0 0 1]	3	2
4	[8 4 7 2 0 6]	3-7	3,5,8
5	March 2017	3-7	1:20

Scenario 1 is used as a simple test case while developing the model. Scenario 2 and 3 are used to see the effect of higher inflow, capacity, and intermittent demand on the workings of the model. Scenario 4 introduces a variable repair time and higher inflow to check the impact of available capacity and repair time on the performance. Finally, scenario 5 uses real inflow data to obtain first rough results on accuracy. For all scenarios the model showed the same results as the ones obtained by manually going through the process, from which can be concluded that the model works well and yields the desired results.

5 Case Study

As mentioned in Chapter 1, this research project is a collaboration between TU Delft and KLM Engineering & Maintenance. Given the methodology and implementation strategy of Phase I, the aim of this chapter is to provide background information on the specific case on which the methodology will be applied: Shop EWF. The chapter starts with general information on the shop and processes in Section 5.1. This is followed by a short inflow analysis to further highlight the specific problem of Shop EWF and by extension the KLM Component Services supply chain in Section 5.2. Finally, more detailed information is provided and the scope of the case study is presented in Section 5.3.

5.1. General Shop Information

As discussed previously, the practical scope of this research project is Shop Hydraulics 1 - Workcentre EWF. This shop is characterised by high business value components; key components being the Integrated Drive Generator (IDG), Back-Up Generator (BUG), and the Variable Frequency Starter Generator (VFSG). These are complex components that require highly skilled personnel and have a relatively long Repair Process Time (RPT) between 2 and 7 days depending on the workscope and type of component. Combining this with the large variation in inflow and limited capacity in terms of manpower, the main challenge is to obtain and maintain a stable shop performance and meet the required shop TAT.

Another important characteristic of Shop EWF is that the majority of the inflow comes from customers with a Time & Material (T&M) contract. For customers from KLM's Component Services Division, two main types of contract exist. The first, 'pool' customers, pay a certain fee for both maintenance of unserviceable components as well as the availability of serviceable components. The pool of spare (serviceable) components is owned and maintained by KLM. The second type of customers pay repair cost per unserviceable component they want KLM to repair, and thus do not pay for availability, meaning the components are not owned by KLM. This complicates the situation for KLM, since these components can not easily be outsourced (permission from the customer is required), and reducing TAT on these components does not yield a direct benefit for KLM. On the other hand, in general, more profit is made on the repair of T&M components compared to pool components.

The shop process is visualised in Figure 5.1. Once the component is handled by the Customer Interface Repair Officer (CIRO), the component enters the shop process, after which it awaits the pre-test. This buffer-time is relatively short; a component on average spends less than half a day here. The pre-test is used to determine the workscope and obtain an indication of the RPT. After the pre-test is completed, the component is put in the main buffer, where it waits until it is picked up by a mechanic to start the repair. This buffer is the 'large' buffer, where components can spend several days up to weeks, depending on the availability of manpower, material, and/or other external factors. Once a component is picked up by a mechanic, the process is fairly continuous up until the repair is completed. After completion, the component undergoes final testing and, in case of successful results, is signed out by the CIRO. This is where the shop process ends and the responsibility of the component is handed over to another entity.

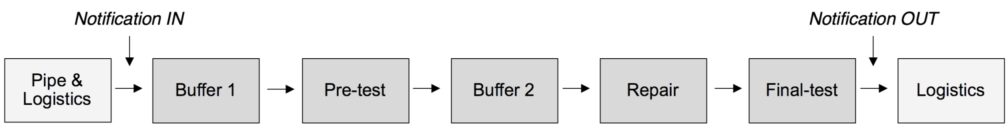

Figure 5.1: Visualisation of Shop Process in Shop EWF

5.2. Analysis of Inflow

Figure 5.2 shows the total number of incoming components per day (excluding weekends when the shop is closed) for the year 2017. For this year (2017), the total number of incoming components per component type are shown in Figure 5.3. The top 5 components are responsible for more than 80% of the total inflow.

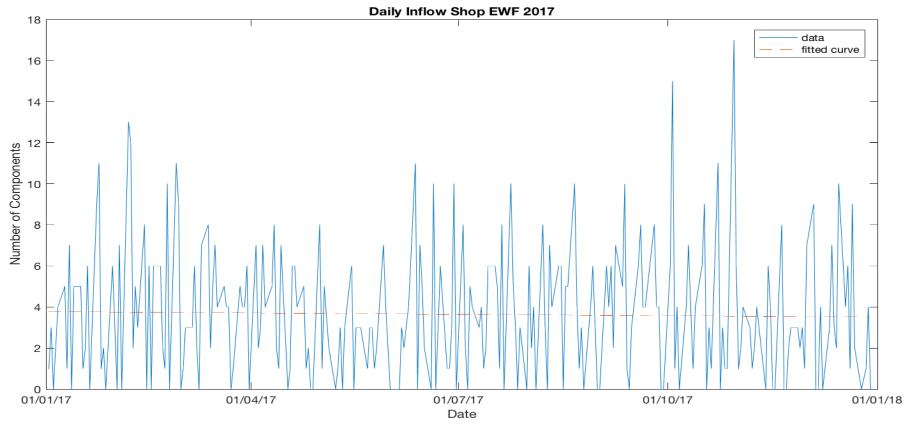

Figure 5.2: Daily Inflow Shop EWF 2017

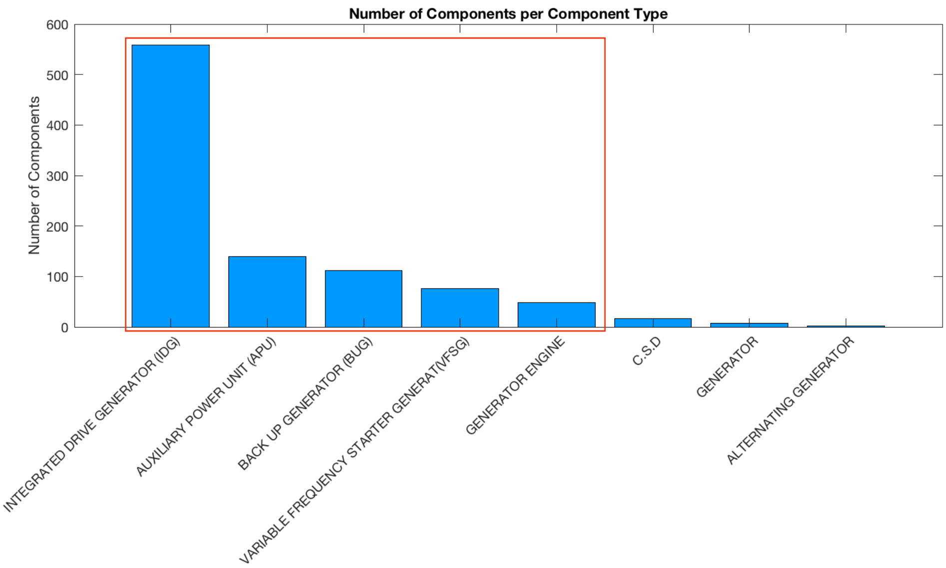

Figure 5.3: Total Number of Incoming Components per Type 2017

Using demand characteristics introduced in Chapter 2, the inflow pattern is categorised as smooth. However, the variation in inflow is relatively high, with several peaks larger than twice the average inflow per day. Variation in itself is not necessarily a problem if there is sufficient buffer in the process to handle the peaks. In shop EWF however, the combination of long RPT (Table 5.1) and high inflow peaks result in large backlogs of work that require a long time to be processed given the limited amount of capacity in terms of manpower of 13 net fte per day. This is represented in Figure 5.4a, in which the weekly inflow together with the number of components in stock is shown. No data from the first 20 weeks was available, therefore the 'work in stock' graph starts in week 21. Similarly, Figure 5.4b shows the weekly inflow and service level over 2017.

Table 5.1: Historical Repair Process Time Shop EWF 2017

RPT	Percentage of Total
0	0
1	17
2	12
3	19
4	17
5	14
6	8
7	6
> 7	5

It can be seen that the number of components in stock - or the backlog - continuously grows until the end of the year when the inflow is slightly lower. In general, a trend is to be seen: after a period of low inflow the number of components in stock decreases slightly, and after a period of high inflow the number of components in stock increases. Assuming a certain minimal and maximal throughput of the shop this is to be expected. When looking at the SL, the first observation is that it varies greatly; while theoretically the shop performance should be stable around 95%, in this case the shop level is highly erratic. Comparing both figures, if the backlog increases in size, the SL decreases.

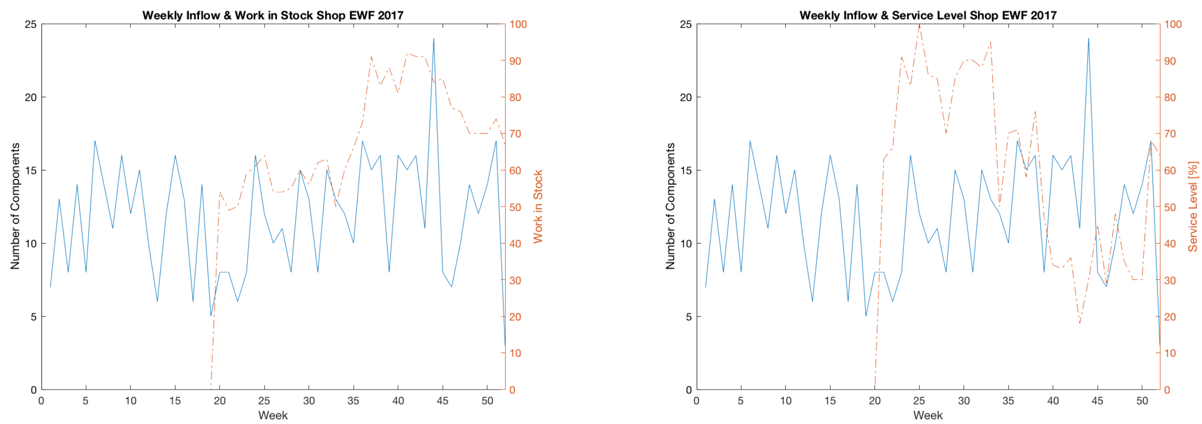

Figure 5.4: Weekly Inflow vs. Work in Stock & Weekly Inflow vs. SL 2017 Shop EWF

As mentioned in Chapter 2, one way to control the incoming flow is to use a forecasting model to predict demand. To test whether this is a possibility for Shop EWF, several time series modeling methods and a stochastic method are implemented, of which the methodology and results can be seen in Appendix A. Using the Mean Average Deviation as an error metric it was found that expecting the average yields minimum deviation between the actual and predicted inflow. This means the use of a complex forecasting tool is unnecessary. However, the minimum deviation is still 3 components on an average inflow of 4 components per day. This suggests that even the best scenario (the average) performs poorly at predicting incoming flow. This supports the conclusion made in Chapter 2: predicting demand in case of highly variable inflow is very complex and rarely yields good results.

Table 5.2 shows the current turnaround time for all Shop EWF components for the year 2017. Separate scenarios are considered including or excluding disrupted flow, and/or Time & Material components, dependent on the application.

Table 5.2: Turnaround Time 2017 for Varying Scenarios Shop EWF

Scenario	TAT reg + dis	TAT reg	TAT pool reg + dis	TAT pool reg
All components 2017	20	17	16	17

5.3. Scope of Case Study

As stated above, the scope for this research project is on Shop EWF at KLM Engineering & Maintenance Component Services Division. Currently, the net average capacity in the shop is 13 fte per day. The contracted TAT depends on the contract type, and varies between 15 and 28 calendar days. In the new (CS2.0) situation, KLM aims to move towards the new industry standard of 14 calendar days. Given the information from Figure 5.3, the scope is limited to the top 5 components in terms of volume. However, the APU is no longer in repair at KLM, and the generator engine is excluded due to lack of incoming components from mid-2017 up until mid-2018. Therefore the scope is limited to the IDG, BUG, and the VFSG, with specific focus on the VFSG due to the expected growth in terms of Boeing 787 contracts. Regarding the time interval used for the case study, 2017 is considered since it is the latest complete calender year since starting this project. It might be possible that one full year results in a long computational time for the models, and for that reason the period January up to and including March 2017 is selected as reference. This period includes high peaks, an average of 4 components per day, and a representative distribution of incoming components, repair types, and regular vs. disrupted flow.

Key Performance Indicators (KPIs) in the shop are the shop performance (or service level), number of components in stock (the buffer), and the distribution regular/disrupted flow. In this research project, part of the objective is to make a link between the shop performance and the entire supply chain. In the supply chain, there are two main KPIs: 1) Customer Service Level (CSL), and 2) financial risk. CSL represents the percentage of on-time deliveries of serviceable components to the customer, and financial risk includes required stock, investments, and lease-in of components. Both these KPIs benefit from lower turnaround time (TAT). The focus in this research project is on shop performance, translated to TAT, and financial risk in terms of required units in stock. Table 5.3 provides additional information regarding in- and outflow of components, that are of importance for the implementation and obtaining of results based on this case study.

Table 5.3: Key Numbers and Percentages Regarding In- and Outflow of Components Shop EWF

Description	**Count**	**Percentage**
Total Incoming Components	1027	100
Regular Flow	752	73
Disrupted Flow	275	27
Outsourced	38	0.04
Pool Components	403	40
Time & Material Components	624	60
Overdue (total)	360	35
Overdue (disrupted)	206	75

6
Results & Discussion Phase I: Optimisation of Shop Capacity

Based on the scenarios discussed in Chapter 3 there are several results to discuss. The chapter starts with the results of the uncertainty scenarios in Section 6.1, in which the effect of multiple parameters on the shop performance are discussed. This is followed by the computation of required capacity in Shop EWF based on historical data, and the possible savings in the supply chain in Section 6.2.

6.1. Results of Uncertainty Scenarios

Starting with Figure 6.1 which provides a simple relationship between available capacity and performance in shop EWF based on March 2017 inflow data. It also includes the difference in performance between regular and disrupted flow. The pattern shown in both plots is that performance increases fairly quickly per additional unit of capacity. Once performance reaches 90% the graph stagnates and relatively many units of capacity are required to reach a performance of 100%. When looking at the disrupted flow a similar pattern can be observed, however due to the fact that 23% (disrupted) * 75% (overdue from disrupted) = 17% of all components ends up overdue regardless of the level of capacity due to them being disrupted. This immediately answers question 6: the effect of disrupted flow, more specifically overdue due to disrupted flow, on the shop performance is significant, and cannot directly be improved by increasing capacity.

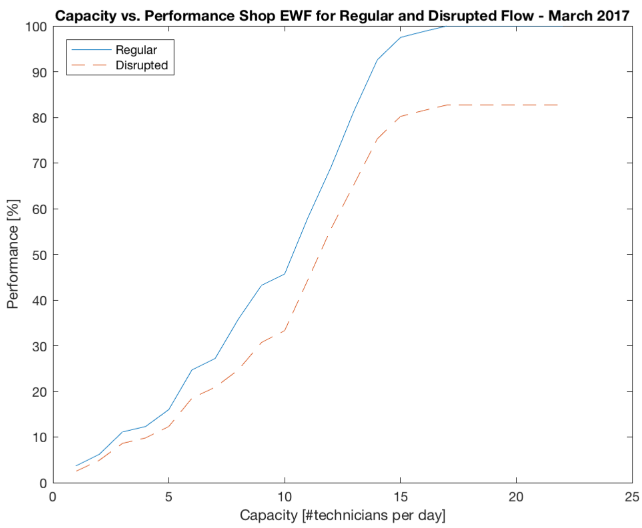

Figure 6.1: Capacity vs. Performance Shop EWF Based on March 2017 Inflow: Regular vs. Disrupted Flow

Figure 6.2 shows similar graphs for inflow of several months in 2017 to observe whether or not the pattern repeats itself. It can be seen that the pattern is similar for every month of inflow, the only difference being in the steepness of the curve, which is dependent on the number and interval of components coming in, as indicated in Table 6.1. The May-curve is steepest, and has the lowest average inflow per day, as well as the lowest maximum inflow per day. October on the other hand has the highest average inflow per day, as well as the highest inflow peak.

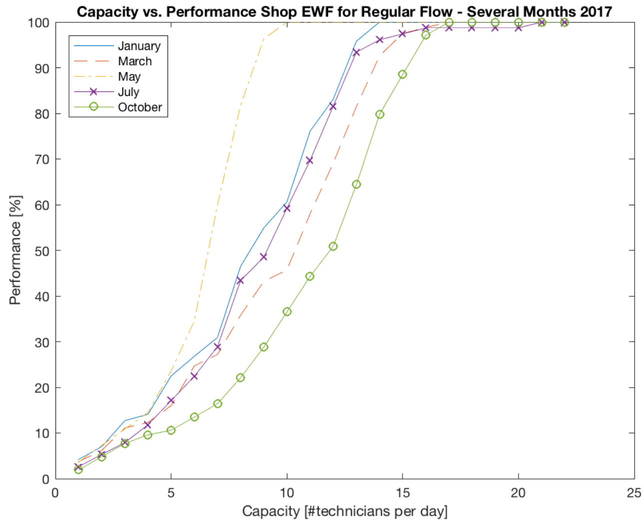

Figure 6.2: Capacity vs. Performance Shop EWF Based on Inflow from Several Months 2017: Regular Flow

Table 6.1: Average Inflow, Maximum Inflow, CV^2, and ADI Values for Varying Inflow Scenarios

Time Interval	**Avg. Inflow/day**	**Max Inflow**	**CV^2**	**ADI**
January 2017	4.3	11	0.46	1.19
March 2017	3.5	8	0.22	1.16
May 2017	2.8	8	0.39	1.8
July 2017	3.6	10	0.32	1.8
October 2017	4.7	17	0.72	1.17

When looking at a longer time interval of three months, in this case from January up to and including March, again a similar pattern can be observed, as shown in Figure 6.3. The main difference compared to the shorter time interval can be found in the steepness of the curve for low capacity. This can be attributed to the larger degree of accumulation in the buffer and overflow buffer in case of low capacity for more incoming components, which is the case when considering a longer time interval of historical data. This also results in a higher capacity required to obtain a shop performance of > 95%, compared to one month of inflow data.

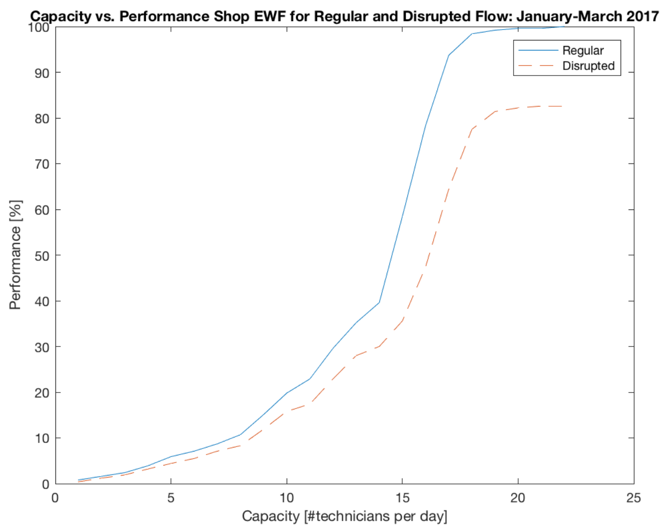

Figure 6.3: Capacity vs. Performance Shop EWF Based on Jan-Mar 2017 Inflow: Regular vs. Disrupted Flow

Moving on to the second question: *"What is the required capacity per month of inflow in 2016 and 2017 - and can a pattern be distinguished?*. Figure 6.4 shows the required capacity per month for a SL of > 95%, both for 2016 and 2017 inflow data. Only regular flow is considered, as with the current distribution of overdue out of disrupted, the shop performance will not exceed 83%.

