INTEGRATED CAPACITY PLANNING AND INVENTORY CONTROL FOR REPAIRABLE SPARE PARTS

MAARTEN DRIESSEN

Integrated Capacity Planning and Inventory Control for Repairable Spare Parts

A catalogue record is available from the Eindhoven University of Technology Library.

ISBN: 978-90-386-4557-5

Printed by Proefschriftmaken || www.proefschriftmaken.nl
Cover Design by Wijnand Driessen - Matte! Nande?

This research has been funded by Gordian Logistic Experts BV.

Integrated Capacity Planning and Inventory Control for Repairable Spare Parts

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de
Technische Universiteit Eindhoven, op gezag van de
rector magnificus, prof.dr.ir. F.P.T. Baaijens, voor een
commissie aangewezen door het College voor
Promoties, in het openbaar te verdedigen
op 11 september 2018 om 16.00 uur

door

Maarten Alexander Driessen

geboren te Valkenswaard

Dit proefschrift is goedgekeurd door de promotoren en de samenstelling van de promotiecommissie is als volgt:

voorzitter:	prof.dr. I.E.J. Heynderickx
1e promotor:	prof.dr.ir. G.J.J.A.N. van Houtum
2e promotor:	prof.dr. W.H.M. Zijm (Universiteit Twente)
co-promotor:	dr.ir. W.D. Rustenburg (Gordian Logistic Experts BV)
leden:	prof.dr.ir. I.J.B.F. Adan
	prof.dr. D.P. van Donk (Rijksuniversiteit Groningen)
	prof.dr.ir. R. Boute (Katholieke Universiteit Leuven)
adviseur:	dr.ir. V.C.S. Wiers

Het onderzoek of ontwerp dat in dit proefschrift wordt beschreven is uitgevoerd in overeenstemming met de TU/e Gedragscode Wetenschapsbeoefening.

Acknowledgements

"Winnen doe je met z'n allen."

Johan Cruijff

Winning a football match requires technical skills, stamina, teamwork, a good strategy and encouragement from the supporters. These ingredients are also required for completing a dissertation. Although a lot of people have helped me during this project, for understandable reasons I will address only a number of them.

First, I would like to express my gratitude to my first promotor Geert-Jan van Houtum for his supervision and support. You learned me how to navigate the academic world and helped me structure my thoughts, that were sometimes biased due to the experiences I gained from practice. Besides this, I also enjoyed our "Theo Maassen"-imitations when we were drinking beers at the ISIR-conference in Budapest.

Second, I would like to thank my second promotor Henk Zijm, who joined the team somewhat late in the project. You appeared to be the right person at the right place at the right time to help me get my research project back on track. From the moment you entered the project team, somehow the momentum was there to successfully complete this dissertation.

Even more important is the support that my co-promotor Jan Willem Rustenburg gave me in this project. You played a crucial role when my research project got stuck and you decided to temporarily stop the project to have it assessed by Henk. Next to this, I would like to thank you for helping me organize the project and I would like to thank you and Gordian for giving me the opportunity to do a part-time PhD project.

This research has both a design and empirical character. I am greatly indebted to KLM E&M, The Dutch Railways, ECT, the Royal Netherlands Airforce, Army and

Navy and to GVB for their willingness to serve as 'research labs'. I thank all employees involved in the case studies for their support. You created the inspiring environment I needed to do my research. I thank Fujifilm for their indirect support.

I thank Joachim Arts for our collaboration on Chapter 2 of this dissertation. I liked our many conversations on how to apply scientific knowledge in practice. I would also like to thank Jürgen Donders and Bob Huisman for their invaluable input on Chapter 2, and Jürgen also for his useful comments and insights on other parts of the dissertation.

I thank Joan van Aken, Jan Fransoo, Simme Douwe Flapper, Nynke Faber and Tiedo Tinga for their help to set up and structure the case study research in Chapter 3. The results in Chapter 5 are based on the research Herman van Walree did in his Master's project, and I thank him for doing the data analysis and providing me with the requested results. I would like to thank Vincent Wiers for our collaboration on what finally became Chapter 6, and for being a member of my thesis committee. I thank Wijnand Driessen for creating the beautiful cover of this dissertation.

I thank Dirk Pieter van Donk, Robert Boute and Ivo Adan for being on my doctoral committee and their valuable feedback on my work. Paul and John, I am very happy that you agreed to stand by me as paranymphs!