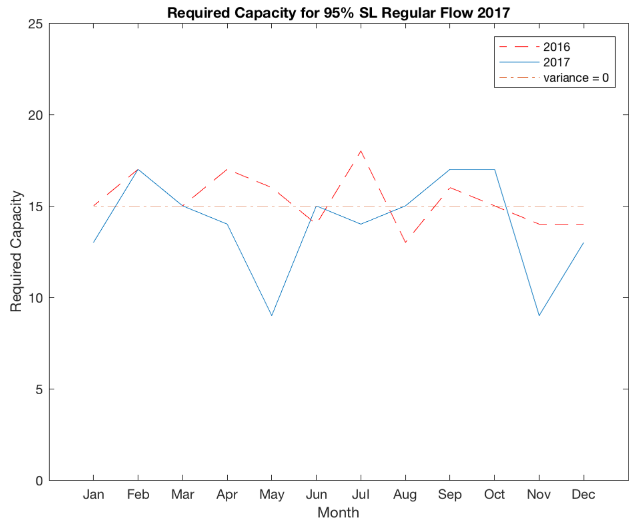

Figure 6.4: Required Capacity per Month for SL > 95%, 2016 and 2017 Inflow Data

The first observation based on Figure 6.4 is the difference in required capacity for the year 2016 vs. 2017 shows that there is no clear pattern or seasonality to be observed. Besides that, there is a large variation in required capacity per month. While in May 2017 the required capacity is only 9, in July of 2016 the required capacity has doubled to 18. In order to obtain information on the correlation between the required capacity and key characteristics of the inflow data, Figure 6.5 shows the correlation between the required capacity and the maximum and mean inflow. Also, Figure 6.4 indicates the required capacity when the demand (or inflow) has no variation: 4 components enter the shop every (work) day. In that case, the required capacity is constant at 15 net fte per day, which is significantly lower than the required 18 in case of highly variable demand. The reason for this is that there are no peaks and/or backlog that needs to be handled, since there is a relatively constant outflow of the shop. This immediately answers the last question: *"What is the impact of highly variable inflow on the required capacity - compared to the inflow with low variance?"*: the effect in this case is 3 fte per day on average.

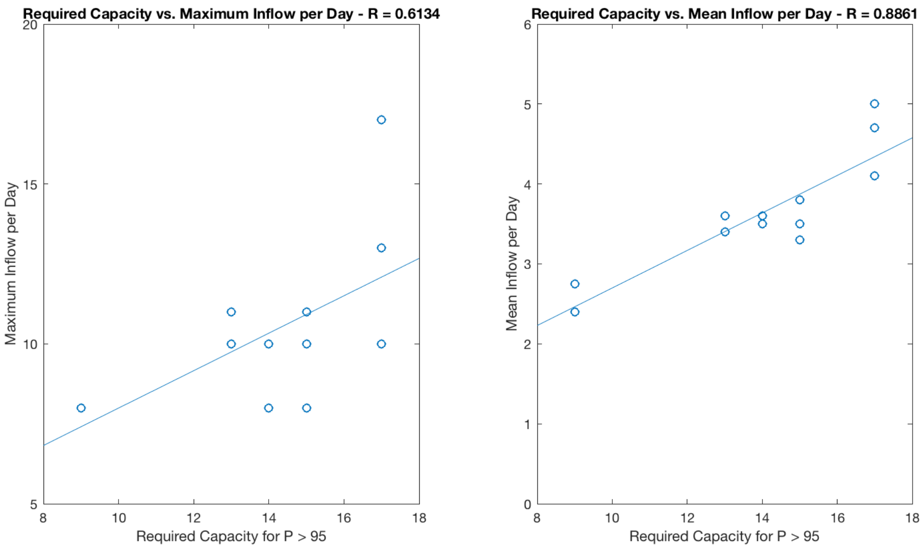

Figure 6.5: Correlation between Required Capacity and Key Characteristics of Inflow Data 2017

Another uncertainty scenario to be evaluated is the required capacity in case of growth or decline of inflow, shown in Figure 6.6. As a baseline the inflow for the time interval between January and March 2017 is taken. To simulate an increase or decrease in inflow, the number of incoming components is multiplied by a factor 0.5 up until 1.5, and rounded up. The next step is to run the optimisation model and determine the required capacity for every scenario. It can be observed that the relationship is almost linear.

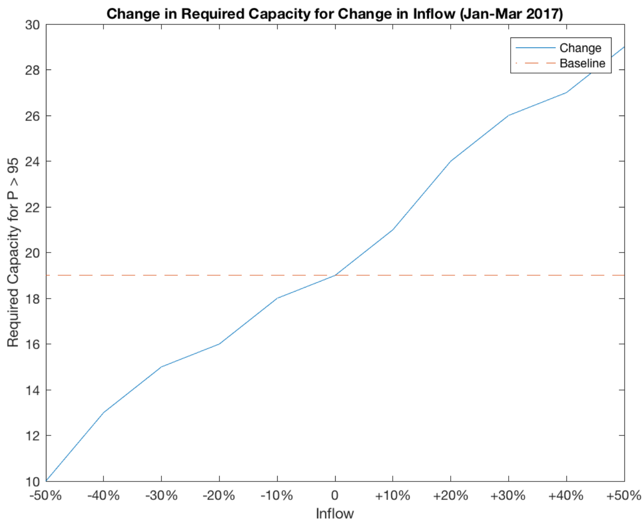

Figure 6.6: Required Capacity for Multiple Growth Scenarios (Baseline Jan-Mar 2017)

The result of the next scenario can be see in Figure 6.7, in which the effect of working weekends is shown compared to the current situation. On average, the performance increases by 10% regardless of the capacity. This is achieved by increasing the contracted TAT from 10 workdays to 14. Similar effects can thus be expected if the repair process time is reduced, or productivity is increased.

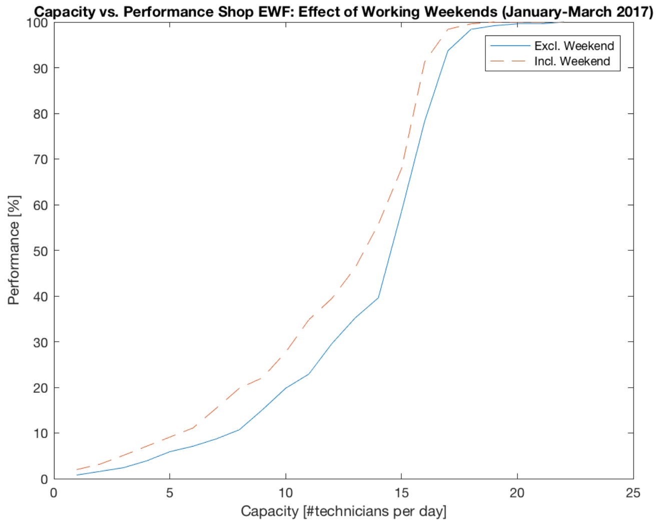

Figure 6.7: Effect of Working Weekends for Jan-Mar 2017 Inflow

Finally, Figure 6.8 shows the capacity versus performance for two different priority procedures in the shop. The blue curve indicates the First-In-First-Out (FIFO) procedure, while the red curve indicates performance in case components are first taken from buffer, regardless of when they entered the shop. As shown in Figure 6.8, the difference in performance is significant. The angle of inclination for the red curve is very large up until a capacity of 12, after which it slowly stagnates and crosses the blue curve at a capacity of 17. Especially in the region of low capacity the red curve is almost opposite to the blue (FIFO) curve. This can be explained by the fact that the theoretical performance is higher as less components enter the overflow buffer, since they are taken from the regular buffer first. However, this yields a significant increase in average days overdue, up to 70 days in case of the Jan-Mar 2017 inflow scenario.

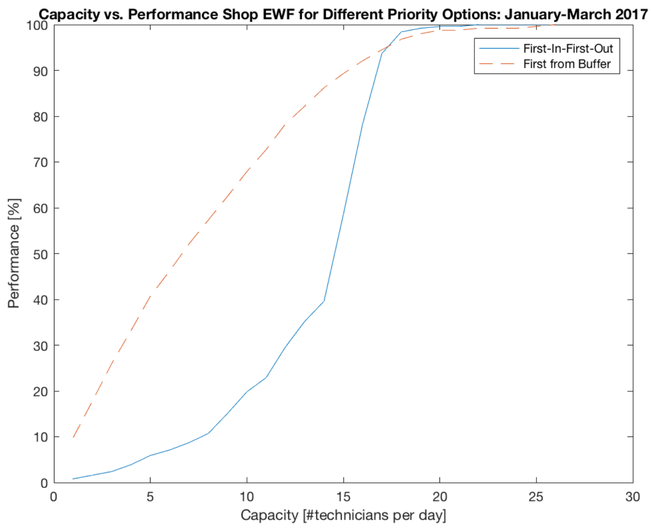

Figure 6.8: Effect of Different Priority Procedures on Performance, Inflow Jan-Mar 2017

6.2. Link to Supply Chain

Now that the relationship between shop capacity and performance has been established, this section is concerned with determining the optimal shop capacity by researching the link between shop performance and required stock in the supply chain. As described in Chapter 3, the following steps are to be taken:

1. Determine required capacity for different time intervals based on historical data
2. Determine relationship between shop SL and shop TAT
3. Compute difference between old and new situation in terms of shop TAT
4. Determine relationship between TAT reduction in the shop and required stock level of top components in shop EWF
5. Compute possible savings

Starting with step 1; the determination of required capacity for multiple time intervals, which is already initiated in Phase I as it is part of the uncertainty scenarios that are researched. Based on the available data, the following time intervals are analysed: 1, 3, 6, 12, and 24 months. The result for required capacity per month can be seen in Figure 6.4, from which the maximum required capacity is 18 in July 2016. Table 6.2 shows capacity and corresponding service level for 1, 3, 6, 12, and 24 month intervals taken between January 2016 and December 2017. It should be noted that for each of the time intervals, the worst case scenario is represented in Table 6.2.

Table 6.2: Required Capacity for Varying Time Intervals

Time Interval	Required Capacity for SL > 95%	Service Level [%]
1 month	18	99.4
	17	94.5
3 months	19	98.7
	18	74.2
6 months	19	99.2
	18	59.7
12 months	19	96.6
	18	89.2
24 months	19	96.6
	18	76.3

It can be concluded that, when considering 3, 6, 12, or 24 month time intervals, the required capacity is 19. The difference compared to the required capacity in case of 1 month intervals can be attributed to the accumulation of components that occurs when considering a longer period of inflow. Another important observation is the difference in performance in case of reducing the capacity by 1. For every time interval (except for 1 month) the performance drops significantly, in the worst case with more than 30%. This implies that the optimal capacity is at least the required capacity of 19. Given the current net capacity in the shop of 13, this means an additional 6 fte is required to meet the new desired capacity. With an average yearly cost per fte of USD 60,000, this sums up to a yearly cost of USD 360,000.

The next step is to provide insight on effects in the supply chain in case the net capacity is increased to 19. While the consequences in the supply chain are indirect, they can accumulate to have a significant impact. The benefits are mostly caused by a high, stable, and reliable shop performance, which leads to shorter shop (and thus end-to-end) TAT. For most components in the KLM pool, the number of required units in the supply chain is partially dependent on the end-to-end TAT. Other parameters are:

- Number of contracted aircraft
- Mean Time Between Removal (MTBR)
- QPA: number of units per aircraft

- Flight Hours (FH): number of flight hours in the pool
- TAT: assumed end-to-end TAT
- Units in stock

The expected service level is based upon the above-mentioned parameters using the relationship in Equation 6.1, based on the Poisson distribution. Using this equation and the requirement of an expected SL of >95%, the required number of units in stock can be determined for varying TAT.

$$SL = \sum_{k=0}^{x} \frac{e^{-\lambda}\lambda^{x}}{k!} \tag{6.1}$$

Where,
x = units in stock-1
λ = number in pipeline = TAT(days) * expected removals
$expected\ removals = \frac{FH*QPA}{MTBR}$

As shown in Chapter 5, the top 3 components (IDG, BUG, VFSG) are responsible for 80% of the inflow. However, the BUG is a Time & Material component only, meaning it does not have any relation to the component pool and thus stock levels of KLM E&M. Therefore, the scope is limited to the IDG and the VFSG. The result can be seen in Figure 6.9, indicating that with an R^2 value > 0.95, linearity can be assumed.

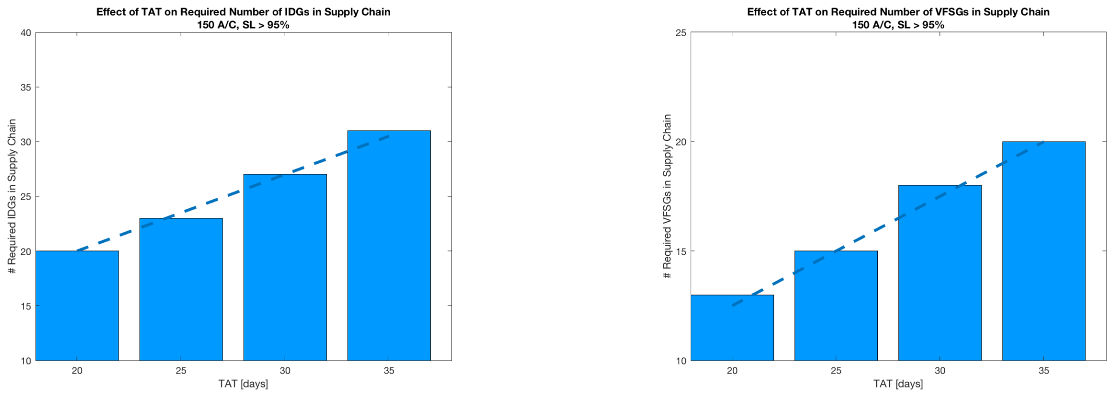

Figure 6.9: Relationship between end-to-end TAT and Required number of IDGs and VFSGs

This results in Equation 6.2 and 6.3 for the IDG and VFSG, respectively. Here, y is the number of required units in the supply chain, and x is the end-to-end TAT in days.

$$y = 0.7x + 6 \tag{6.2}$$

$$y = 0.5x + 2.5 \tag{6.3}$$

In order to determine possible savings in stock levels, the difference in shop TAT is computed between the current and new situation. Shop data is used from SAP that indicates the average shop TAT per component type - for regular flow and pool components only, excluding outliers. The new expected shop TAT is 14 days, as this is the new shop standard that the greedy optimisation algorithm assumes. This results in an average saving of 2 days for the VFSG, and 4 days for the IDG, as shown in Table 6.3. Using Equations 6.2 and 6.3, this results in savings of 2 IDGs, and 1 VFSG. Given the list price of both components (USD 315,853 and USD 502,380), this results in a total expected saving of USD 1,134,086.

Table 6.3: Savings in Required Stock per Component Type in TAT and USD

Component Type	Current Shop TAT	Expected Shop TAT	Savings [days]	Savings [Units]
IDG	18	14	4	2
VFSG	16	14	2	1

6.3. Conclusions Phase I

This section aims to summarise and discuss the results stated in this chapter. Section 6.3.1 covers the results from the uncertainty scenarios and effect of certain parameters on the relationship between available capacity and shop performance. Section 6.3.2 contains the conclusion and discussion of results with regards to the supply chain.

6.3.1. Conclusions: Uncertainty Scenarios

Based on the results shown in the previous section, several conclusions can be drawn. First of all, a clear pattern can be identified between the available capacity and shop performance, regardless of the considered time interval. The curve starts with a small angle of inclination. This is caused by the high degree of accumulation that occurs if the capacity is far below the required capacity; increasing capacity by one makes little to no difference. The level of accumulation is also dependent on the size of the considered time interval, and/or the amount of incoming components within that interval. Once the capacity reaches a certain point, the net effect of increasing capacity by one results in significant increases in performance. When performance is close to the 95 mark the curve stagnates, meaning that a large amount of additional capacity is required to reach a small increase in performance. The level of stagnation is related to the amount and size of high inflow peaks, which disrupt the balance in the process. While the above-described pattern is present for the First-In-First-Out principle, a different pattern is observed for other priority scenarios. In case of 'First from Buffer', the curve behaves almost opposite when considering low capacity. This is explained by the fact that there is still a large amount of accumulation of components in the overflow buffer, but less components enter the overflow buffer since they are immediately taken from the regular buffer. Even though it yields a higher performance, it also results in components staying in the overflow buffer for a significant amount of time; in case of Jan-Mar 2017 inflow up to 70 days). This immediately highlights an issue in the current KPI measurement, in which focus is on shop performance, rather than on average days overdue. In many situations it might be beneficial to have a stable procedure in which every component is 1 day overdue compared to a highly erratic procedure resulting in a higher shop performance, but a very high average in days overdue.

Another conclusion to be drawn is that the impact of the percentage of overdue due to disrupted on the shop performance, is large. If no progress is made in the on-time performance of components in the disrupted flow, the desired SL of > 95% will not be met. When looking at the required capacity per month for both 2016 and 2017, it can be seen that there is a large variation in required capacity per month. As shown from the correlation plots, the required capacity is mostly dependent on the average inflow per day. Besides that, no pattern or seasonality can be seen in the required capacity when comparing 2016 and 2017, which was also concluded based on inflow patterns and seasonality analysis in Chapter 5. The effect of highly variable inflow is significant: when considering constant inflow of 4 components per day, the required capacity is equal to 15 instead of the 18 in case of actual inflow (considered monthly). Regarding the effect of growth on the required capacity, it can be concluded that there exists a fairly linear relationship between the two. Finally, in case of working weekends (corresponding to an increase of 4 days in contracted TAT), the performance is increased by approximately 10%. Similar effects are expected if process time is decreased, which can be obtained by removing waste from the process and/or eliminating waiting/buffer time.

6.3.2. Conclusions: Required Capacity and Effect on Supply Chain

Regarding the required capacity based on historical inflow data from 2016 up to and including 2017, the minimal required capacity is 19 net fte per day, yielding a SL > 95%. This only assumes regular flow for a period of more than 3 months, as it has been established that a large part of the disrupted flow is overdue regardless of the available capacity in the shop. By using the assumed shop TAT of 14 calender days and comparing this to the actual shop TAT in shop EWF over 2017 per component type, it was shown that for the VFSG the shop TAT is reduced by 2 days, and for the IDG by 4 days. Using a linear approximation of the user function based on KLM data, it was found that this results in savings of 1 VFSG and 2 IDGs in the supply chain, summing up to a total saving of approximately USD 1.2 million. While this is quite significant, it should be noted that there are also additional cost for adding capacity, a total of 6 additional fte's resulting in a yearly cost of approximately USD 360,000. This case study only focuses on the top 2 components. In reality, shop EWF maintains a larger pool of components, for which similar computations can be made, expected to yield similar results and thus additional profit. Besides that, the scope is limited to focus only on required stock, but another important source for additional cost is the leasing of components which occurs in case of no available serviceable components in the pool. Due to complexity in linking shop TAT to lease-in, this parameter was not taken into account, however it is expected that with a higher and mostly stable shop TAT, the number of lease-ins will decrease. Further research should be done to analyse possible savings. Finally, it is expected that KLM's Component Services Division will grow significantly in the coming years. For that reason it is important to strategically upscale the available capacity in the shop to be able to handle the increase in incoming components. This also affects the required stock levels and thus investment to be made in the next years.

7
Methodology Phase II: Decision Support System

In Phase II, the operational parameters such as capacity are treated as a given, and the inflow (or demand) is to be controlled. The aim of this chapter is to discuss the methodology used to apply a MCDM method to make a decision on which components can enter the shop with the aim of ob- and maintaining a sufficient service level. The objective of the decision support model is to be able to make quick decisions based on data rather than human experience. In the context of the case at KLM E&M, this translates to choosing between two alternatives: repair a component in-house, or outsource it to another repair station. This decision is made based on several criteria, on which both alternatives are graded. The manner in which alternatives are graded and/or chosen, depend on the type of multi-criteria decision method used, as described in Chapter 2. The chapter starts with the argumentation behind the chosen MCDM in Section 7.1. This is followed by the approach taken, in which the determination of the three criteria is explained in Section 7.2. Section 7.3 describes the approach taken in terms of scaling and grading of each of the alternatives. This is followed by the assumptions and data gathering and reliability in Sections 7.4 and 7.5. The solution techniques and approach are discussed in Sections 7.6 and 7.7. Finally, the strengths, weaknesses, and limitations of the methodology are covered in Section 7.8.

7.1. Choosing a Multi-Criteria Decision-Making Method

In order to choose between MCDM models, several factors are to be taken into account: the desired objective of the model, problem characteristics, and user requirements. The objective is discussed previously and can be summarised as follows: *"Choosing the best of two alternatives for each incoming component, taking into account multiple criteria in order to reduce TAT and obtain a high (>95%) and stable shop service level"*. When looking at the problem at hand, it has the following characteristics:

- Multi-objective: maximise service level, while minimising cost
- Finite and limited number of alternatives and criteria
- Includes uncertainty

There are several user requirements: transparency, simplicity, and objectivity. The interpretation of these requirements differs depending on perspective. Academically, the desire is to be able to generalise results and use the discussed methodology and implementation for other applications with similar characteristics. In general, higher complexity requires more assumptions and thus makes the technique more specific for a certain application. For that reason, in this case simplicity is desired to avoid limitations in use. If the foundation is relatively simple, the model can be tuned for specific applications by expanding the base model. Also, if a model is built specifically for one application, accuracy and possibilities for use in real-life applications is reduced due to the large amount of required assumptions. When looking at the practical application, transparency is critical. Two other requirements are ease-of-use and simplicity. Since the decision support tool will be used by several employees in different layers of the organisation, it is of importance that it can easily be understood why a certain alternative is to be chosen. This not only gains trust of the employees, it also provides insights in possible fundamental issues and creates opportunities for continuous improvement of both the model as well as the process. Finally, it is important that the model has a high degree of objectivity. In the current situation the decision to outsource is based solely on experience or opinion by one, or a handful of employees. While these people posses a large amount of information on the dedicated process, they are often unable to look at the overall picture. Therefore, it is important to reduce the subjectivity in grading the alternatives. Taking into account the above-mentioned factors, the choice is made to prioritise simplicity and transparency and choose the Weighted Sum Method.

7.2. Approach

As stated in earlier in this chapter, there are two alternatives to choose from: repair a component in-house, or outsource it to an outside repair station. The two main KPIs are used as criteria, in which service level is divided into in-shop service level, and expected days overdue. This results in the following three criteria, on which both alternatives will be graded: 1) direct cost, 2) effect on in-shop service level, and 3) expected number of days overdue. Figure 7.1 shows a flowchart for the approach taken to obtain a solution.

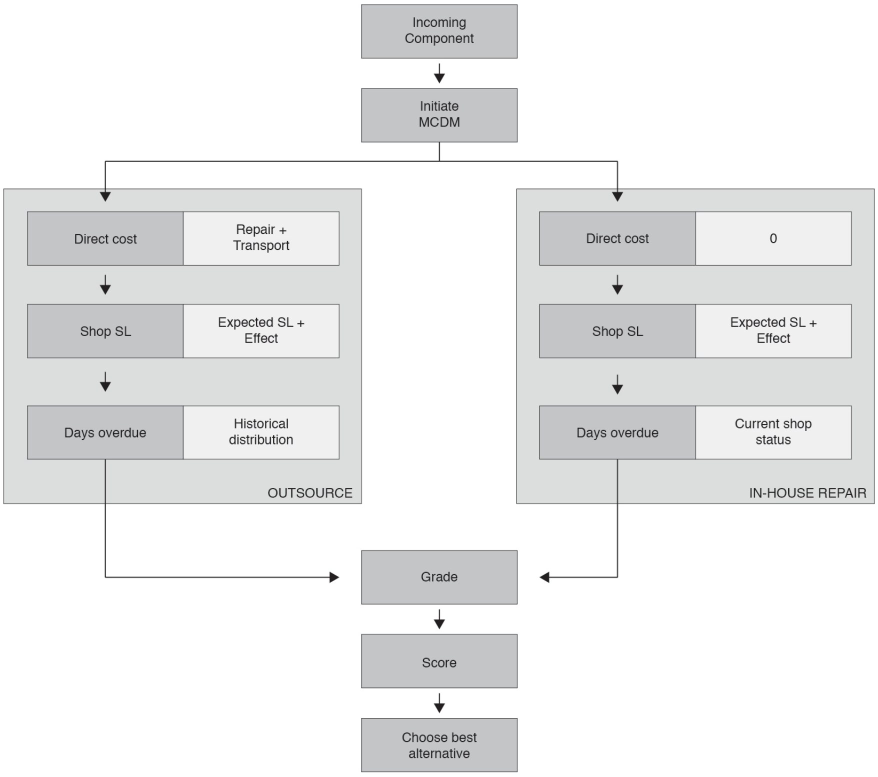

Figure 7.1: Flowchart Weighted Sum Method

7.2.1. Determination of Direct Cost

The direct cost consist of the repair cost of a component and a standard transport fee. The transport fee consists of a standard export fee, a transportation fee, and a standard outstation fee including customs and administrative tasks. In order to compute the repair cost per component and repair type, historical data is required to determine the cost of repairs, as well as the distribution of incoming components and types of repair. This is followed by randomising the occurrences of component and repair type.

7.2.2. Determination of the Effect on the In-Shop Service Level

The second criterion that each alternative is graded on is the effect that it has on the in-shop performance. This criterion consists of two sub-criteria: 1) the current service level, and 2) the effect of the incoming component on the service level. If the current service level is (far) below the required service level, it should increase the likelihood to decide to outsource a component, regardless of the effect that one additional component has on the SL. In general, this is positive in case of outsourcing, as it creates space in the shop process and allows for flexibility and a small buffer for components in the shop. If a component enters the shop it will be more likely to have a negative effect on the in-shop service level, as it creates additional work in the shop, and thus has a chance of increasing accumulation of existing work in the shop, increasing chances of an overdue component. In order to identify the effect that each incoming component has on the in-shop service level, both the current shop performance as well as the expected performance for that specific component are to be known. The current service level is fairly easily obtained by using greedy algorithm as described in Chapter 3. To obtain the expected service level, three steps are to be taken, of which the implementation is found in Chapter 8:

1. Determine the expected number of days until completion of the component
2. Determine whether or not this component is expected to be completed in time, and by how many days
3. Compute the effect on the service level of a certain component

7.2.3. Determination of Expected Days Overdue

The third and final criterion is that of the expected number of days overdue. This is important, as the in-shop service level itself does not provide any information on how long a component remains overdue. As shown in the results of Phase I of the research project, having several components overdue for over a large amount of days can the same impact on the supply chain as having many components overdue for 1 day. While the in-shop performance does not have a direct impact on the supply chain unless accumulated over time and/or volume. The expected days overdue on the other hand has a direct link to the supply chain, and is thus more directly related to other KPIs: customer service level and cost in the supply chain. For the in-house repair alternative, the expected days overdue can be taken from the operational model described in Section 7.2.2. When looking at the outsource alternative, the distribution of expected number of days overdue is taken from historical data and randomised to link the expected number of days overdue to the specific component.

7.3. Set-up of Scales & Grading of Alternatives

As discussed in Chapter 2, a major limitation of the Weighted Sum Method is its inability to handle data expressed in different units. Therefore, the grading of alternatives for each criteria will be done using normalised scales. Regarding cost, there is a pre-defined set of possibilities, dependent on the type of component and the corresponding repair cost. This set is normalised between 0 and 1, where 0 indicates no additional cost, and 1 indicates the highest possible amount of additional cost for a component-repair mix. When looking at the service level and effect on service level, there is no pre-defined set of possibilities, rather it can be any percentage between 0 and 100. For the current performance the 1 corresponds to 80%, and 0 corresponds to 100% (the lower total score the better). Anything below 80% has a grade higher than 1 for in-house repair (vice versa for outsourcing), since the likelihood of outsourcing should increase drastically in case the shop performance drops below 80%. For the effect on the shop service level, the scale is from -1% to 1%, where the negative effect is in case of in-house repair, and positive in case of outsource. Another aspect to be taken into account is the fact that the larger the chosen time interval (or inflow scenario), the smaller the effect of one component on the SL. For that reason, only the last x components are taken into account in the computation of the (expected) performance and effect on SL. Finally, for the expected days overdue, the lower limit is 0 and the upper limit is 30 days.

Besides the scaling of criteria and the subsequent grading of alternatives, the third parameter required to determine the total score are the weights corresponding to each of the criteria. These weights are often based on experience, current situation, or the opinion of the stakeholders. To limit the subjectivity in determining the weights, an approach is chosen in which a large variation of possible weights are considered. For each of the combination of weights, the best alternative is chosen. This is done using Eq. 7.1 until 7.3. Besides limiting subjectivity, this approach also allows for a much higher degree of flexibility, as it is possible (and highly realistic) that the importance of each of the criteria changes depending on the current state of the operations.

$$w_1 = rand(1, n) \tag{7.1}$$

$$w_2 = (1 - w_1) * rand(1, n) \tag{7.2}$$

$$w_3 = 1 - w_1 - w_2 \tag{7.3}$$

Using the above-mentioned procedure, the final step is to compute the total score for each of the alternatives, using Eq. 7.4 and 7.5. The best alternative for each set of weights is the alternative with the lower score.

$$A_{in} = w_1 * c1_{in} + w_2 * c2_{in} + w_3 * c3_{in} \tag{7.4}$$

$$A_{out} = w_1 * c1_{out} + w_2 * c2_{out} + w_3 * c3_{out} \tag{7.5}$$

7.4. Assumptions

Since the operational model is an adaption of the greedy algorithm from Phase I, the same assumptions apply to the MCDM. However, additional assumptions are made for the remainder of the MCDM model, which are discussed below.

1. Direct cost are zero for in-house repair
2. All used distributions are based on historical data
   (a) Type of incoming component
   (b) Required repair
   (c) Expected days overdue for outsourced components
3. IDG, BUG, and other repair is assumed major if RPT > 4 days
4. Repair and transportation cost for 'other' components is taken as the average of the IDG and BUG
5. All components can be outsourced
6. The three criteria *direct cost, shop service level, and expected days overdue* are the only factors influencing the decision to outsource a component

Assumption 1 states that the repair cost for in-house repair are zero. In reality this is not true, since cost include material, personnel, tooling, etc. to repair the component. However, both personnel as well as tooling are a fixed cost in the shop, as these factors are present regardless of the number of incoming components. Material is not, and therefore this will be the cause of the largest discrepancies for this assumption. The implications on the final result is that the difference in cost for both alternatives is smaller than implied by the result from the DSS. However, it is expected that cost will not drastically change the best alternative, unless disproportionally large. The second assumption might also cause discrepancies with the real situation, as it is not up-to-date. It is however required to initialise the model with information, and using a historical distribution will yield more accurate results compared to simply using an average. There is a risk that in some instances an extreme will be used by the model while in the actual situation the average occurred. This might yield inaccurate results on a given day, but over time the extremes will cancel out. However, in case the distributions change over-time, which is to be expected given the dynamic nature of aircraft maintenance, they need to be updated in the model in order to cause possibly large deviations from the actual situation. Assumption 3 is used to make a distinction between minor and major repair. Whereas for the VFSG this data is readily available, for the remaining components this is not the case. This assumption is discussed with, and set-up using the help of the shop leaders of EWF and can be taken as a rule-of-thumb. Regarding the repair and transportation cost for components outside of the top 3, they are assumed to be the average of the IDG and BUG due to lack of available data. The impact can be quite significant if certain components have extremely high or low repair cost and it is therefore advised to further research this to obtain more accurate results. However, given the low percentage of 'other' components, the total effect is expected to be minor. Assumption 5 states that all components can be outsourced. In reality, the majority of incoming components has a Time & Material contract, meaning they are not part of the pool of components owned and maintained by KLM E&M. In many cases these customers pay for the quality and reliability of maintenance performed by KLM, which cannot be guaranteed in case of outsourcing. For that reason, many T&M components will not be outsourced. This introduces a higher degree of complexity in terms of preferred components to outsource. The net effect is expected to be small, however it is possible that on a given day a large inflow peak is present consisting only of T&M components, which cannot be outsourced. This limits the benefit of using a decision tool, as it is not effective in some cases. The final assumption limits the scope and effect of parameters on the decision to outsource a component. In reality, many factors have to be taken into account, such as practical and regulatory considerations.

7.5. Data Gathering & Reliability

Input:

- Repair cost
- Transportation cost
- Inflow scenario as baseline
- Inflow today in number of components per day
- Shop capacity in net fte per day
- Distribution of component-repair combination
- Distribution of days overdue for outsourced components

Output:

- Grade for each alternative per criterion
- Total score, leading to best alternative

The direct cost consist of repair cost and transportation cost for certain component. The repair cost is taken from current contracts with outside vendors, which is up-to-date and is thus reliable. Discrepancies are only introduced for the 'other' components, which are not considered in the Top 3, for which the repair cost is taken as the average of the Top 3. While the percentage of 'other' components is relatively small, this might introduce significant differences in results. For the transport cost a standard fee is identified, which is taken from a standard equation using the weight of the component. The inflow scenario is taken from *SAP*, consisting of the number and type of components entering the shop on a daily basis. As discussed in Section 3.3, the reliability of this data is high. The same is the case for the inflow of today, however ideally one would want to predict the number of incoming components in order to make pro-active decisions to save valuable transportation time. Another input is shop capacity, which is chosen manually, based on the results from Phase I of this research project. Finally, regarding the distribution of component-repair combinations, and the expected days overdue for outsourced components, the historical data from 2016-2017 is used. This data is highly reliable, as it is monitored continuously by the shop. To verify the obtained distribution, it is tested on multiple time intervals, yielding discrepancies no larger than 10%. However, the model might output a different sequence of incoming components compared to the actual inflow, due to randomisation. This can cause discrepancies between model outputs and the actual situation on a daily level, but will level out over time. It is important however, to update this distribution with time, given possible phase-out or increase in flow.

7.6. Solution Technique for Analysing the Effect of Parameters on Best Alternative

To obtain valuable information and insights on the effect of different parameters on the best alternative (in-house repair or outsource), several scenarios are to be run. There are five key parameters that influence the result:

- The component-repair combination, which corresponds to a certain repair- and transportation cost
- The current status of the shop, yielding the expected days on-/over-time, and information on the current and expected shop performance
- The available capacity in the shop
- The time interval taken as the initial conditions, which corresponds to a certain inflow scenario
- The number of components entering the shop on day t

The first step is to create a baseline scenario. For this case the initial time interval is taken as the inflow between January up to and including March 2017. This inflow scenario is also used in Phase I of the research, and is presumed to be representative of the recent inflow distribution. A period of three months is used, as this has proven to be sufficient in order to determine the required capacity, and to limit the computational time. For the baseline scenario, the component-repair combination is an IDG with a minor repair, as the probability of combination is largest compared to others. Besides that, the repair cost of this combination is near average. For simplicity, the 1st incoming component on day t is considered in the baseline scenario. Regarding the available capacity in the shop, the baseline scenario considers the required capacity computed in Phase I of this research project, which is equal to 19. However, it is expected that with a capacity of 19, there should be sufficient space in the process to always repair a component in-house. For that reason a second baseline capacity of 15 is used, the main objective being to show the effect of varying parameters in choosing one of two alternatives. Table 7.1 provides a summary of the parameters used as the baseline scenario.

Table 7.1: Baseline Scenario 1 and 2 Decision Support System

Historical Inflow Scenario	***n*-th Component**	**Component-Repair Mix**	**Capacity**
Jan-Mar 2017	1	IDG minor repair	19
Jan-Mar 2017	1	IDG minor repair	15

Using the baseline, several scenarios are developed to test the impact of individual parameters on the final score - and thus the choice to in- or outsource a component, of which the results are presented in Chapter 9.

1. Varying time interval
   (a) 1 month - capacity 15, 19
   (b) 3 months - capacity 15, 19
   (c) 6 months - capacity 15, 19
   (d) 12 months - capacity 15, 19
2. Varying component-repair combination
   (a) IDG minor - capacity 15, 19
   (b) VFSG major - capacity 15, 19
3. Varying capacity
   (a) Capacity = 10
   (b) Capacity = 15
   (c) Capacity = 19
4. Varying number of incoming components
   (a) From 1 to 30 for capacity between 15 and 19

7.7. Approach for Analysing Result of Using MCDM Method on Historical Inflow Scenario

Similar to Phase I, the above-mentioned scenarios provide insight in the effect of certain parameters and the relationship between different variables and the total score – and thus the best alternative. However, these insights do not provide information on the possible contribution to TAT reduction, in this case focused on waiting-, transport-, and buffer-time. For that reason, the decision support tool will be used over a limited time interval. For this time interval, three things are monitored: 1) shop service level, 2) average shop TAT, and 3) outsourcing cost. These parameters are then compared to results from the actual situation. Direct repair and transportation cost are fairly straightforward to compare, and will most likely be higher when using the DSS, since more components will be outsourced. However, it is expected that the shop SL will be higher, and the average shop TAT per component-type will be lower. The shop TAT of the actual situation will be compared to the simulated scenario using the DSS. With the user functions for the top components determined in Phase I, the shop TAT can be linked in changes in required number of units in stock, which corresponds to a certain reduction in financial risk. Summarising, the following steps are taken to analyse the result of the use of an MCDM method on actual historical inflow scenario:

1. Run model with DSS for specific time interval
2. Monitor key parameters
   (a) Number and type of outsourced components per week
   (b) Shop SL per week
   (c) Average shop TAT per component type (regular + outsource) for the chosen time interval
3. Obtain actual historical information for the chosen time interval on key parameters (see points 2a-2c)
4. Compare results from actual situation to simulated scenario

7.8. Strengths, Weaknesses & Limitations of Approach

The main strength of the model is firstly that it uses the weighted sum method, which is simple and transparent. Based on the inputs, it is easily understood for the user why a certain alternative is preferred. Besides that, the use of varying weights in the output significantly limits the level of subjectivity, and in many cases provides a clear 'best choice', without the need to think about which criterion is most important. Besides limiting the level of subjectivity, it also provides a clear overview of the best alternative subject to uncertainty. Another benefit of the model is that the grading is done objectively based on actual values; a major part obtained from the operational model. The operational model allows for a more accurate and direct link to the current state in the shop. This immediately introduces a weakness of the approach, namely the accuracy and practicality of the operational model. As stated before, the operational model is based on historical data and distributions, and does not project the real-time situation in the shop. Therefore, errors can accumulate and resets are necessary. Another weakness of the approach is that the weighted sum method is very simple and requires normalisation to obtain the same units for all criteria. This might lead to reduction in accuracy and reliability and introduces additional uncertainty. Also, only three criteria are considered, while in reality the decision to outsource a component is dependent on many parameters, often linked by certain relationships. This level of complexity is not taken into account in the weighted sum method. Finally, a limitation of the approach is found in the translation to supply chain, which is only based on one time interval of 3 months. Also, given the assumptions, the focus is on one part of the operation, rather than taking the entire supply chain into account. Similar to the approach in Phase I, another limitation can be found in the link to the supply chain, which is limited to the effect on TAT and stock levels.