I thank Marijcke and Claudine for their creativity and ability to solve the challenging planning problems with busy people like Jan Willem, Geert-Jan en Henk.

I would like to thank my parents, family, friends and colleagues for supporting me during this long project. You were really enthusiastic about it at the start, fairly critical in middle (*"ben je nu nog niet klaar?"*) but complimentary at the end (*"knap gedaan, ik was veel eerder al gestopt"*). With the right statements (*"afzien duurt maar even, opgeven is voor altijd"*) you really motivated me to finish this dissertation, after doubting for the hundredth time whether I should stop. Without your support I would not have completed it, and I like to thank you for that.

Finally, Chantal, you really entered my life at the right moment and I am very happy with the love that you gave me during the last couple of months, and all the love that you will give me during the next 150 years!

Maarten Driessen
Valkenswaard, July 2018

Contents

Glossary

"Als ik zou willen dat je het begreep, dan had ik het wel beter uitgelegd."

Johan Cruijff

The following abbreviations are used throughout this dissertation:

Abbreviation	Description
ADI	Advance Demand Information
AM	Additive Manufacturing
EBO	Expected number of backorders
ECT	Europe Container Terminals
FCFS	First-come, first-served
GVB	Gemeentelijk Vervoersbedrijf Amsterdam
KLM E&M	KLM Engineering & Maintenance
KPI	Key Performance Indicator
LRU	Line Replaceable Unit: spare part used at first level maintenance
MLO	Maintenance Logistics Organization: department responsible for supplying all resources required to conduct asset maintenance
MO	Maintenance Organization: department responsible for maintaining the assets
MTBF	Mean Time Between Failures
NRFU	Non-Ready-For-Use: characteristic of a part that indicates that a part is not immediately available for maintenance
NS	Dutch Railways
OEM	Original Equipment Manufacturer: supplier of the capital assets (in this dissertation)
RFU	Ready-For-Use: characteristic of a part that indicates that a part is immediately available for maintenance
RNLA	Royal Netherlands Army
RNLAF	Royal Netherlands Airforce
RNLN	Royal Netherlands Navy
RSPL	Recommended Spare Parts List
SRU	Shop Replaceable Unit: component of an LRU that is replaced at the repair shop

Chapter 1

Introduction

> *"We zijn op zoek gegaan naar de overwinning en dan kom je hem vanzelf tegen."*
>
> *Johan Cruijff*

In many industries, companies depend on the availability of high-value capital assets to provide their services or to manufacture their products. Examples include aviation, oil refineries, public transportation (e.g. trains and subway carriages) and military organizations (e.g. frigates, tanks or combat helicopters). Periods in which the asset is not available for use (i.e. *downtime* of the asset) may among others result in lost revenues, customer dissatisfaction and possible claims or a public safety risk. For example, the (unscheduled) downtime costs of an airplane are estimated to be around $ 36,000 per hour (Knotts, 1999). For optimal functioning and availability of the assets, timely maintenance of the assets is essential. Usually a maintenance concept is developed by the supplier of the capital assets, i.e. the Original Equipment Manufacturer (OEM), who prescribes how to efficiently and effectively conduct maintenance on the asset.

To maintain the high-value capital assets, a service network is set up that includes amongst others maintenance and repair facilities, resources and service tools required to conduct the maintenance and stock locations to store spare parts. In general, we can distinguish two 'extreme' types of service networks for high-value capital assets (Basten and Van Houtum, 2014): *OEM networks* and *user networks*. An OEM network consists of an OEM or system integrator who services (a part of) the assets that were sold earlier, usually at a lower cost than the user can achieve individually. An OEM often has more knowledge about the asset than an individual user and can benefit from the scale of operations (share resources), because an OEM typically services multiple users.

Companies in user networks both use and maintain their high-value capital assets. In such networks, it is often easier to synchronize the operating and maintenance schedule of the assets, because resources required to conduct maintenance are not shared with multiple users. Next to that, the users of the assets often have a high knowledge level about the asset technologies. Therefore, they are able to function as a partner to the OEM in developing technical improvements of the asset (modifications), and they can be a 'smart buyer' in the design phase of new assets.