8

Implementation Phase II: Decision Support System

Following the methodology of Phase II in Chapter 7, this chapter aims to describe the steps taken for the implementation of the decision support model. Section 8.1 covers the initialisation of the model, followed by discussion of the operational model and weighted sum method in Sections 8.2 and 8.3. Finally, Section 8.4 contains the verification strategy and results for the MCDM.

8.1. Initialisation

During the initialisation of the model, the model structure is set-up and the initial conditions are determined. For this model the initialisation consists of two parts. Part 1 is the initialisation of the operational model, which is equal to that of the greedy algorithm disucssed in Phase I of the research project. This includes the set-up of multiple distributions and the initial conditions with which to start the model. Part 2 includes the initialisation of the MCDM, which contains the set-up of the distribution of the component-repair combination and the determination and normalisation of scales, as discussed in Chapter 7.

8.2. Operational Model

In order to obtain the expected number of days until completion, number of days on/over-time, and the effect on the shop SL, an operational model is to be developed, which is an adaption to the greedy algorithm from Phase I. The adaption is found in the initial conditions; while for Phase I the model was initiated from 0, for the operational model it is desired to start each day using the end-of-day results from the day before. This is achieved by running the greedy algorithm from Phase I as a baseline inflow scenario, and save the results in arrays. The operational model starts with the inflow of today, and runs the greedy algorithm with as initial conditions the previously saved arrays. The expansion is found in the three above-mentioned points: expected number of days until completion, the expected number of days on-time (or overdue), and the effect on the in-shop service level. The pseudo-algorithm of the expansion is shown below.

Here,
t_{compl} = time to completion
$P_{current}$ = current performance
$P_{expected}$ = expected performance
SL_{effect} = effect on SL
$n_{overflow}$ = number of components in overflow
$n_{exp_{overflow}}$ = number of components expected to overflow
n_{total} = number of total components

Algorithm 3 Operational Model

for time = 1:length(time) **do**
  for component = 1:sum(inflow) **do**
    procedure SIMULATION
      if component = completed **then**
        remain completed
      else if component is in repair **then**
        t_{compl} = *RPT - time in repair*
      else
        $t_{compl} = t_{compl}$ *(i - CAP) + RPT*
      if component = completed **then**
        remain completed
      else
        expected days on-time = contracted TAT - (length(inflow) - j) - t_{compl}*(i)*

      procedure OUTPUT
        $P_{current} = (1 - (n_{overflow})/(n_{total})\ *100\%)$
        $P_{expected}\ (j-1) = (1 - (n_{overflow} + n_{exp_{overflow}}\ (1:(j-1))/(n_{total}\ )\ *100\%)$
        $P_{expected} = (1 - (n_{overflow} + n_{exp_{overflow}})/(n_{total}\ )\ *100\%)$
        $SL_{effect} = P_{expected}\ (i) - P_{expected}\ (i-1)$

As stated above, this is an expansion of the greedy algorithm from Phase I, with the main purpose to obtain the required parameters for the MCDM. To compute the time to completion for each component, two options are possible. If the component is in repair, the time to completion is simply the expected RPT - time spent in repair. In all other cases the time to completion is the time to completion of the component that entered the shop exactly *capacity* components + expected RPT. The time to completion can then be used to compute whether or not a component will be completed on time, and by how many days. This is done by using the contracted TAT - time between current component and inflow date - time to completion. Finally, the effect on the in-shop service level is computed by taking the difference in performance excluding and including the component at hand.

8.3. Multi-Criteria Decision-Making: Weighted Sum Method

After the MCDM model is initiated and the operational model is implemented, the next step is to set-up the complete model. For every incoming component (or expected incoming component), the first step is to assign it a component and corresponding repair, based on historical distribution. Once the component-repair combination is known, the direct repair cost is known, both in case of in-house repair (= 0) as well as for the outsource scenario. The next step is to implement the operational model which provides the MCDM model with the current state of the shop. This yields the in-shop service level, effect on the service level, and expected days overdue in case of in-house repair. The expected days overdue in case of outsourcing are based on the historical distribution, as discussed in Section 8.1. Using this information, the next step is to grade both alternatives on each of the criteria, using the normalised scales as discussed in the initialisation. This is followed by the set-up of weights to ensure they do not exceed 1, and the final grading per alternative. The final step is to plot all possible weights and results in a 3D scatterplot.

Algorithm 4 Multi-Critera Decision Model Algorithm

```
for components = 1:max(inflowToday) do
    obtain component type and corresponding repair
    direct cost in-house repair = 0
    direct cost outsource = repair cost + transport cost
    procedure OPERATIONAL MODEL
        implementation operational model
        obtain SL, effect on SL and expected days overdue
        expected overdue outsource = taken from distribution
        procedure DETERMINE SCORE
            directCostIn = 0
            directCostOut = normalise(direct cost)
            slIn = normalise(SL + effect_S L)/2)        slOut = normalise((SL − effect_S L)/2)
                edoIn = normalise(expected days overdue in-house)
                edoOut = normalise(expected days overdue outsource)
                procedure SET-UP WEIGHTS
                    w1 = rand(1,1000)
                    w2 = (1 - w1)*rand(1,1000)
                    w3 = 1 - w1 - w2
                    procedure DETERMINE FINAL SCORE PER ALTERNATIVE
                        for i = 1:length(w1) do
                            aIn(i) = w1(i) * directCostIn + w2(i) * slIn + w3(i) * edoIn
                            aOut(i) = w1(i) * directCostOut + w2(i) * slOut+ w3(i) * edoOut
                            if aIn(i) < aOut(i) then
                                result(i) = 0
                            else
                                result(i) = 1
                        procedure PLOT GRAPH
                            scatter3D (w1, w2, w3, result)
```

8.4. Verification of Multi-Criteria Decision-Making Method

To check if the MCDM model yields the correct results, verification is performed on the following parameters:

- Computation of direct cost
- Implementation of operational model
- Computation of expected days on/over-time
- Scaling
- Computation of weights
- Computation of total score

For the computation of direct cost the verification strategy is to compare the results of the model to the results when computed manually. In order to verify the implementation of the operational model, the results of the effect on the service level is compared to the effect on shop service level of the operational model as described above. The same approach is taken for the computation of the expected days on/over-time, as this is also directly taken from the operational model. Regarding the scaling, the only checks to be performed are the normalisation of the baseline scale and the position of the grade on the scale. Both will be verified using manual computation. As described in Section 7.3, the best alternative will be computed for a large range of different combinations of weights in order to limit the degree of subjectivity in determination of weights. In order to verify the distribution of weights, two main steps are to be taken. Firstly, it should be checked that the weights of the three criteria add up to (and do not exceed) 1, are randomised, and cover many possibilities. Secondly, to ensure accuracy a sufficiently large data-set of samples is to be chosen.

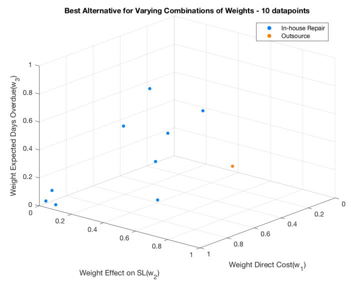

Figure 8.1: Verification Best Alternative Varying Weights: 10 data points

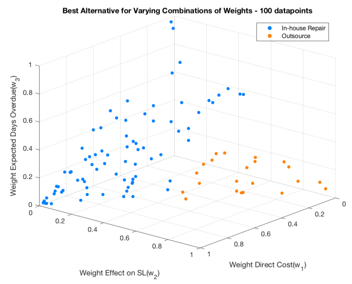

Figure 8.2: Verification Best Alternative Varying Weights: 100 data points

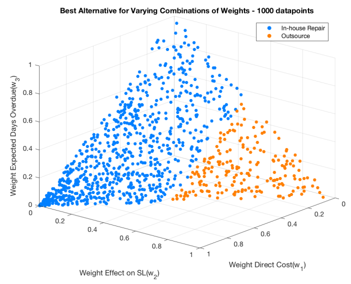

Figure 8.3: Verification Best Alternative Varying Weights: 1000 data points

Figures 8.1 until 8.3 shows the difference between 10, 100 and 1000 data points, from which it can be seen that the latter provides a much better overview and boundary of each of the choices. Figures 8.4 until 8.6 shows the distribution and coverage of the 1000 data points, which proves that there is sufficient coverage of different possibilities and combinations of weights.

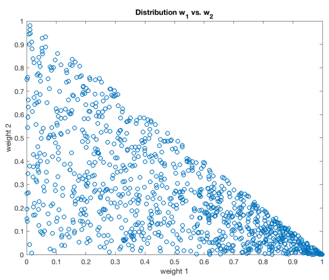

Figure 8.4: Verification Best Alternative Varying Weights: w1 vs. w2

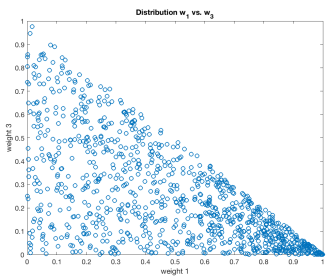

Figure 8.5: Verification Best Alternative Varying Weights: w1 vs. w3

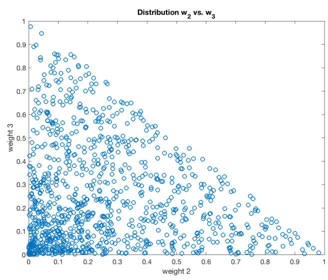

Figure 8.6: Verification Best Alternative Varying Weights: w2 vs. w3

Finally, the computation of the total score must be verified, which is done by checking boundary conditions (e.g. w1 w2 w3 = [1 0 0], [0 1 0], [0 0 1]), and a selected number of other random scenarios. The results from the model will then be compared to the results which are computed manually. Table 8.1 shows the results, indicating the model performs as expected.

Table 8.1: Verification of Total Score Multi-Criteria Decision-Making Model

Scenario	**Weights**	**Test Result**	**Model Result**
[0 0.4 0.7; 0.2 0.3 0.7]	[1 0 0]	In-house	In-house
	[0 1 0]	Outsource	Outsource
	[0 0 1]	In-house	In-house
	[0.33 0.33 0.33]	In-house	In-house

9
Results & Discussion Phase II: Decision Support System

Implementation of the methodology discussed in Chapter 7 on the case study introduced in Chapter 5, yields several results to be analysed. Section 9.1 discusses the effect of varying key parameters on the decision to outsource a component or repair it in-house. Besides that, the effect on the shop and supply chain is analysed by running the model on actual data, which can be found in Section 9.2.

9.1. Effect of Varying Individual Parameters on Decision-Making

As mentioned before, there are four main parameters that influence the choice for either one of two alternatives:

- The component-repair combination, which corresponds to a certain repair- and transportation cost
- The current status of the shop, yielding the expected days on-/over-time, and information on the current and expected shop performance
- The available capacity in the shop
- The time interval taken as baseline, which corresponds to a certain inflow scenario

To identify and analyse the effect of these parameters on the outcome of the multi-criteria decision method, several scenarios are developed, each varying at least one of these parameters. This section discusses the effect of the above-mentioned parameters on the best alternative. Starting with the effect on the time interval in Section 9.1.1, followed by capacity in Section 9.1.2, component-repair mix in Section 9.1.3, and finally the effect of the inflow peak size in Section 9.1.4.

9.1.1. Effect of Time Interval/Inflow Scenario

Starting with the time interval, inflow scenarios of 1, 3, 6, and 12 months are considered, each starting from January 2017. The results can be seen in Figure 9.1 for a capacity of 19 fte, and in Figure 9.2 for a capacity of 15.

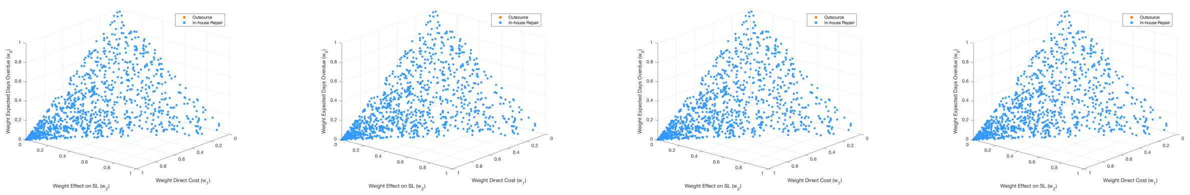

Figure 9.1: Best Alternative for Varying Weights (Capacity = 19) - Time Intervals: 1, 3, 6, and 12 Months

As seen from Figure 9.1, given the initial conditions stated above, the chosen time interval does not have any effect on the decision to repair a component in-house. This can partially be explained by the fact that a fixed number of components are taken into account in order to normalise the computation of performance. Therefore, if the chosen time interval contains more than 300 components, only the last 300 are taken into account in the performance computation. While this decreases the effect of longer time intervals, it still shows the impact of different inflow scenarios and initial conditions. From this it can be concluded that when the available capacity is sufficient, the chosen time interval does not affect the decision to in- or outsource a component.

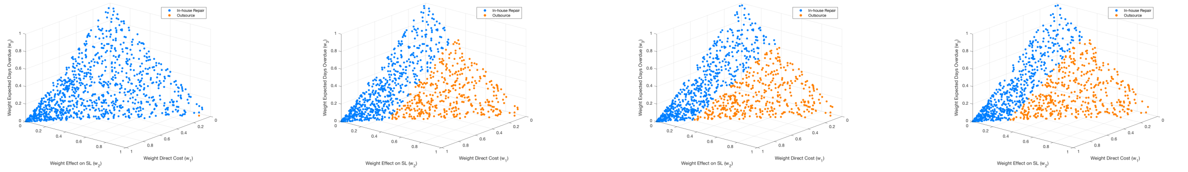

Figure 9.2: Best Alternative for Varying Weights (Capacity = 15) - Time Intervals: 1, 3, 6, and 12 Months

However, when looking at Figure 9.2, the effect of the chosen time interval is evident. The first graph in which only one month of historical data is used advises to repair the component in-house, which can be explained by the fact that the accumulation of components in the buffer and overflow buffer has not yet matured due to short time interval. In the 3-month scenario the situation has matured and it can be seen that it is advised to outsource the component in the majority of weighing combinations. For both the 6 and 12-month scenario there are slight differences to be seen, which can be attributed to the difference in initial conditions. This suggests that if the time interval is longer than 3 months, it does not have a major impact. Again, this is explained by the normalisation factor which means only the 300 most recent components are taken into account for the service level computation. However, the other criteria are not affected by this assumption. There are slight changes to be seen in the distribution of best alternatives, which are most likely to be present due to different initial conditions of the model rather than be affected by the length of the chosen time interval.

9.1.2. Effect of Capacity

When looking at the effect of varying capacity, three options are considered: 10, 15, and 19 available fte per day. As discussed in Phase I, a capacity of 19 is the advised capacity given the historical inflow data. Lowering the capacity results in accumulation of components in the buffer and overflow buffer and thus a large backlog of work, which directly has a large impact on the shop performance. The results can be seen in Figures 9.3 until 9.5.

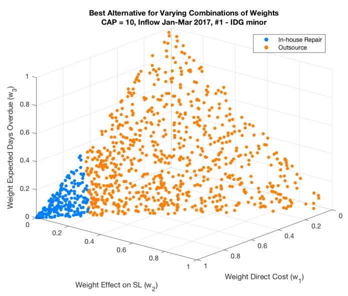

Figure 9.3: Best Alternative for Varying Weights: Capacity = 10

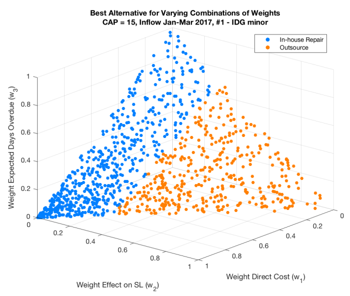

Figure 9.4: Best Alternative for Varying Weights: Capacity = 15

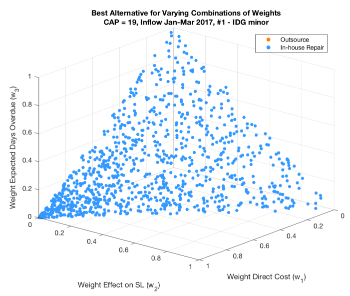

Figure 9.5: Best Alternative for Varying Weights: Capacity = 19

From this it can be concluded that the available capacity has a significant impact on the choice to in- or outsource. When looking at Figure 9.3, outsourcing is preferred unless cost are considered a highly important criterion. This makes sense, given the fact that the cost for in-house repair are zero and thus always lower than the cost in case of outsourcing. For a higher capacity of 15 (Figure 9.4, in most cases it is preferred to repair the component in-house, except when considering shop service level as the most important criterion. This is as expected, since outsourcing a component has a positive effect on the shop service level, while repairing that component in-house has a negative impact on the shop service level. Compared to the scenario with capacity of 10, the top part of the graph (high weight on expected days overdue) has also changed. In the low capacity situation, there is likely a large buffer in the shop, and a large number of components overdue. Adding another component to the buffer means it will be the last in line, often yielding a large amount of expected days overdue. In the medium-capacity situation, the buffer and overflow buffer will be smaller, resulting in less expected days overdue. Given the distribution of expected days overdue for outsourced components, resulting in a high probability of expected days overdue > 20, in most cases the in-house repair performs better on this criterion. Finally, for the high capacity scenario, the advice is to repair the component in-house, regardless of the weights of the considered criteria. In this case again the cost and expected number of days overdue are lower, but also the service level is in favor of the in-house repair. This can be explained by two things, first the current shop performance is above the required SL of 95%. Secondly, the negative effect of a component entering the shop is zero. This means that the required SL is not in danger, yielding the advice to repair a component in-house.

9.1.3. Effect of Component-Repair Combination

The third parameter to be varied is the component-repair combination, which mostly has an effect on the cost. It is expected that, the more expensive the repair, the higher the likelihood for in-house repair. The results for the baseline scenario can be seen in Figure 9.6, in which IDG minor repair is compared to a VFSG major repair. It can be seen that there is no difference between both graphs, which is expected when using a capacity of 19 as there should be sufficient space in the process to repair most incoming components in-house.

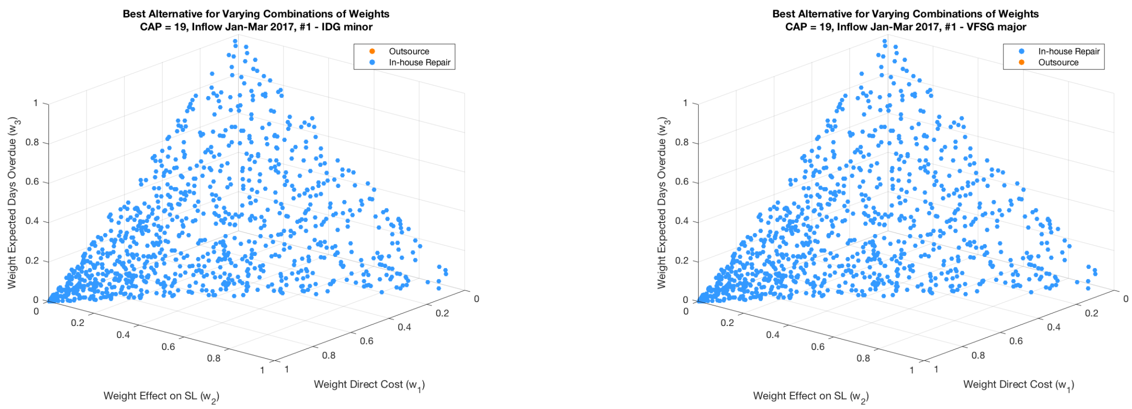

Figure 9.6: Best Alternative for Varying Weights - Effect of Component/Repair: Capacity = 19

Figure 9.7 shows the same graphs, this time for a capacity of 15. A large difference is seen compared to Figure 9.6 regarding the presence of outsourcing alternatives. This is expected given the lower capacity. Also, in this case a clear difference is seen in the distribution between in- and outsourcing between the IDG minor and VFSG major repair. This shows the impact of cost on choosing the best alternative; the higher the cost, the higher the likelihood of repairing a component in-house.

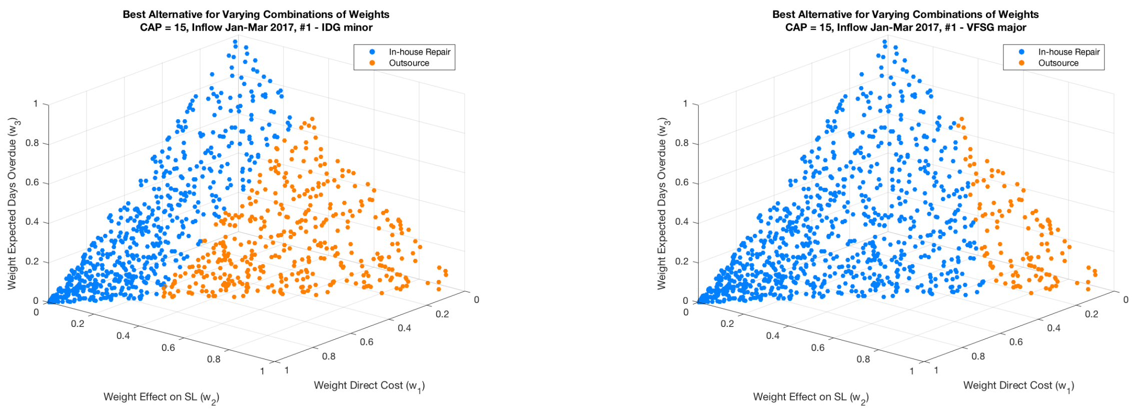

Figure 9.7: Best Alternative for Varying Weights - Effect of Component/Repair: Capacity = 15

9.1.4. Effect of Inflow Peak Size

For the previous scenarios only the first incoming component is considered, but the effect of a large inflow peak has not yet been analysed. For that reason, the model is run for the first until the 30th incoming component, and the best alternative is shown. As a reference point, it is assumed that all criteria are considered of equal importance ($w_1 = w_2 = w_3 = \frac{1}{3}$) in order to provide a better overview. Again, the baseline scenario is considered, with inflow = Jan-Mar 2017, and all incoming components are IDG with minor repair. Figure 9.8 shows the effect of the number of incoming components on one day for capacity between 15 and 19. Here, 0 indicates in-house repair, and 1 indicates outsource.

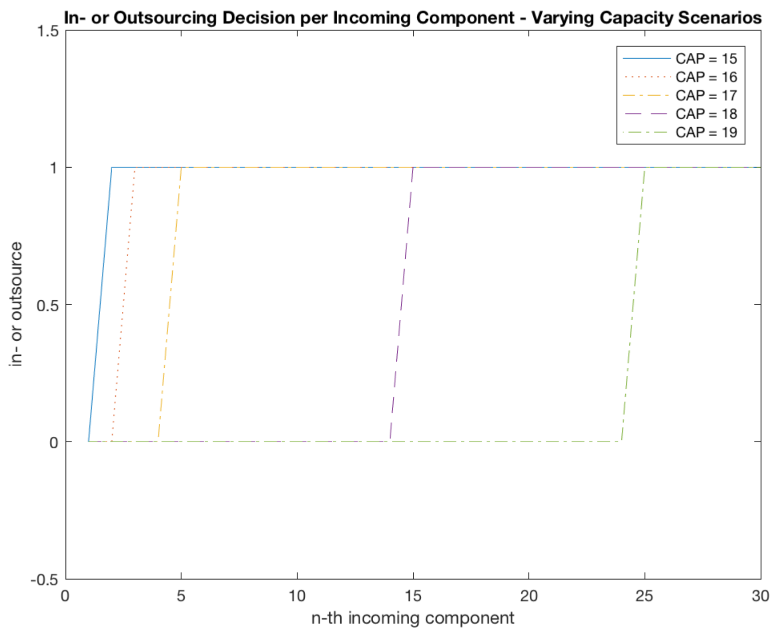

Figure 9.8: Effect of Incoming Component on In- or Outsourcing Decision for Multiple Capacity Scenarios

It can be seen that for low capacity, it is advised to outsource the first and all following components. The reason for this is that the baseline scenario results in high accumulation of components in the buffer, meaning that the shop performance is poor and the negative effect of a component entering the shop is relatively large. For higher capacity situations such as those with 18 and 19 fte per day, the situation is much better, but still it is advised that in case of a large inflow peak some components are to be outsourced in order to maintain a stable performance in the shop and eliminate the risk of building up a backlog.

9.2. Effect of MCDM on Shop and Supply Chain Based on Historical Inflow Scenario

While the results shown and discussed in the previous sections provide insights in the effect of key parameters on the choice to in- or outsource a component, it does not yet create a link to the supply chain and possible benefits regarding TAT reduction focused on waiting- and transport-time. This section aims to do the latter. As discussed in Chapter 7, the followings steps are carried out:

1. Run model with DSS for time interval from Jan-Mar 2017
2. Monitor key parameters
    (a) Number and type of outsourced components per week
    (b) Shop SL per week for Shop EWF
    (c) Average shop TAT per component type (regular + outsource) for the chosen time interval
3. Obtain actual historical information from January up to and including March 2017 on key parameters (see points 2a-2c)
4. Compare results from actual situation to simulated scenario

For the period between January and April 2017, the distribution of outsourced components can be seen in Table 9.1 for both the actual situation as well as the simulation, which is based on historical data.

Table 9.1: Distribution of Outsourced Components in Shop EWF: January-March 2017

Component Type & Repair	Count - Actual	Count - Simulated
IDG minor	2	20
IDG major	0	10
BUG minor	0	3
BUG major	0	3
VFSG minor	7	3
VFSG major	1	1
Other minor	0	6
Other major	0	6

Figure 9.9 shows the cumulative number of outsourced components over the 13 weeks in the chosen time interval. Figure 9.10 shows the comparison between the actual situation and the modelled situation regarding change in in-shop service level.

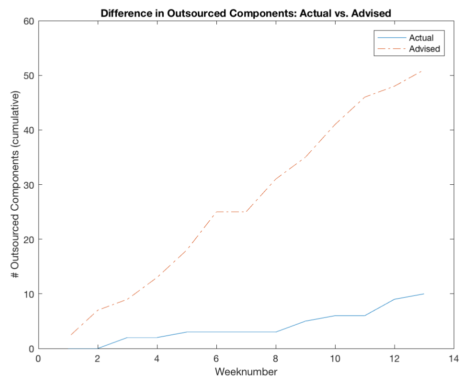

Figure 9.9: Number of Outsourced Components (cumulative) for Actual Situation and Simulation over Time

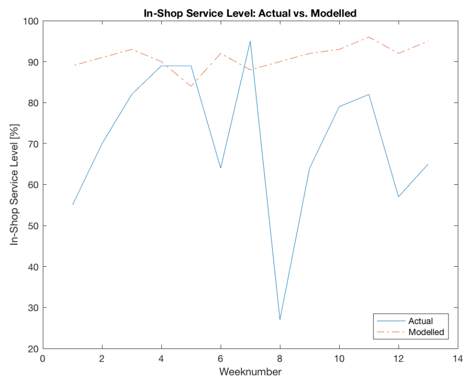

Figure 9.10: In-Shop Service Level over Time for Actual Situation and Simulation

First, when looking at Figure 9.9 it can be seen that the advised number of outsourced components is 5 times as large as the actual number of outsourced components in the same period. Looking at Figure 9.10 it can be understood why. In the actual situation the in-shop service level varies greatly, is highly unstable, and mostly below the target of 95%. On the other hand, in the simulated scenario, the service level is stable and consistently above 85%. It should be noted that this is the in-shop SL only, not taking into account the outsourced components.

Table 9.2 shows the difference in expected repair (& transportation) cost between the actual situation and the modelled scenario. While the number of outsourced components is 5 times higher, the expected cost is only 4 times as high. This can be explained by the fact that in reality many VFSGs were outsourced due to warranty issues. The model does not take this preference into account and assumes the distribution as initiated in Chapter 8.

Table 9.2: Difference in Repair Cost: Actual vs. Modelled Scenario

Total Repair Cost Actual [USD]	Total Repair Cost Modelled [USD]
427,500	1,600,900

Regarding the shop TAT for both scenarios, the results can be seen in Table 9.3. Here it can be seen that in case of regular flow (not taking into account outsourced components), the shop TAT is reduced by 2 days for the modelled scenario. This can be explained by the fact that more components are outsourced, meaning the inflow into the shop is lower, and there is more room in the shop to process the incoming components on time. The average shop TAT including outsourced components is as expected higher, due to the longer TAT in case of outsourcing. However, the difference between the actual (10 outsourced components) and modelled (52 outsourced components) is relatively small.

Table 9.3: Average Shop TAT for Several Scenarios - Actual vs. Modelled

Scenario	Shop TAT
Actual average shop TAT (excl. outsource)	18
Actual average shop TAT (incl. outsource)	20
Modelled average shop TAT (excl. outsource)	16
Modelled average shop TAT (incl. outsource)	23

The effect of the shop TAT on the CSL and financial risks in the supply chain have been analysed and discussed in Chapter 6. In this case however, the use of a MCDM method on average yields a longer shop TAT by 3 days compared to the actual situation. Therefore, no quantitative benefits are to be found based on the chosen scenario.

Two additional scenarios are to be tested in which the outsource TAT cannot exceed the contracted TAT agreed upon between KLM and the outside vendor. This means that if the outside vendor is not able to repair the component within the contracted time, the customer (in this case KLM) is entitled to a replacement, which means a serviceable component is added to the KLM stock. Currently, that is 28 days, which is significantly lower than the actual average outsource TAT of 41 days. Unfortunately, KLM currently does not take sufficient advantage of this agreement, meaning they do not demand a serviceable component in case the TAT exceeds 28 days; resulting in the relatively high current shop TAT. The second scenario is taking into account the new contracted outsourcing TAT of 15 days, which will become the new industry standard. The results can be seen in Table 9.4.

Table 9.4: Average Shop TAT in Case of Varying Outsourcing Contracted TAT

Scenario	TAT in days	IDG	VFSG
Actual average shop TAT - no limit	20	16	26
Actual average shop TAT - 28 day limit	19	16	22
Actual average shop TAT - 15 day limit	18	16	18
Modelled average shop TAT - no limit	23	22	22
Modelled average shop TAT - 28 day limit	19	19	19
Modelled average shop TAT - 15 day limit	16	16	16

From Table 9.4 it can be seen that when KLM consistently reclaims a serviceable component in case the outsourced repair exceeds the contracted TAT of either 28 (current) or 15 (new) days, the average shop TAT is reduced both for the actual situation and the modelled scenario. Similarly to Phase I, the TAT reduction for the top 2 pool components (IDG and VFSG) is linked to savings in required number of units in stock. The last two columns in Table 9.4 show the difference in TAT for the VFSG and IDG for varying contracted TAT agreements.

It can be seen that for the IDG the actual average shop TAT is 16 days and thus within the current contracted TAT. This means that no TAT reduction is achieved, which can be attributed to the fact that only 2 IDGs were outsourced in the actual situation, thus barely impacting the TAT. Therefore, the reclamation of serviceable components in case of outsourcing has limited effect. For the modelled scenario, several IDGs are outsourced with an average TAT of 41 days, thus resulting in the longer average shop TAT of 22 days. This is slightly reduced when considering the contracted TAT of 28 days in case of outsourcing, and even further for a limit of 15 days. For the VFSG the situation is very different: in the actual situation the average shop TAT was 26 days, which can be explained by the large number of outsourced VFSGs in the chosen time period. In the modelled scenario, the number of outsourced VFSGs is smaller than the actual situation, which explains the lower shop TAT. When considering the 28 and 15 day limit, the TAT is reduced significantly.

Savings are computed by comparing the actual average shop TAT to the modelled scenarios with a 28 and 15 day limit. For the IDG this does not yield any savings. However, for the VFSG this results in a TAT reduction of 3 and 2 days, both yielding a saving of 1 VFSG.

9.3. Conclusions Phase II

This section briefly discusses the results presented in this chapter. Section 9.3.1 covers the effect of the key parameters on the decision to in- or outsource a component, which is followed by Section 9.3.2 in which the results of the actual case study are discussed.

9.3.1. Conclusions: Effect of Parameters on Decision to In- or Outsource

Based on the results presented in this chapter, several conclusions can be drawn. Firstly, when considering a capacity of 19, the only parameter that has significant impact on the decision to in- or outsource a component is a dis-proportionally large inflow peak. Otherwise, the situation in the shop is robust enough to handle changes of initial conditions, and any type of component-repair mix. When looking at the inflow scenario it can be concluded that it has no effect if the capacity is sufficient (19). In case of lower capacity (15), the chosen time interval has limited effect if longer than 3 months: only minor effect of inflow scenario on decision to in- or outsource a component. Varying the component-repair combination does not have a significant effect on the decision to in- or outsource a component, especially in case of high capacity. For major repair and high-value components the decision will be in favor of in-house repair, but the difference is minimal, and only present if cost is the most important criterion, which in reality will not occur often. Finally, the change in capacity (taken over the entire time interval) has the largest impact on the decision to in- or outsource a component. This is expected, since it has a significant impact on many parameters taken into account in the MCDM: current- and expected performance, effect on SL, and expected days overdue.

9.3.2. Conclusions: Case Study and Effect on Supply Chain

The case study itself shows that use of the decision support system leads to a 500% increase in outsourced components, with an additional cost of USD 1.2 million. This results in an average increase in shop service level of 30%. This means that the model performs well in increasing service level and smooths out the incoming inflow to a level that the shop is capable of handling with a capacity of 13 fte per day. Also, if KLM starts reclaiming serviceable components from outside vendors if the TAT exceeds 28 days (current situation) or 15 (new situation), the average TAT can be reduced by 2 or 3 days respectively. When looking at the top pool components, this results in savings of 1 VFSG when looking at the actual situation between January and March 2017.

Besides the quantitative benefits obtained when contract agreements are met, there are several qualitative benefits. Firstly, having a high and stable shop performance results in a relatively constant shop output. This is beneficial for the entire supply chain, as it indicates high reliability. High shop reliability suggests that there is little variance in the output of the shop. This means that if a component is notified as unserviceable, it can be fairly accurately predicted when the component will be serviceable again, and thus returned to the customer or back in stock; ready for use. This increases the likelihood of making well-founded decisions regarding possible lease-in or buy-in of additional components, or the necessity to outsource. Besides that, benefits are obtained by being able to quickly make decisions on outsourcing, preferably before the component enters the shop. Unfortunately, that option is not taken into account in these results, as it required previous data on the component, for example the removal notification at the customer. This type of data is difficult to obtain given the current IT infrastructure and of poor quality. However, operating with that data will allow for direct outsourcing. Direct outsourcing suggest the direct shipment from the customer to an outside vendor instead of having to send the component to Amsterdam. In the current situation it is required to follow the entire internal process in Amsterdam until it enters the shop, where it is determined that it needs to be outsourced. Dependent on the origin and destination of the customer and the vendor, the total TAT can be reduced by multiple days in terms of waiting- and transport time. While saving 2 days on a total TAT of 40 might seem insignificant, these 2 days can be the difference in the pool between having a spare part available and having to lease or buy a component.

10 Validation

Regarding validation of the model, the objective is to ensure valid and realistic outputs that represent the actual situation in the shop. Unfortunately, it is very difficult to test scenarios in the shop environment for multiple reasons. First, the performance in the shop is currently highly unstable and far below the required SL of 95%. Besides that, the model is developed using the new situation of CS2.0, in which different assumptions are made that will significantly impact the results. Moreover, creating a test set and scenario requires multiple resources, both practically as well as financially, which given the operational nature of the shop is not a priority. For that reason, two alternative validation methods are used: 1) expert validation (Section 10.1), and 2) sensitivity analysis (Section 10.2). Finally, limitations of the validation strategy are discussed in Section 10.3.

10.1. Expert Validation

For expert validation, several scenarios are developed in which the key parameters of the model are to be validated. Two experts are asked to make a decision for each scenario on whether they would repair the component in-house or outsource it. The experts in this case are the ex shop leader and current shop leader of Shop EWF, as they are responsible for making that decision. The focus is on varying: 1) component type (and thus repair cost), 2) available capacity, 3) current service level, 4) effect on service level, and 5) the expected days overdue. The complete scenarios can be found in Appendix C, and results are shown in Table 10.1.

Table 10.1: Results of Expert Validation

Scenario	Result Expert Validation	Result Model
1	Outsource	Outsource
2	Outsource	Outsource
3	Outsource	Outsource
4	Outsource	Outsource
5	In-house repair	In-house repair
6	In-house repair	In-house repair
7	In-house repair	In-house repair
8	In-house repair	In-house repair
9	In-house repair	In-house repair
10	In-house repair	In-house repair

From the results it can be concluded that the model yields the same decision as the one made by experts. It should be noted that for the modeled scenarios the alternative is chosen that covered the majority of the surface of the 3D plot for varying weights. There were some doubts regarding the decision to make for scenario 2 and 4, given the very high cost in case of outsourcing. However, given the very low shop performance and expected number of days overdue the preferred alternative was still to outsource.