In this dissertation, we restrict ourselves to user networks. In such a network, a *Maintenance organisation* (MO) is responsible for maintaining the assets. An MO synchronizes the maintenance and usage requirements of the assets, following from the maintenance concept and the operating schedule respectively, and schedules maintenance for all assets. Maintenance on an asset can only be conducted when the required resources (a.o. technicians, service tools and equipment, maintenance manuals and spare parts) are available. An MO is responsible to have all the required resources available in time, except spare parts. The required spare parts are requested at a *Maintenance Logistics Organisation* (MLO).

Obviously, spare parts have to be stocked in order to be able to quickly conduct (unplanned) asset maintenance. Hence, the time that an asset is not available for use clearly depends on the availability of the required spare parts. At the same time, stocking spare parts is costly as total stock investments exceed a few up to tens or even hundreds of millions of euros. An MLO is responsible for all activities that are required to have the right spare parts available in time. Hence, an MLO is concerned with matching supply and demand of spare parts, and its aim is to find the optimal balance between the availability of spare parts and the associated costs.

In the area of spare parts planning and control, much research has already been conducted. Nevertheless, application of this stream of literature in practice is limited and spare parts are often being managed using techniques from classical inventory control theory (such as single-item, single-location models with fill rate targets, sometimes combined with an ABC-classification). This observation was the main motivation for our research approach: by taking practice 'as is' and based on a systematic and detailed description of reality, concepts and models are developed that aim to have both a high practical as well as scientific relevance.

Following this line of reasoning, we start with a broad view on planning and control of MLOs by presenting a *normative framework* in Chapter 2, in which we give an overview of all relevant decision functions including their mutual relationships. This framework is used to cluster available literature, and to identify open research topics by comparing the framework to practice in case studies in Chapter 3. In this case study research we closely collaborated with the Dutch Railways (NS), KLM Engineering & Maintenance (KLM E&M), Royal Netherlands Airforce (RNLAF), Royal Netherlands Army (RNLA), Royal Netherlands Navy (RNLN), Europe Container Terminals (ECT) and Gemeentelijk Vervoerbedrijf Amsterdam (GVB), who all served as 'research labs' throughout the entire research.

From the case studies, we found that *repairable* parts represent the major part of the capital invested in spare parts inventories (70-85%) and that about 60-70% of all repair jobs for spare parts are executed by internal repair shops. Furthermore, in a case study at the RNLN we found that in 80% of all (sub)system downtime occasions due to unavailability of spare parts, it was due to the unavailability of a *repairable* spare part. Hence, in terms of availability and costs of spare parts, it seems worthwhile to focus on parts that are repaired by internal repair shops. Moreover, planning and control of internal repair shops are within the span of control of the MLO, in contrast to parts that are repaired by external repair shops.

Repairable parts inventory systems[1] are systems in which a repair shop returns (failed) repairable parts to ready-for-use (RFU) stock (Hausman and Scudder (1982)). In such systems, decisions regarding repair capacity and inventory interrelate; e.g. the required amount of repairable spare parts depends on the rate at and sequence in which the repairs are conducted. How to integrate or synchronize decisions regarding repair capacity and inventory control is far from clear. In Chapters 4-7 of this dissertation, we develop new concepts and models for integrating decisions in repairable parts inventory systems with *internal* repair shops.

This introductory chapter describes the environment for which we developed concepts and models and further motivates the need for this research. In Section 1.1 we characterize MLOs in more detail, which is the scope of Chapters 2 and 3. In Section 1.2 we give a description of repairable parts inventory systems, the scope of Chapters 4 to 7. Section 1.3 presents our research questions and methodology, and Section 1.4

[1] Actually denoted *repairable inventory systems* by Hausman and Scudder (1982).

concludes with an outline of this dissertation.

1.1 Maintenance Logistics Organisation

Maintenance of the assets generates demand for spare parts, and an MLO is responsible for timely supplying these spare parts. In general, there are multiple sources for spare parts demand and supply and, consequently, there are multiple ways to match the demand and supply of spare parts. In Sections 1.1.1 and 1.1.2, we describe these sources for demand and supply and motivate which sources play a dominant role in this dissertation. In Section 1.1.3 we describe the options to match demand and supply, which option is assumed in this dissertation and how this translates into assumptions regarding performance metrics and the service network.