10.2. Sensitivity Analysis

To limit subjectivity in validation, sensitivity analysis is performed as a second means of validation. The objective is two-fold. Firstly, the impact of individual parameters on the final result is to be analysed. Similar to the expert validation, the focus is on the 5 key parameters. This type of sensitivity analysis is shown in Section 9.1, where the effect of individual parameters is analysed on the best alternative. Given the solution set-up discussed in the approach, the resulting graphs are 3D scatterplots, in which 1000 combinations of varying weights are plotted. This means that an increase or decrease of x% in weights corresponds to a different set of weights, meaning it is already incorporated in the plot. Another sensitivity can be found in increasing or decreasing the grading of alternative for each criterion. However, for this problem and approach, the grading itself is objective as it is based on actual data. Based on the results discussed in Section 9.1 a brief summary of the effect of parameters on the choice of best alternative:

- Chosen time interval and length of said interval does not affect the choice of alternative if chosen interval is longer than 3 months
- Available capacity over the chosen time interval has major effect on best alternative due to changing initial conditions which have an impact on SL and expected days overdue
- The effect of component-repair mix (repair cost) is minor and only impacts decision-making in case of limited capacity
- The effect of the inflow peak is highly dependent on the initial conditions and thus the available capacity, but has a significant impact if the peak is larger than 14 incoming components

The second objective of performing a sensitivity analysis is to test the robustness of the results given the uncertainty in decision-making. Here, the focus is on the results obtained in Section 9.2 in which the MCDM model is applied to the actual inflow data from January up to and including March 2017. The results are compared to those of the actual situation in terms of in-shop service level, number and type of outsourced components, cost related to outsourcing, and shop TAT. To test the robustness of the results, the weights of the three criteria are altered. In the baseline scenario, all weights are considered equal ($w_1 = w_2 = w_3 = \frac{1}{3}$). For the sensitivity analysis, three scenarios are developed in which each criterion is given a significantly higher weight than the other two: $[0.7, 0.15, 0.15]$, shown in Table 10.2.

Table 10.2: Validation Scenarios with Varying Weights

Scenario	**Weight 1 (Direct Cost)**	**Weight 2 (Effect on SL)**	**Weight 3 (Expected Days Overdue)**
1	0.7	0.15	0.15
2	0.15	0.7	0.15
3	0.15	0.15	0.7

Figure 10.1 shows the results of the three scenarios compared to the baseline for both the cumulative number of components to outsource as well as the resulting in-shop service level. Looking at scenario 1 in which the weight of direct cost is larger than the two others, it can be seen that it is advised to outsource less components due to the higher cost in case of outsourcing. The difference compared to the baseline in terms of total number of components is 3. The in-shop service level is lower due to the higher number of incoming components, which has an impact of 6% on average. For the second scenario, in which the in-shop service level is considered a top priority, the total number of outsourced components increases by 3. This results in an increase in service level up to 4%. Finally, scenario 3 (high weight on number of days overdue) results in larger degree of in-house repairs, as outsourced components on average yield a longer TAT. This has a similar effect as scenario 1, with a decrease in number of outsourced components and in-shop SL. For all scenarios however, it can be seen that the difference compared to the baseline scenario does not exceed 10%, from which it can be concluded that the model is fairly robust.

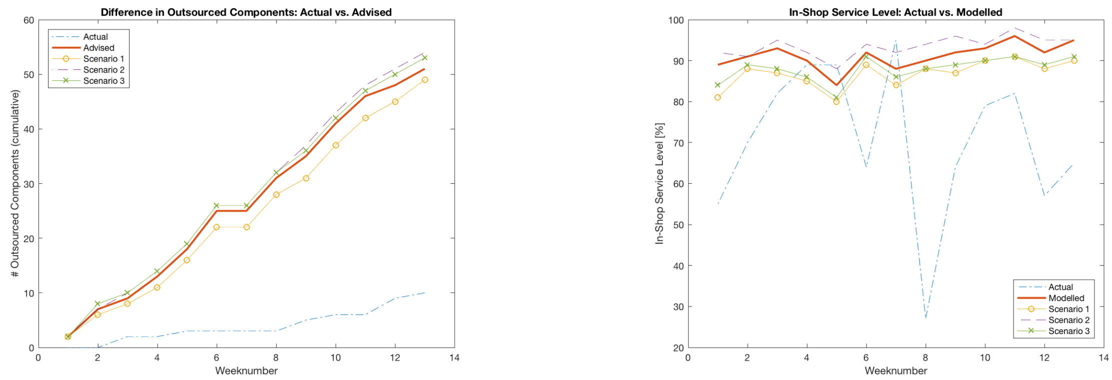

Figure 10.1: Cumulative Number of Outsourced Components and SL for Various Validation Scenarios

Table 10.3 shows the repair cost corresponding to each of the validation scenarios. The difference compared to the baseline scenario is smaller than 5%. Since the component-repair mix is based on the historical distribution randomly distributed over the outsourced components, it is possible that for one scenario there are only minor repairs while for the other a major repair is included. This explains some differences between the repair cost.

Table 10.3: Repair Cost and Difference from Baseline for Validation Scenarios

Scenario	Repair Cost [USD]	Difference [%]
Baseline	1,600,900	0
Scenario 1	1,522,690	4.9
Scenario 2	1,634,580	2.1
Scenario 3	1,604,310	0.2

The final check is with regards to the shop TAT corresponding to each of the scenarios. Table 10.4 shows the average shop TAT including and excluding outsourced components. The shop TAT including outsourced components is equal for all scenarios. While there are some minor changes due to differences in shop SL (lower SL yields more components overdue) the average TAT does not change. However, when looking at the shop TAT including the outsourced components there are some differences which can be explained by the relatively large influence of a component being outsourced. Given the historical distribution it is possible that one additional component being outsourced is the one component that has a TAT of 80 days, which impacts the average TAT.

Table 10.4: Average in-shop TAT Including and Excluding Outsourced Components for Validation Scenarios

TAT	Baseline	Scenario 1	Scenario 2	Scenario 3
Average in-shop TAT (excl.)	16	16	16	16
Average in-shop TAT (incl.)	23	22	24	23

In conclusion, the effect of changing weights and thus relative importance of the three criteria of the MCDM model do not have significant (>10%) effects on the results. This suggests the model is robust and can handle uncertainty well.

10.3. Limitations of Validation Strategy

Even though the validation strategy and results discussed above provide insights in the robustness of the model and its results, there are several limitations. Starting with the expert validation, this type of validation provides some insight in the accuracy of decision-making, but it is still very subjective. Some level of subjectivity is eliminated by systematically building up the scenarios and separately interviewing two shop-leaders, however the choice made by the experts remains dependent on experience and personal preference. Besides that, the sensitivity analysis focuses mainly on the effect of varying weights. There are many other parameters that might influence the decision-making which are not considered in this analysis. Finally, the validation is only performed on a limited time interval in history for a specific application. While this time interval is relatively representative of the last 18 months, it will yield discrepancies to the actual situation.

The ideal validation strategy for this model would include a test in Shop EWF only considering regular flow and components that can be outsourced. The model would then have to be initialised using the begin state of the shop in terms of capacity, buffer size, and service level. Based on the incoming flow the model would adapt the expected service level of the shop and advise on in- or outsourcing decision-making. The aim is then to follow the advice provided by the DSS and observe the effect on service level, outsourced components, and eventually the supply chain. Ideally, this would be performed for at least 3 months, as this is the minimum interval after which the situation has matured. In case the pilot would be performed for at least 3 months, the results can be extrapolated to the supply chain by means of TAT analysis in which the actual TAT is compared to the modelled TAT.

11
Comparison Application and Results Phase I and II, & Possibilities for Generalisation

The previous chapters have covered the methodology, implementation, and results of both approaches. Where Phase I focuses on implementing flexibility into the operation by determining the optimal capacity in the presence of highly variable demand, Phase II aims to control the incoming flow by means of an in- and outsourcing decision-making tool. The objective of this chapter is to summarise key similarities and differences in both the application and limitation (Section 11.1) as well as the obtained results (Section 11.2). Finally, possibilities for generalisation are discussed in Section 11.3.

11.1. Similarities and Differences in Application Phase I & Phase II

Phase I focuses on increasing flexibility in the operational processes by determining the required capacity, assuming and accepting the highly variable demand. Phase II on the other hand, aims to do the opposite: control the incoming flow by means of a decision support system, while keeping the operational parameters fixed. This immediately highlights a major difference between both approaches. For Phase I the inflow scenario is a given en fixed parameter used as input for the greedy algorithm, whereas for Phase II the inflow is an output of the DSS. A similar analysis can be down for the capacity, which in Phase I is the output, but in Phase II is part of the set of input parameters.

Regarding the choice of technique, similar requirements are used for Phase I and II: computational speed, simplicity, possibilities for generalisation, transparency, and flexbility. For both phases this results in the choice of a relatively simple method, which although not academically challenging by itself, provides sufficient scientific novelty in combination with the other research areas and works well in practice. For both phases the reliability of the inputs (and thus the output) is highly dependent on the assumptions and use of historical data. As stated above, several assumptions are made for both approaches that reduce the accuracy of the results. For both models several distributions are used, whether it be RPT, type of incoming component and repair, or overdue from disrupted flow. Given the highly unpredictable nature of aircraft maintenance, there is always a possibility (or risk) that any given time period varies significantly compared to the chosen baseline.

When looking at the implementation of both approaches, the greedy algorithm in Phase I is initiated from 0, while for the adapted greedy algorithm in Phase II the initial conditions are updated at every time step to allow for a near real-time current state of the operational process, which is then used as input for the multi-criteria decision-making model. Regarding the link to the supply chain, Phase I focuses on the reduction in shop TAT obtained by meeting the 95% service level. By means of a user function, this shop TAT (per component type) can be translated to a reduction in required units in stock. For the case study, the focus is on determining the required capacity in the new CS2.0 situation, with a contracted shop TAT of 14 calender days, in order to strategically make decisions on capacity. For Phase II, the objective of the case study is to compare the effect of using the decision support tool to the actual situation, rather than simulating the new CS2.0 scenario. By monitoring key parameters such as number and type of outsourced components, shop SL, and average shop TAT in both situations, the user function from Phase I can be used to determine the effect on the required stock.

Finally, when considering the strengths and weaknesses for both approaches, the first observation is that a key strength for both approaches is the simplicity; both in implementation as well as use considering the case study. Besides that, both approaches can be adapted and expanded to meet the ever-changing operational environment and thus requirements - this will be elaborated on more in Section 11.3. The development of a decision support tool in Phase II introduces more parameters and thus more uncertainty. However, by utilising 1000 combinations of weights, the level of subjectivity is significantly reduced, as in many instances it is expected that one of the two choices will clearly be the better one. The main weakness for both models is that the output is as good (or accurate) as the input, and given the methodology for both approaches in which historical data is used, this can introduce significant risk in case the actual input deviates from the assumed input. For the WSM in Phase II, normalisation of scales and thus grades is required due to the inability of the method to handle differences in units. This might lead to a reduction in accuracy and requires continuous monitoring to check of the results are still valid. A major limitation is that only one area of the supply chain is considered, while in reality there are many inter-dependencies between the different business entities in a supply chain. Also, for both phases only regular flow is taken into account, which results in significant discrepancies from the real-life situation depending on the distribution and effect of disrupted flow. Finally, the focus in this research project is on TAT reduction and the effect on required number of units in stock. In reality, there are many other parameters that affect the total performance of the supply chain, such as lease-in cost and customer service level.

11.2. Similarities and Differences in Case Study Results Phase I & Phase II

In Phase I, the direct relationship between available capacity and shop service level is identified. The pattern between both parameters is dependent on the length of the chosen time interval (the longer, the slower the increase in performance per added unit of capacity), and the average number of incoming components per day (the higher, the slower the increase in performance per added unit of capacity). Another key observation is that the net effect of adding (or subtracting) one unit of capacity differs depending on the current available capacity, but can be significant (>30%).

Based on historical data from 2016 and 2017, the required capacity in Shop EWF is 19 fte per day, as determined in Phase I. Using a capacity of 19 as a baseline for Phase II, this is confirmed by multiple inflow scenarios indicating that any incoming component can be repaired in-house, implying sufficient space in the process. Using the results from Phase II it is concluded that the available capacity in the shop - and thus the current state in the shop in terms of SL and backlog - has the largest effect on the decision to outsource a component, which is consistent to results found in Phase I.

From Phase I is was concluded that the disrupted flow has a large impact on the possibility of meeting the desired SL. This is partially supported by the findings from Phase II, which shows the impact of outsourcing components on the TAT (and thus SL). In the current situation, outsourcing components leads to a much longer TAT due to poor reclamation of components and monitoring of outsourcing TAT. This significantly impacts the shop performance and thus the average shop TAT. It should be noted that outsourcing only takes up a small percentage of the total number of disrupted components, meaning the effect of outsourced components on the TAT is only a relatively small part of the effect of disrupted flow on both the SL as well as the average TAT.

Regarding the link to the supply chain, Phase I results in a 2 day average shop TAT reduction for the VFSG, and 4 for the IDG (considering only regular flow and pool components), resulting in a saving of 1 VFSG and 2 IDGs. This difference is taken between the actual situation and the expected new situation in which the contracted TAT is reduced to 14. For Phase II, as stated above, the case study focused on determining the differences in outsourced components between using and not using the decision support model. Therefore, the expected reduction in contracted TAT was not taken into account in the initial computations. However, if KLM starts reclaiming serviceable components from outside vendors if the TAT exceeds 28 days (current) or 15 (new situation), the average TAT can be reduced by 1 or 4 days respectively. When looking at the top pool components, this results in savings of 1 VFSG compared to the actual situation between January and March 2017. Comparing this to the results from Phase I, the savings from Phase II are slightly less in terms of stock reduction. Besides that, outsourcing 50+ components per quarter (given the current capacity of 19), also yields additional cost of approximately USD 1.2 million.

11.3. Possibilities for Generalisation

While the case study application requires a significant amount of specialisation, the methodologies presented for both approaches allow for some generalisation of the results. The methodology itself is fairly basic; specialisation is introduced by means of 1) assumptions, and 2) inputs. Regarding the assumptions, a key parameter for both phases is the structure of the operational process and procedure that could influence the path-simulation of the components. For Phase II, an important assumptions is the choice of criteria. For example, if the expected days overdue is not a key criterion for a specific application, the entire model structure must be adjusted. When looking at the inputs, the use of historical data and distributions for both approaches allow for easily adjustable models which can be tuned to a specific application by simply using the corresponding inflow data and distributions.

Regarding the link to the supply chain, for both approaches the focus is on the combination of business entity service level and TAT reduction, which is then translated to reduction in required stock. In the KLM case study, a direct relationship between TAT and required units in stock per component type is determined by means of the user function. The benefit of the use of a user function is that it is highly dependent of the specific application, meaning that if the circumstances change, so should the user function. The disadvantage of this is that there is no one-fits-all relationship and it should thus be developed per case.

Another benefit regarding generalisation is that both models (greedy algorithm as well as the decision support system) is that they essentially consist of 'blocks', meaning that the greedy algorithm itself can be used without the link to the supply chain, and the decision support system can be used without the operational model as one of the key inputs. Of course, this lowers the complexity and novelty of the models, but the possibility of separating the 'blocks' allows for broader use.

Overall, both approaches are successful in increasing service level and decreasing TAT in the presence of highly variable inflow. The specific assumptions and operational processes can differ depending on the application, but no fundamental changes are required if the problem has the following characteristics:

- Operational business entity in an aircraft maintenance supply chain - repair shop
- Service level problems and/or structural backlog of work
- Insufficient flexibility in the process
- Time step of 1 day yields sufficient accuracy
- Possibilities for outsourcing
- No more than 3 key criteria on which the decision to in- or outsource is dependent - service level, expected days overdue, and direct cost

12
Conclusion & Recommendations

This chapter contains the conclusions obtained from both Phases of the research project in Section 12.1, and recommendations for further research in Section 12.2.

12.1. Conclusions

Based on the results discussed in the previous chapters there are several conclusions to be drawn for both phases of the research project. Phase I aims at service level optimisation by means of capacity optimisation in the presence of highly variable demand, while Phase II focuses on service level optimisation by controlling incoming demand by use of a multi-criteria decision-making (MCDM) model that assists in in- or outsourcing decision-making. Both approaches aim to make a link to the aircraft maintenance supply chain by translating service level optimisation to TAT reduction, and eventually savings in required stock.

Phase I:

For Phase I, a greedy algorithm is used which simulates the path that each component follows in steps of 1 day for the chosen time interval. The capacity is increased by 1 after each iteration of the time interval until the service level is sufficient. This methodology is applied to the case study at KLM E&M, from which several conclusions can be drawn. The greedy algorithm is preferred over other exact techniques such as linear programming, due to its simplicity, transparency, high computational speed, and flexibility regarding generalisation. The main risk of a greedy algorithm (obtaining a local optimum rather than a global) is avoided by initiating the model from 0 and increasing capacity by 1 until the desired performance is obtained. The lowest capacity at which this occurs is automatically the best.

Firstly, there exist a direct relationship between available capacity in terms of manpower and service level (or performance). This relationship is dependent on several parameters: distribution regular/disrupted flow, productivity, demand scenario, priority procedure, and personnel scheduling. The implementation of disrupted flow assumes a certain percentage of incoming components will be overdue, regardless of the available capacity. For the case study based on historical data from 2017 this results in a maximum performance of 83%, far below the desired 95%. The higher the productivity of technicians, the lower the waste in the repair process, implying lower required capacity. Regarding the demand scenario, the most important parameters are the variance in incoming inflow and the average number of incoming components per day. Increasing (and decreasing) demand by steps of 10%, results in an almost linear relationship between required capacity and incoming demand for the case study at KLM. Perhaps surprisingly, the effect of high peaks on the required capacity is limited, suggesting the process itself contains a certain degree of flexibility. The priority procedure in many cases is a strategical choice, but has shown to impact the performance significantly. While FIFO is the most logical procedure, another possibility is 'first-from-buffer', which provides the illusion of a significantly higher performance, but yields a major increase in average days overdue (in case of insufficient capacity). Finally, working multiple shifts, weekends, and/or nights, results in a different relationship between available capacity and service level, as it effectively yields a lower RPT or longer contracted TAT. In the case study the scenario of working weekends is tested, resulting in an average increase in performance of 10%, regardless of the current capacity.

For the case study at hand, the required capacity is computed based on historical inflow data from 2016-2017, taking account only regular flow, and the new industry standard of 14 calendar days as contracted TAT. This yields a required capacity of 19 net fte per day: an increase of 6 fte from the current 13. The effect of highly variable flow can be seen in the difference in required capacity between the actual situation (19 fte), and a simulated situation in which the variance is 0 (15 fte). Implementing a capacity of 19 in the shop yields on average 4 days TAT reduction for the IDG, and 2 days for the VFSG. By means of a component-specific user function, this is translated to a reduction in required units in stock of 2 for the IDG, and 1 for the VFSG; yielding a total saving of USD 1,134,086 based on current latest list price of the components.

Besides the quantitative savings described above, there is one major qualitative benefit: the shop service level is increased significantly, and perhaps more importantly, stable. This yields higher reliability of the shop, which extrapolates to the entire supply chain. To conclude, the greedy algorithm discussed in Phase I can be used as a tactical model to determine the required capacity based on historical demand in an operational environment, and provide insights and information on the effect of multiple parameters on the service level. By implementing this model, the service level will be increased, resulting in reduced TAT and thus savings in the required stock.

Phase II:

For Phase II, the Weighted Sum Method (WSM) is used as an MCDM to assist in in- or outsourcing decision-making for each incoming component. The WSM is simple, transparent, fast, and allows for generalisation of the model for other applications. The main weakness of the WSM is the inability to incorporate different units, which is eliminated by normalisation of scales. Using the three criteria *(effect on) shop service level*, *direct cost*, and *expected days overdue*, for each incoming component the model decides which option is best: to repair the component in-house, or outsource it to an outside vendor. To limit subjectivity in determining the level of importance of each of the criteria, 1000 combinations of weights are used to visualise the effect on the best alternative.

Several parameters are tested to analyse the possible impact on the best alternative: chosen time interval, available capacity, peak size, and component-repair combination. Using data from the KLM case study, it can be concluded that the available capacity has the largest impact on the decision to in- or outsource. The reason for this is that available capacity over a certain time interval affects the entire shop state: buffer-size, work in stock, shop performance, and shop TAT. For a capacity of 19, the decision support system advises to repair all incoming components in-house, up to a maximum peak of 24 incoming components in one day. This supports the conclusion from Phase I that using a capacity of 19 fte creates sufficient flexibility, and thus robustness, in the process to maintain the desired service level of >95%. For a capacity of 10, or even 15, the advice is to outsource the first (and following) incoming components, in order to restore a stable and sufficient shop service level. The component-repair combination might impact the decision to in- or outsource a component, only in case of insufficient capacity and high repair cost. Maximum peak size to be handled in-house is also highly dependent on the available capacity; for a capacity of 15 or lower, the maximum peak size is 1. On the other hand, for a capacity of 18 the maximum peak size is 14, and for a capacity of 19 the maximum peak size is 24.

When testing the MCDM model on a historical inflow scenario, and comparing the results to the actual situation, the first conclusion is that the model advises to outsource more than five times as many components as the amount actually outsourced: 10 vs. 52. This results in a USD 1,200,000 increase in repair cost, but also an increase in service level of 30% on average: from a 60% average with high variation, to a stable 90%. Making the link to the supply chain by means of TAT and required stock, this translates to an *increase* in TAT of 3 days. This is explained by the significantly longer TAT in case of outsourcing, which is five time higher in the modelled scenario. However, if KLM follows up on contract agreements with outside vendors (maximum TAT of 28 (current) or 15 (new)), the shop TAT is reduced by 2 and 3 days, respectively. This yields a saving of 1 VFSG in required stock for the analysis between January and March 2017, resulting in approximately USD 500,000 savings in required stock.

Similar to Phase I, a major qualitative benefit is the higher and stable service level obtained by implementation of the MCDM. This yields higher reliability and lower variance in the output. Besides that, by quickly making in- or outsourcing decisions (preferably when the component is still at the customer), valuable transportation time can be saved by 'direct outsourcing'. To conclude, the MCDM is an operational model that can be used day-to-day to assist in in- or outsourcing decision-making, resulting in a manageable demand for the shop. This yields a higher service level, which (following contract agreements) results in a TAT reduction and thus savings in required stock.

Combination, Limitations, & Generalisation:

Both approaches are successful in increasing and maintaining a high and stable service level in the presence of highly variable demand. The generally introduced assumptions on demand characteristics are eliminated by using a simulation of the operational process for incoming flow. The result of higher service level is translated to the turnaround time of components, which in both cases is reduced. By translating the TAT reduction to the required stock by means of a user function, savings in terms of cost are obtained.

Comparing the results from Phase I and Phase II, the use of a MCDM (over a 3-month period) yields a larger saving in terms of stock, however also increases repair cost significantly. Phase I yields less direct savings in terms of stock, but is more cost-effective and provides a more strategical decision to upscale capacity to be able to repair components in-house. In reality, a combination of approach I and II will most likely result in the optimal combination of maximising service level and minimising cost, which will be elaborated on further in Section 12.2.

Main limitations of the models include the inability for flexible capacity and priority procedures for the greedy algorithm, and the use of only three criteria for the weighted sum method. When looking at the relationship between TAT and required stock, this is highly dependent of the number of contracted tails, which can therefore significantly impact the results. Regarding the case study, major limitations are the focus on regular flow and the use of historical distributions which are likely to change in the future. Besides that, only the top two components are taken into account for supply chain analysis, and the focus is on reducing the required units in stock. In reality there are other parameters with significant impact on the supply chain, such as Customer Service Level and cost related to additional lease-in of components. Finally, the case study for Phase II only takes into account a limited time interval of 3 months.

While a large degree of specialisation is introduced by introduction and analysis of the case study, the methodology for both phases allow for generalisation. The main reason for this is that both approaches discussed in Phase I and II are straightforward, transparent, and relatively simple. Besides that, they consist of 'building blocks', which can be adapted and/or expanded depending on the application. Specialisation is introduced mainly by use of application-specific inputs and assumptions, which therefore allows for significant flexibility in application. Overall, both approaches are successful in increasing SL and reducing TAT in the presence of highly variable flow, focusing on applications in an operational business entity in an aircraft maintenance supply chain that have problems in terms of service level, flexibility, and/or available capacity, with a process that allows for outsourcing demand, and for which a time step of 1 day is sufficiently accurate.

12.2. Recommendations

While the previous section provides clear conclusions based on the research performed, there are several possibilities for further research, which will be discussed in this section. First, specific recommendations for Phase I and II are covered, after which recommendations for the entire project are discussed.

Phase I:

For the operational process FIFO is assumed, and the performance of the FIFO priority procedure is compared to another common procedure: 'first-from-buffer'. However, there are many in-between scenarios that have not been analysed in this research project, that could affect that relationship between available capacity and performance. Another shortcoming of the greedy algorithm, related to the above-mentioned problem, is that it is not able to take into account component (or customer) preferences. For example, in some cases it might be beneficial to prioritise a Time & Material component instead of a pool component. To increase accuracy of the model this preference should be considered.

Another possibility for further research is the implementation of higher complexity in terms of material and/or personnel scheduling. Regarding material planning, for this model it is assumed that material is always available, which in reality is often not true: availability of material can cause significant problems. For the current approach and case study it is sufficient to assume one shift per day in which all technicians work full-time and possess the required skills to repair all incoming components. For other applications, or to analyse the effect on the service level, the implementation of multiple shifts, contract-types, flexible skills, and process sequences should be researched.

For the case study at hand, the optimal capacity is taken as the minimum required shop capacity to handle historical inflow data. Dependent on the strategy at KLM, the optimal capacity could be either higher or lower than the required capacity. This possibility was not part of the scope of this research project, but can be of interest for the business. Besides that, the capacity is taken as a fixed parameter, while in reality capacity is uncertain due to unexpected illness and holidays. To increase accuracy, it is advised to take this uncertainty into consideration for future research.

Phase II:

One of the main limitations of the approach taken in Phase II is that the inter-dependencies between the different criteria of the multi-criteria decision-making model are not taken into account. In reality, the service level is affected by the expected days overdue, and vice versa. While in most cases it does not lead to a different 'best alternative', it would improve accuracy of the model.

In order to determine the effect of different parameters on the best alternative, for the case study one baseline scenario was used. To increase the accuracy, it would be advised to test the effect on multiple scenarios. Besides that, the effect of capacity uncertainty should be taken into account; similarly to Phase I, the capacity is assumed to be fixed, which decreases accuracy of results.

For the research project at hand the MCDM model is validated using expert validation and a sensitivity analysis performed on a case study. In order to validate the results further it would be advised to use the model in the shop to see the effect on the service level and cost.

General Recommendations:

Firstly, while the effect of disrupted flow is taken into account in the relationship between available capacity and performance, it is not taken into consideration when considering the required capacity. As shown in the case study, the impact of disrupted flow on the possibility of meeting the required service level is significant. For that reason it is important to further analyse the effect of disrupted flow on the service level, and thus on the TAT and stock levels.

When looking at the link to the supply chain, it is limited to TAT reduction, translated to possible savings in required units in stock. It is advised to provide a more complex relationship between TAT and required stock. Besides that, only the top 2 components are considered in this research project. To create a more accurate representation of the possible savings in terms of cost, other parameters should also be taken into account such as lease-in cost of additional components, or indirect cost due to overdue components, and a larger number of components.

This research project contains two approaches to service level optimisation in the presence of highly variable inflow from an aircraft maintenance supply chain perspective. The next step would be to combine both approaches to determine the optimal capacity. For the case study at hand, implementing the required capacity of 19 fte net per day, means the use of a decision support system is not necessary, since there will be sufficient flexibility in the process to handle almost all inflow scenarios - assuming similar demand. Decreasing the capacity to 18 would already yield different results, in some cases a service level drop of 30%. Combining this information with the MCDM to outsource certain components to control the inflow and maintain a stable and sufficient SL could result in an overall better performance, and lower cost. The next step would therefore be to combine both approaches and analyse possible effects on service level, cost, and TAT.

Bibliography

[1] A. Bacchetti, N. Saccani. Spare parts classification and demand forecasting for stock control: Investigating the gap between research and practice. *Omega*, 40:722–737, 2012.

[2] A. Baghalian, S. Pezapour, R.Z. Farahani. Robust supply chain network design with service level against disruptions and demand uncertainties: A real-life case. *European Journal of Operational Research*, 227(1):199–215, 2013.

[3] A. Callegaro, M. Faccio, F. Sgarbossa. Forecasting Methods for Spare Parts Demand. *Universita' Degli Studi di Padova*, pages 1–80, 2010.

[4] A. Charnes, W.W. Cooper. *Management Models and Industrial Applications of Linear Programming*. Wiley, New York, 1961.

[5] A. Gupta, C.D. Maranas. Managing demand uncertainty in supply chain planning. *Computers and Chemical Engineering*, 27(8-9):1219–1227, 2003.

[6] A. Regattieri, M. Gamberi, R. Gamberini, R. Manzini. Managing lumpy demand for aircraft spare parts. *Journal of Air Transport Management*, 11:426–431, 2005.

[7] A.A. Ghobbar, C.H. Friend. Sources of intermittent demand for aircraft spare parts within airline operations. *Journal of Air Transport Management*, 8:221–231, 2002.

[8] A.A. Ghobbar, C.H. Friend. Evaluation of forecasting methods for intermittent parts demand in the field of aviation: a predictive model. *Computers and Operations Research*, 30:2097–2114, 2003.

[9] A.H.C. Eaves, B.G. Kingsman. Forecasting for the ordering and stock-holding of spare parts. *Journal of the Operational Research Society*, 55(4):431–437, 2004.

[10] B. Cardoen, E. Demeulemeester, J. Beliën. Operating room planning and scheduling: A literature review. *European Journal of Operational Research*, 201:921–932, 2010.

[11] B.M. Beamon. Supply chain design and analysis: Models and methods. *International Journal of Production Economics*, 55(3):281–294, 1998.

[12] C. Romero. *Handbook of Critical Issues in Goal Programming*. Pergamon, 1991.

[13] D. von Winterfeldt, W. Edwards. *Decision Analysis and Behavioural Research*. Cambridge University Press, 1986.

[14] D.J. Power. What is a dss? *The On-Line Executive Journal for Data-Intensive Decision Support*, 1(3).

[15] E. Loken. Use of multi-criteria decision analysis methods for energy planning problems. *Renewable and Sustaintable Energy Reviews*, 11(7):1584–1595, 2007.

[16] E. Triantaphyllou, S.H. Mann. An examination of the effectiveness of multi-dimensional decision making: a decision-making paradox. *Decision Support Systems*, 5:303–312, 1989.

[17] E.S. Gardner Jr. Exponential smoothing: The state of the art - Part II. *International Journal of Forecasting*, 22:637–666, 2006.

[18] F. You, I.E. Grossman. Optimal design and operational planning of responsive process supply chains. *Supply Chain Optimisatio: Part I, Wiley-VCH*, I:107–134, 2008.

[19] G. Philips van Buren, A. Gortenmulder. Process excellence klm e&m, 07 2017.

[20] G.M. Marakas. *Decision Support Systems: In the 21st Century, Second Edition*. Prentice Hall, 2003.

[21] Idea Group, Inc. What is a supply chain. Retrieved online, 07 2018. `https://www.igi-global.com/dictionary/the-role-of-information-technology-in-supply-chain-management/28762`.

[22] J. Beliën, H. Forcé. Supply chain management of blood product: A literature review. *European Journal of Operational Research*, 217:1–16, 2012.

[23] J. Van den Bergh, J. Beliën, P. De Bruecker, E. Demeulemeester, L. De Boeck. Personnel scheduling: A literature review. *European Journal of Operational Research*, 226:367–385, 2013.

[24] J.D. Croston. Forecasting and stock control for intermittent demands. *Journal of the Operational Research Society*, 23:289–303, 1972.

[25] J.E. Boylan, A.A. Syntetos. On the bias of intermittent demand estimates. *International Journal of Production Economics*, 71:457–466, 2001.

[26] J.E. Boylan, A.A. Syntetos, J. Croston. On the categorization of demand patterns. *Journal of the Operational Research Society*, 56(5):495–503, 2005.

[27] J.E. Boylan, J.E., A.A. Syntetos. Spare parts management: a review of forecasting research and extensions. *IMA Journal of Management Mathematics*, 21:227–237, 2010.

[28] J.R. San Cristobal Mateo. Multi-criteria analysis in the renewable energy industry. *Green Energy and Technology*.

[29] K. Mela, T. Tiainen, M. Heinisuo. Comparative study of multiple criteria decision making methids for building design. *Advanced Engineering Informatics*, 26:716–726, 2012.

[30] K. Ramaekers, G.K. Janssesns. On the choice of a demand distribution for inventory management models. *European Journal of Industrial Engineering*, 2(4), 2008.

[31] L.A. Zadeh. Optimality and non-scalar-valued performance criteria. *IEEE Transactions on Automatic Control*, 8:59–60, 1963.

[32] L.C. Edie. Traffic delays at toll booths. *Operations Research*, 2:107–1389, 1954.

[33] M. McFadden, D. Worrells. Global Outsourcing of Aircraft Maintenance. *Journal of Aviation Technology and Engineering*, 1(2):63–73, 2012.

[34] M. van Herwijnen. *Spatial Decision Support for Environmental Management*. Vrije Universiteit Amsterdam, 1999.

[35] M. Velasquez, P.T. Hester. An Analysis of Multi-Criteria Decision Making Methods. *International Journal of Operations Research*, 10(2):56–66, 2013.

[36] M.A. Hariga, A. Al-Ahmari, A.A. Mohamed. A joint optimisation model for inventory replenishment, product assortment, shelf space anc display area allocation decisions. *European Journal of Operational Research*, 181:239–251, 2007.

[37] M.K. Zanjani, M. Nourelfath. Integrated spare parts logistics and operations planning for maintenance providers. *International Journal of Production Economics*, 158:44–53, 2014.

[38] N. Subramanian, R. Ramanathan. A review of applications of analytic hierarcy process in operations management. *International Journal of Production Economics*, 138(2):215–241, 2012.

[39] P. Fishburn. Conjoint measurement in utility theory with incomplete product sets. *Journal of Mathematical Psychology*, 4(1):104–119, 1967.

[40] P. Konidari, D. Mavrakis. A multi-criteria evaluation method for climate change mitigation policy instruments. *Energy Policy*, 35(12):6235–6257, 2007.

[41] P. Zhang, S.W. Zhao, B. Tan, L.M. Yu, K.Q. Hua. Applications of decision support systems in aviation maintenance. Retrieved online, 2012.

[42] P.K. Naraharisetti, I.A. Karimi, R. Srinivasan. Chemical supply chain redesign. *Supply Chain Optimisatio: Part I, Wiley-VCH*, I:245–300, 2008.

[43] Q. Xu, N. Wang, H. Shi. A Review of Croston's Method for Intermittent Demand Forecasting. *Conference on Fuzzy Systems and Knowledge Discovery*, 9:1468–1472, 2012.

[44] R. Marti, G. Reinelt. *The Linear Ordering Problem; Exact and Heuristic methods in Combinatorial Optimization.* Springer, 2011.

[45] R. Teunter, B. Sani. On the bias of Croston's forecasting method. *European Journal of Operational Research*, 194(1):457–466, 2009.

[46] R.H. Teunter, A.A. Syntetos, M.Z. Babai. Intermittent demand: Linking forecasting to inventory obsolescence. *European Journal of Operational Research*, 214:606–615, 2011.

[47] R.K. Mobley. *An Introduvtion to Predictive Maintenance.* Butterworth-Heinemann, 2002.

[48] R.L. Ferguson, C.H. Jones. A computer aided decision system. *Management Science*, 15(1):B550–B561, 1969.

[49] R.L. Keeney, H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs.* Wiley, New York, 1976.

[50] R.N.P. Atmojo. Simulation modelling of tablet pcs selection using weighted product algorithm. *Applied Mathematical Sciences*, 8:5705–5719, 2014.

[51] R.S. Gutierrez, A.O. Solis, S. Mukhopadhyay. Lumpy demand forecasting using neural networks. *International Journal of Production Economics*, 111(2).

[52] R.T. Marler, J.S. Arora. The weighted sum method for multi-objective optimization: new insights. *Structural and Multidisciplinary Optimisation*, 41:853–862, 2010.

[53] S. Lui, L.G. Papageorgiou. Multiobjective optimisation of production, distribution and capacity planning of global supply chains in the process industry. *Omega*, 44:149–187, 2014.

[54] S. Opaloglu, I. Ozkarahan. An implicit goal programming model of rht etour scheduling problem considering the employee work preferences. *Annals of Operations Research*, 128:135–158, 2004.

[55] S. Tzafestas, G. Kapsiotis. Coordinated control of manufacturing/supply chains using multi-level techniques. *Computer Integrated Manufacturing Systems*, 7(3):206–212, 1994.

[56] S.M. Lee. *Goal Programming for Decision Analysis.* Auerback, Philadelphia, 1972.

[57] T. Wang. The interactive trade decision-making research: An application of novel hybrid mcdm model. *Economic Modelling*, 29(3):6926–935, 2012.

[58] T.H. Cormen, C.E. Leiserson. *Introduction to Algorithms.* The MIT Press, 2001.

[59] T.L. Saaty. *The Analytic Hierarchy Process.* McGraw-Hill, New York, 1980.

[60] T.L. Saaty. *Decision Making with Dependence and Feedback: Analytic Network Process.* RWS Publications, Pittsburgh, 1996.

[61] T.M. Williams. Stock control with sporadic and slow moving demand. *Jounral of the Operational Research Society*, 35(10):939–1948, 1984.

[62] T.R. Willemain, C.N. Smarta, J.H. Shockor, P.A. DeSautels. Forecasting intermittent demand in manufacturing: a comparative evaluation of Croston's method. *International Journal of Forecasting*, 10:529–538, 1994.

[63] V. Belton, T. Stewart. *Multiple criteria analysis: An integrated approach.* Kluwer Academic Publishers, 2002.

[64] W. Tsai, J. Leu, J. Liu, S. Lin, M. Shaw. A mcdm approach for sourcing strategy mix decision in it projects. *Expert Systems with Applications*, 37(5):3870–3886, 2010.

[65] W. Walker, P. Harremoes, J. Rotmans. Defining uncertainty: A conceptual basis for uncertainty management in model-based decision support. *Integrated Assessment*, 4(1):5–17, 2003.

[66] W. Wang. A stochastic model for joint spare parts inventory and planned maintenance optimisation. *European Journal of Operational Research*, 216:127–139, 2012.

[67] W. Wang, A.A. Syntetos. Spare parts demand: linking forecasting to equipment maintenance. *Transportation Research Part E*, 47:1194–1209, 2011.

[68] W.J. Kennedya, J.W. Pattersonb, L.D. Fredendall. An overview of recent literature on spare parts inventories. *International Journal of Production Economics*, 76:201–215, 2002.

[69] Y. Yang, H. Shieh, J. Leu et al. A novel hybrid mcdm model combined with dematel and anp with applications. *International Journal of Operations Research*, 5(3):160–168, 2008.

A

Appendix: Visualisation Greedy Algorithm

This appendix covers the day-to-day simulation performed by the greedy algorithm in FIFO priority procedure (Figure A.1), and an example of a simulated day in case of 'first-from-buffer' in Figure A.2.