1.1.1 Spare parts demand

When an asset is maintained, parts are taken out and replaced by RFU parts based on their condition. This principle is called *repair-by-replacement* (Muckstadt, 2005). Spare parts used at the first level maintenance are called *Line Replaceable Units* (LRUs) (Muckstadt, 1973, 2005), and the decision to designate a part as an LRU lies with the MO. The replaced parts cannot be used as *spare* parts, and are denoted *non-ready-for-use* (NRFU). They are either scrapped or returned to a repair shop. In this dissertation, we distinguish two types of parts:

1. Repairable parts: parts that are maintained (repaired or overhauled) after replacement, i.e. parts that are technically and economically maintainable. After maintenance, the part becomes RFU again.

2. Non-repairable parts or consumables: parts that are scrapped after replacement, because maintenance is either technically not possible or economically not viable.

Repairable parts are usually maintained by replacing one or more lower level parts. These lower level parts are called *Shop Replaceable Units* (SRUs). SRUs, like LRUs, are either scrapped or maintained after replacement, and hence there are multiple maintenance levels (*indenture levels*).

An MO consists of multiple maintenance depots, that are responsible for first level maintenance of the assets. An MLO usually consists of multiple internal repair shops, that are responsible for lower level maintenance, and a department that subcontracts lower level maintenance to external repair shops. Maintenance on (a part of) a capital asset is conducted according to a maintenance policy/program or a modification plan, and generates demand for LRUs (SRUs). In total, we can distinguish four types of maintenance (Arts et al. (to appear in 2018)):

- Usage (or age) based maintenance: maintenance that is conducted in order to prevent failure, by replacing parts based on usage or age (time). Since the operational usage of the assets is more or less planned, also the moment that maintenance is to be performed can be planned some time in advance.
- Predictive (or condition based) maintenance: maintenance that is conducted in order to prevent failure, by replacing parts based on condition. The condition of the parts is measured periodically (via inspections) or continuously (remote monitoring). Maintenance is planned based on a prediction of the time-to-failure of parts.
- Corrective maintenance: maintenance that is conducted after a failure has occurred. Corrective maintenance can be partially planned when it involves a non-critical part whose maintenance can be delayed.
- Modificative maintenance: maintenance conducted to improve the performance of the capital asset. This maintenance can usually be delayed until all resources are available.

Each type of maintenance generates demand for spare parts in a different way. Table 1.1 summarizes the different maintenance types, their drivers and spare parts demand characteristics related to timing and content. Spare parts demand of which both timing and content is known upfront is denoted *planned*, all other demand is denoted *unplanned*. The maintenance depots (internal repair shops) request the required LRUs (SRUs) at the MLO by creating spare parts orders. These spare parts orders have a desired lead time, i.e. *delivery lead time*, and the desired delivery date is equal to the desired start date of the maintenance.

Typically, preventive maintenance activities are substantial in user networks. Despite, a data analysis on the VIRM train series (from February 2009 to May 2010) at the NS indicated that 61% of train maintenance (in % of train arrivals) is conducted when

Table 1.1 Maintenance types, including drivers and spare parts demand characteristics.

Maintenance type	Driver(s)	Timing	Content
Usage based maintenance	Time, running hours (age/usage based)	Known	Known
Predictive maintenance	Remote condition monitoring	Unknown	Known
	Periodic inspections	Known	Unknown
Corrective maintenance	Non-deferable defects	Unknown	Unknown
	Deferable defects	Known	Known
Modificative maintenance	Reliability/safety issues, cost reductions	Known	Known

the train enters the maintenance depot unplanned. Moreover, in our case studies we found that 85% (NS) and 82.6% (KLM E&M) of the requests for (repairable) LRUs are unplanned. In Chapters 4-7 of this dissertation, it therefore seems convenient to assume that demand for (repairable) spare parts is unplanned.

1.1.2 Spare parts supply

OEMs and system integrators play an important role in the supply of spare parts that are required to maintain the assets they have sold, and MLOs have little buying power compared to other customers of these companies. Typically the OEMs and system integrators are monopolists (or oligopolists) that dominate the market, and production of spare parts (especially those with low revenues) does not always have the highest priority. Consequently, supply lead times can be long and unreliable, and there are high supply availability risks (suppliers suddenly stop to supply certain spare parts).

The production of RFU spare parts can start from scratch using manufacturing techniques, or it can start with a NRFU part that is maintained. In general, MLOs do not or only rarely produce spare parts from scratch themselves, and consumable parts are mainly replenished from new buys at external suppliers. Repairable parts are maintained by an internal or external repair shop, and differ from consumable parts in two ways. First, repairable parts circulate in a closed-loop (as long as repair of the part is possible): the number of parts is limited (*circulation stock*) and short-term supply (repair) of a part can only start after a part is replaced. Second, repairable parts have two sources of supply. Besides repair, the stock of RFU parts can also be replenished from new buys at external suppliers.

For strategic and economic reasons, our case study companies have internal repair shops that conduct a subset of the repairs of LRUs and SRUs. A strategic reason is that companies want to increase or maintain their (high) knowledge level about the asset technologies. An economic reason is that internal repair lead times can be shorter, because long transportation times to external repair shops are eliminated. Such shorter lead times are only beneficial if the corresponding reduction of stock and (external) repair costs outweigh the costs of the internal repair shops.

1.1.3 Matching spare parts demand and supply

Users are interested in the *uptime* of capital assets and therefore often specify *targets on the availability of the assets*. The targets on the availability of the assets are commonly divided over the time components related to the downtime of the asset. That is, separate targets are set on the (expected or maximum) downtime due to diagnosis and maintenance activities, and on the (expected or maximum) downtime waiting for each of the required resources (including spare parts).

A request for an unavailable (critical) LRU in principle results in an unavailable asset. Because the costs of (unplanned) downtime are in general high, the MLO should be able to quickly match supply and (unplanned) demand for LRUs. Therefore, either part supply (production/repair) lead times should be (very) short (i.e. smaller than the delivery lead time of the request), or spare parts stocks should be high. Because lead times are in general long, we assume that spare parts should be stocked and this can either be done by the MLO or by external parties.

These external parties are usually the external suppliers (repair shops) of the spare parts, or other users with which spare parts are shared (*pooling*). In the case study companies, the MLOs and external suppliers sometimes agree on *parts availability contracts*, i.e. stock owned by a supplier, that is paid by the MLO to realize a certain parts availability with a risk for the supplier to be penalized for (too many) late deliveries. Also sometimes stock, legally owned by the MLO, is stored by the suppliers due to storage limits at the MLO. In that case, the supplier is also responsible for timely supply of the parts and is paid for this service. Although these types of agreements are present in practice, the MLO in general mostly stocks the parts itself. In the remainder of this dissertation, we assume that the MLO stocks all spare parts in its own service network.

Spare parts are stocked in local stock locations close to the maintenance depots, which are replenished from a central stock location. The central stock location is replenished from external suppliers/repair shops and from internal repair shops, that are located close to the central stock location. The central stock location is used to stock spare parts that can be used by multiple maintenance depots (pooling), but also for consolidating replenishment orders placed at external suppliers. In this way, a chain is constructed that supplies spare parts from suppliers/repair shops, via the central stock location and one of the local stock locations (*multi-echelon*), to the maintenance depots: a *spare parts supply chain*. When demand for an LRU cannot be met from the local stock location, emergency procedures such as lateral transshipments or emergency shipments from upstream stocking locations may be applied.

Figure 1.1 presents a typical example of a spare parts supply chain in user networks. The spare parts supply chain is a multi-echelon divergent supply chain with multiple repair shops. Besides the aforementioned stock locations, there is also a stock location for parts that still need to go to repair (NRFU parts) and a stock location for parts required for new projects and modifications. In practice, these stocking points are often in one and the same warehouse but for control reasons these stocks are distinguished.