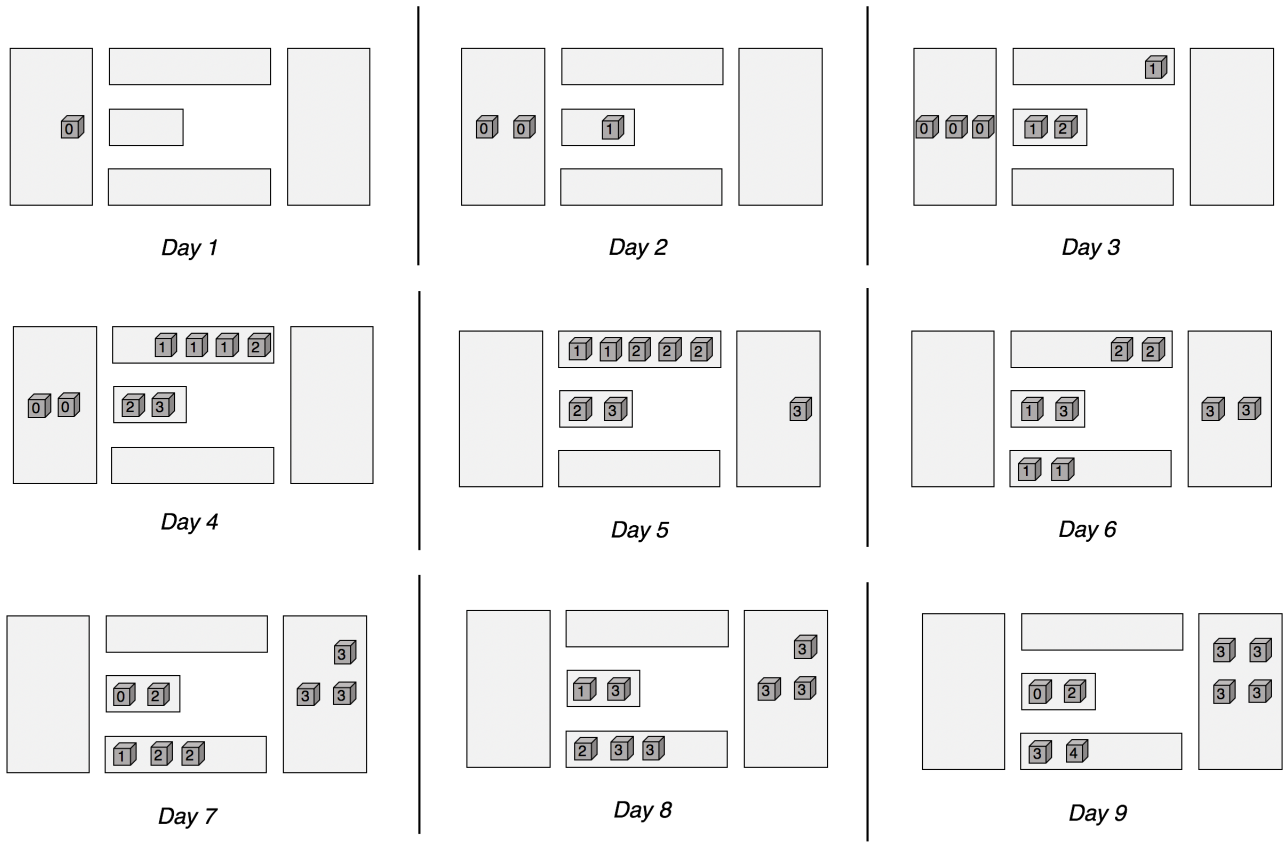

Figure A.1: Visualisation of Greedy Algorithm

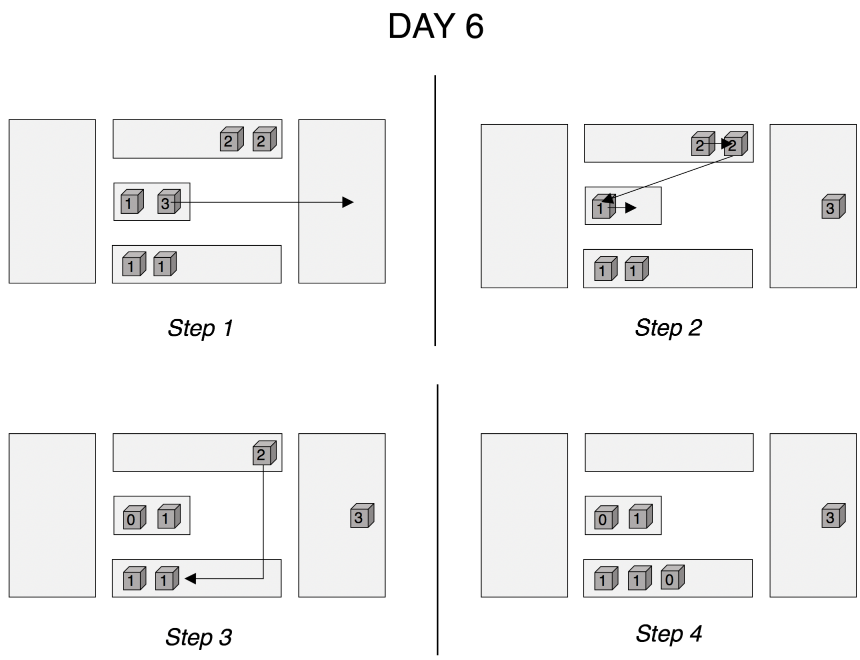

Figure A.2: Visualisation of Key Simulation Actions Greedy Algorithm: Day 6 First-from-buffer

B
Appendix: Demand Forecasting

This chapter provides the methodology and results of the implementation of forecasting techniques for the inflow of Shop EWF for 2017. For the application at hand the focus is on moving average techniques and exponential smoothing techniques. Other techniques include bootstrapping and neural networks, which are eliminated due to the high level of complexity and the need for large amounts of training data. Section B.1 covers a brief theoretical introduction to the techniques used. Using historical data from Shop EWF in 2017, the results are presented and discussed in Section B.2.

B.1. Theoretical Background Demand Forecasting Techniques

Naive forecasting method. This is a primitive forecasting method in which the expected demand is equal to the last observed demand. While this method is incredibly simple, it is highly inaccurate and therefore often only used as a benchmark.

$$\hat{y_{t+1}} = y_t \tag{B.1}$$

Simple Average. In this method the average of all previous values is taken as the new expected demand. This method generally performs better than the naive forecasting method, but is much too simple to be able to take into account complexities such as seasonality and trends.

$$\hat{y_{t+1}} = \frac{1}{n}\sum_{i=1}^{n} y_i \tag{B.2}$$

Moving Average (MA). This method is similar to the simple average method, but instead of taking the average of all previous demand points, it only considers the last n demand points. According to Callegaro et al. [3], this method is only suited for smooth and/or slow moving demand, and has few fields of applicability due to its simplicity.

$$\hat{y_{t+1}} = \frac{1}{n}\sum_{i=0}^{n-1} y_{t-i} \tag{B.3}$$

Weighted Moving Average (WMA). This is a moving average method in which specific values within the last n demand points are given different weights w. In the end, the weights should add up to 1. In general, more recent points are given a higher weight. This technique is again fairly easy to implement, as argued by Ghobbar [8] and Callegaro et al. [3]. However, similar to MA, WMA is only applicable in case of low lumpiness.

$$\hat{y_{t+1}} = \sum_{i=0}^{n-1} w_{t-i} y_{t-i} \tag{B.4}$$

Single Exponential Smoothing (SES). This method can be seen as a weighted average method in which all previous data points are assigned exponentially smaller weights, eventually approaching zero. This method is also known as the exponential weighted moving average (EWMA), and can thus also be categorised as a moving average technique.

$$\hat{y}_t = \alpha y_t + (1-\alpha)\hat{y_{t-1}} \tag{B.5}$$

Here, α is the smoothing factor and is $0 < \alpha < 1$. The higher α, the faster the model 'forgets' older values; meaning higher weights are given to the more recent values. Exponential smoothing was first introduced by Brown in 1956, and expanded by Holt in 1957. Gardner [17] states that SES has been used extensively in a wide variety of fields, including spare part demand, due to its simplicity. However, SES does not perform well in case of intermittent of lumpy demand, as noted by Croston [24]. Also, similar to the moving average techniques, SES introduces a lag in predicting future demand. Another disadvantage of SES is that it does not handle trends or seasonality well. Besides that, it is mostly suitable for short-term forecasting [17].

Double and Triple Exponential Smoothing. As discussed previously, SES does not handle trend and/or seasonality well. Double exponential smoothing applies ES to both level and trend, and triple exponential smoothing applies ES to level, trend, and seasonality. Level is now indicated by l, trend by b, and seasonality by s.

$$l_t = \alpha(y_t - s_{t-L}) + (1-\alpha)(l_{t-1} + b_{t-1}) \tag{B.6}$$

$$b_t = \beta(l_t - l_{t-1}) + (1-\beta)b_{t-1} \tag{B.7}$$

$$s_t = \gamma(y_t - l_t) + (1-\gamma)s_{t-L} \tag{B.8}$$

$$y_{\hat{t+m}} = l_t + mb_t + s_{x-L+1+(m-1)modL} \tag{B.9}$$

Here, α, β, γ indicate smoothing factors determined by fitting, m is any integer, and the index of the seasonal component indicates the offset into the list of seasonal components from the last set from observed data. It should be noted that the above stated equations are additive rather than multiplicative, but both are valid. Both double as well as triple exponential smoothing methods are also known as the Holt-Winters forecasting method. Gardner [17] provides a detailed overview of standard exponential smoothing methods, including damping, multiplicative and additive expansions to incorporate trend and/or seasonality. As to be expected, double and triple ES methods are more accurate and can include trend and seasonal effects. Both are still relatively simple to implement and robust in results, as argued by Regattieri [6]. However, as stated by Bacchetti [1], there is still a problem when dealing with intermittent or lumpy demand.

Croston's method (CR). Since a major problem with previously discussed methods is the low performance in case of intermittent or lumpy demand (something that is very likely in the airline maintenance industry), Croston [24] developed a method based on ES to deal with intermittent demand. Two options are considered, option 1: if $y_t = 0$:

$$p_t = p_{t-1} \tag{B.10}$$

$$z_t = z_{t-1} \tag{B.11}$$

$$q_t = q_{t-1} + 1 \tag{B.12}$$

Option 2: if $y_t \neq 0$:

$$p_t = p_{t-1} + \alpha(q_{t-1} - p_{t-1}) \tag{B.13}$$

$$z_t = z_{t-1} + \alpha(y_t - z_{t-1}) \tag{B.14}$$

$$q_t = 1 \tag{B.15}$$

$$\hat{y}_t = \frac{z_t}{p_t} \tag{B.16}$$

Here, p_t is the demand interval at time t, z_t is the mean demand size at time t, q is the zero demand interval counter, and y_t is the demand at time t. The main advantage of CR compared to other exponential smoothing methods, as stated by Wang et al. [67] is that it takes into account the nature of the relevant demand pattern, meaning demand arrivals as well as demand sizes. Croston's method is currently widely used in industries that deal with intermittent demand, such as aviation, automotive, military and IT sectors [43][62]. Willemain et al. [62] concluded that Croston's method is superior to ES under intermittent conditions. Syntetos and Boylan argued that Croston's method is biased and proposed an adjusted method, proving its superior performance, as discussed further below.

The Syntetos-Boylan Approximation Method (SBA). In 2001, Syntetos and Boylan showed that Croston's method is postively biased by identifying an error in the mathematical derivation of the expected demand

[25], and with a follow-up paper in 2005 proposed an adjusted method: the Syntetos-Boylan Approximation [26]. IN SBA the estimator of mean demand set by Croston is deflated by a factor of $1-\frac{\alpha}{2}$:

$$\hat{y}_t = (1-\frac{\alpha}{2})\frac{z_t}{p_t} \tag{B.17}$$

Several studies (Eaves and Kingman [9], Gutierrez et al. [51], and Syntetos and Boylan [26]) have shown SBA to outperform Croston's method. Due to this, its relative simplicity and proof of concept in the spare part demand research area, SBA is considered a benchmark approach. Other improvements to Croston's method have appeared, e.g. by Teunter et al. [45], who argued that the Syntetos-Boylan Approximation had a negative bias. However, SBA is still the method with the most empirical support. It should be noted however, that SBA is biased for non-intermittent demand.

Syntetos method (SY). The previously discussed bias for intermittent demand when using the SBA method is removed by Syntetos [25], who introduced the following factor to forecast demand. It should be noted that even though the bias on non-intermittent demand is removed, introducing the SY factor increases forecast variance.

$$\hat{y}_t = (1-\frac{\alpha}{2})\frac{z_t}{p_t-\frac{\alpha}{2}} \tag{B.18}$$

Teunter-Syntetos-Babai method (TSB). The problem with many of the above mentioned methods (CR, SBA, SY) is that they do not handle obsolescence well. This means that in case obsolescence occurs, these forecasting methods continue to forecast a fixed nonzero demand. Compared to both Croston as well as SBA, the difference lies in the fact that instead of estimating interval size, TSB estimates the probability of non-zero demand [46]. Besides that, the estimates are updated every period rather than only when demand occurs. Again, two situations are assessed, using d_t to indicate demand occurrence, where $d_t = 1$ in case of nonzero demand, and $d_t = 0$ otherwise. Also, v is the estimate of demand for time t.
For $d_t = 0$:

$$v_t = v_{t-1} - \beta v_{t-1} \tag{B.19}$$

$$z_t = z_{t-1} \tag{B.20}$$

$$\hat{y}_t = v_t z_t \tag{B.21}$$

For $d_t = 1$:

$$v_t = v_{t-1} + \beta(1-v_{t-1}) \tag{B.22}$$

$$z_t = z_{t-1} + \alpha(y_t - z_{t-1}) \tag{B.23}$$

$$\hat{y}_t = v_t z_t \tag{B.24}$$

B.2. Results Demand Forecasting Shop EWF

Implementing the techniques discussed above to the inflow of Shop EWF for 2017, the results for Moving Average and Weighted Moving Average techniques are shown in Figure A.1. For Single Exponential Smoothing, Croston's Method, Syntetos-Boylan Approximation, Syntetos' Method, and Teunter-Syntetos-Babai Method, the results are shown in Figure A.2.

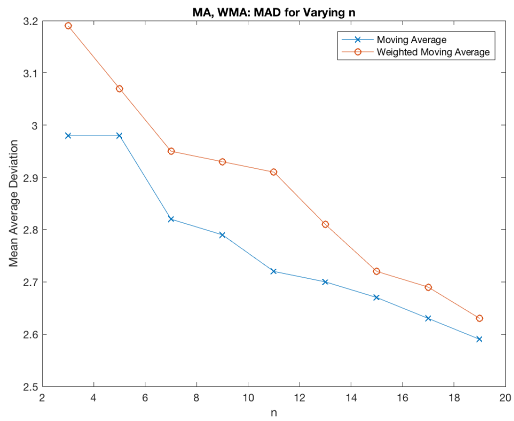

Figure B.1: Mean Average Deviation for Varying n: MA, WMA

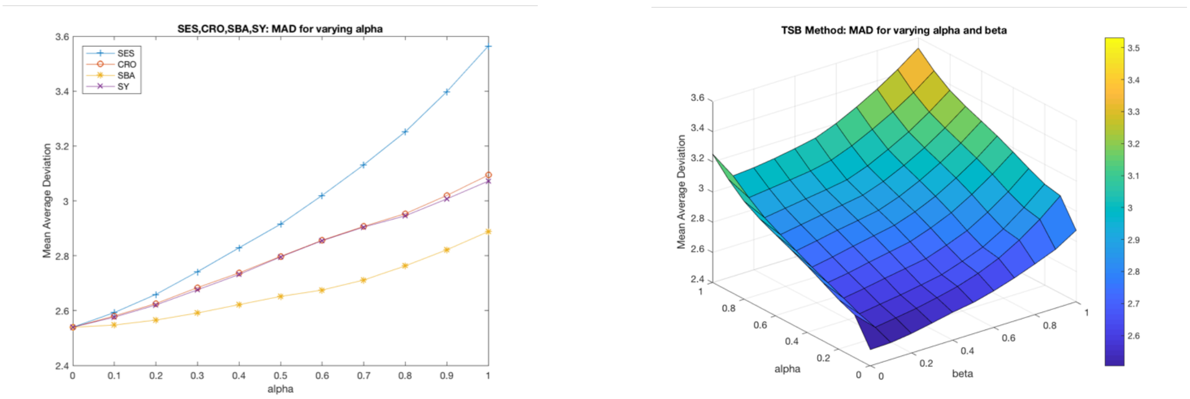

Figure B.2: Mean Average Deviation for Varying Alpha: SES, CRO, SBA, SY, and TSB

From the results shown above, it can be seen that for the moving average techniques, the mean average deviation decreases with increasing n; where n is the number of data points taken into account for the moving average. For the exponential smoothing techniques, the mean average deviation decreases with decreasing $alpha$ (and $beta$). Both indicate that the performance of the forecasting model increases when getting closer to the average; with the optimal performance in case the average is expected.

C

Appendix: Expert Validation Scenarios

Table C.1: Scenario 1

Current shop SL	22%
Incoming component	IDG minor repair
Scaled repair cost if outsourced	0.1358
Effect on shop SL	0.33%
Expected days overdue if outsourced	30
Expected days overdue in-house repair	55

Table C.2: Scenario 2

Current shop SL	22%
Incoming component	VFSG major repair
Scaled repair cost if outsourced	0.9657
Effect on shop SL	0.33%
Expected days overdue if outsourced	30
Expected days overdue in-house repair	55

Table C.3: Scenario 3

Current shop SL	42%
Incoming component	IDG minor repair
Scaled repair cost if outsourced	0.1358
Effect on shop SL	1.33%
Expected days overdue if outsourced	30
Expected days overdue in-house repair	13

Table C.4: Scenario 4

Current shop SL	42%
Incoming component	VFSG major repair
Scaled repair cost if outsourced	0.9657
Effect on shop SL	1.33%
Expected days overdue if outsourced	30
Expected days overdue in-house repair	15

Table C.5: Scenario 5

Current shop SL	70%
Incoming component	IDG minor repair
Scaled repair cost if outsourced	0.1358
Effect on shop SL	0.33%
Expected days overdue if outsourced	30
Expected days overdue in-house repair	1

Table C.6: Scenario 6

Current shop SL	70%
Incoming component	VFSG major repair
Scaled repair cost if outsourced	0.9657
Effect on shop SL	0.33%
Expected days overdue if outsourced	30
Expected days overdue in-house repair	1

Table C.7: Scenario 7

Current shop SL	93%
Incoming component	IDG minor repair
Scaled repair cost if outsourced	0.1358
Effect on shop SL	0%
Expected days overdue if outsourced	30
Expected days overdue in-house repair	1

Table C.8: Scenario 8

Current shop SL	93%
Incoming component	VFSG major repair
Scaled repair cost if outsourced	0.9657
Effect on shop SL	0%
Expected days overdue if outsourced	30
Expected days overdue in-house repair	0

Table C.9: Scenario 9

Current shop SL	98%
Incoming component	IDG minor repair
Scaled repair cost if outsourced	0.1358
Effect on shop SL	0%
Expected days overdue if outsourced	30
Expected days overdue in-house repair	0

Table C.10: Scenario 10

Current shop SL	98%
Incoming component	VFSG major repair
Scaled repair cost if outsourced	0.9657
Effect on shop SL	0%
Expected days overdue if outsourced	30
Expected days overdue in-house repair	0