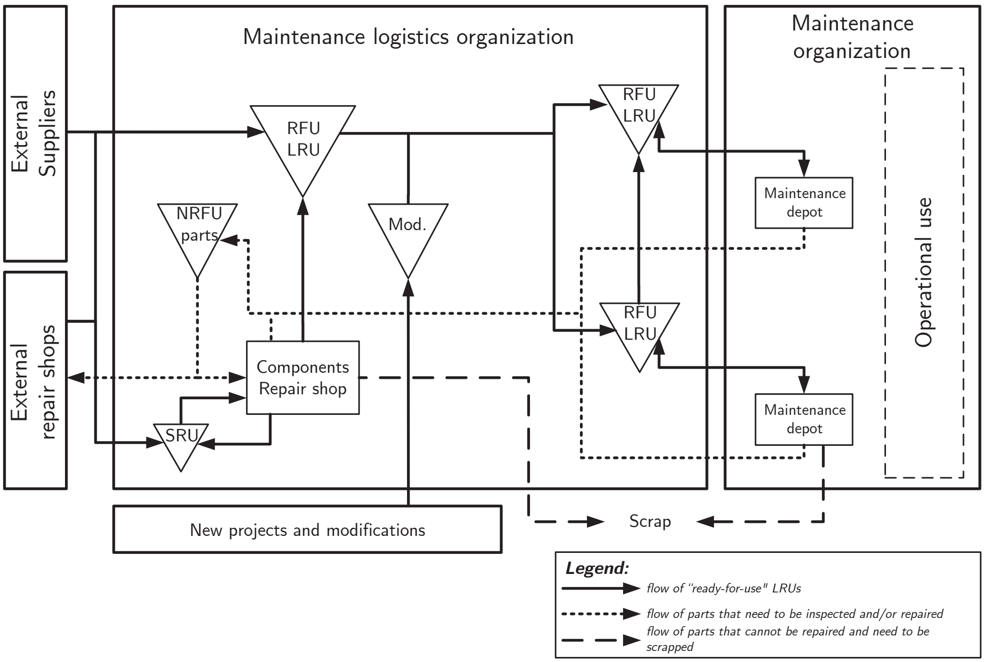

Figure 1.1 *Example of a spare parts supply chain*

MOs and MLOs typically make agreements on specified upper/lower bounds for key performance indicators such as (1) the number of unavailable assets due to unavailable LRUs or (2) the stock availability of (critical) LRUs. From the financial department, on the other hand, the MLO obtains targets on the costs related to the supply of spare parts, such as inventory holding costs and costs related to the procurement, repair and distribution of spare parts. The MLOs try to find the optimal balance between the availability and costs of spare parts, within their span of control.

1.2 Repairable parts inventory systems

Typically an MLO consists of multiple internal repair shops that supply ready-for-use LRUs. Next to this, ready-for-use LRUs are supplied by external supply sources. An MLO faces a large optimization problem in which a trade-off is made between the availability and costs of all spare parts. In Section 1.2.1 we describe this optimization problem and show how and under what conditions this problem can be decomposed into a problem for all spare parts of each supply source separately. Hence this also holds for a repairable parts inventory system with *one internal repair shop*, and this setting is the scope of Chapters 4 to 7. Section 1.2.2 presents an overview of the relevant decisions and a description of (the interface between) the two relevant decision makers in such systems.

1.2.1 Definition and objective

MLOs stock LRUs at multiple locations in order to achieve a certain availability of the high-value capital assets (from a spare parts availability point of view) against acceptable logistics costs. In this section, we pretend that the network has a single ready-for-use LRU stock location. In the user networks we consider, the transportation times between the central and the local stock locations are in general small compared to the replenishment lead times of the central stock location. Therefore, we now consider the entire network as one 'virtual' stock location from a planning and control point of view. We however emphasize that this assumption is not required to make the decomposition described in this section. That is, this decomposition also holds under e.g. a two-echelon network with a central stock location and multiple local stock locations.

In general, MLOs stock a set of LRUs I required to maintain a set of asset types (or systems) K. Each asset type $k \in K$ consists of a subset of LRUs $I^{(k)} \subseteq I$. Demands for LRU i for an asset of type k occur according to a Poisson process with constant rate $\mu_{i,k}$; $\mu_i = \sum_{k \in K} \mu_{i,k}$ denotes the total demand rate for parts of LRU i.

The LRU stock is replenished from a set of supply sources J. The set of LRUs that is replenished by source j is denoted by $I_j \subseteq I$. Each LRU has one (preferred) supply source, i.e. $I_a \bigcap I_b = \emptyset$, $\forall a \neq b$ and $a, b \in J$. Hence the LRUs are partitioned over the supply sources, and $j^{(i)}$ represents the (preferred) supply source for LRU i. A graphical representation of this spare parts supply chain is given in Figure 1.2.

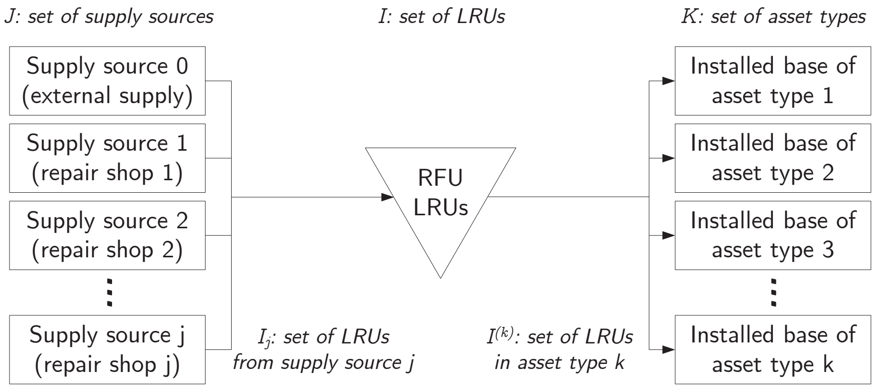

Figure 1.2 *Graphical representation of a single-location spare parts network.*

We assume that each LRU $i \in I$ is replenished according to a base stock policy with base stock level S_i. Let $\mathbf{S}^{(j)} = \{S_i\}_{i \in I_j}$ be the vector of base stock levels for the LRUs replenished from supply source j, and let $\mathbf{S}$ be the vector of all LRU base stock levels.

Let ψ_j represent the repair/production policy of supply source j, and also includes the choice for working overtime, extra capacity for work stations, etc. We assume that $\psi_j \in \Psi_j$, where Ψ_j is a class of repair policies that may depend on the total state in supply source j, that is, it may depend on the actual stock levels of all LRUs $i \in I_j$ and on the base stock levels of these LRUs ($\mathbf{S}^{j^{(i)}}$).

The total costs for supply source j are denoted by $C_j(\mathbf{S}^{(j)}, \psi_j)$, and consist of the holding costs for all LRUs $i \in I_j$ and the costs for applying repair/production policy ψ_j. The total costs (over all LRUs/supply sources) are denoted by $C(\mathbf{S}, \psi)$, where $\psi = \{\psi_j\}_{j \in J}$.

Because of the direct link between the availability of LRUs and the availability of the assets, it is common to use system-oriented service targets (Basten and Van Houtum, 2014). We denote $EBO_i(\mathbf{S}^{j^{(i)}}, \psi_{j^{(i)}})$ as the mean number of backorders for LRU i, and set a constraint on the aggregate mean number of backorders for each asset type k: EBO_k^{obj}. We then obtain the following optimization problem (P):

$$
\begin{aligned}
(P) \qquad & \min \qquad C(\mathbf{S}, \psi) \\
& \text{subject to} \sum_{i \in I^{(k)}} \frac{\mu_{i,k}}{\mu_i} EBO_i(\mathbf{S}^{j^{(i)}}, \psi_{j^{(i)}}) \leq EBO_k^{obj} \quad \forall k \in K \\
& \mathbf{S} \in \mathbb{N}_0^{|I|} \\
& \psi \in \Psi^{|J|}
\end{aligned}
$$

The Lagrange relaxation technique (see e.g. Porteus (2002)) can be used to solve problem (P), and its Lagrange function is defined as:

$$
L(\mathbf{S}, \psi, \lambda) := C(\mathbf{S}, \psi) + \sum_{k \in K} \lambda_k \left(\sum_{i \in I^{(k)}} \frac{\mu_{i,k}}{\mu_i} EBO_i(\mathbf{S}^{j^{(i)}}, \psi_{j^{(i)}}) - EBO_k^{obj} \right)
$$

where $\lambda = (\lambda_1, ..., \lambda_{|K|})$ is the vector of Lagrange multipliers λ_k, with $\lambda_k \geq 0, \quad \forall k \in K$. The Lagrange function can be rewritten as:

$$
L(\mathbf{S}, \psi, \lambda) := \sum_{j \in J} L_j(\mathbf{S}^{(j)}, \psi_j, \lambda) - \sum_{k \in K} \lambda_k EBO_k^{obj}
$$

where

$$
L_j(\mathbf{S}^{(j)}, \psi_j, \lambda) := C_j(\mathbf{S}^{(j)}, \psi_j) + \sum_{i \in I_j} EBO_i(\mathbf{S}^{j^{(i)}}, \psi_{j^{(i)}}) \sum_{k \in K} \left(\frac{\lambda_k \mu_{i,k}}{\mu_i} \right)
$$

is the decentralized Lagrangian function and in which $\sum_{k \in K} \left(\frac{\lambda_k \mu_{i,k}}{\mu_i} \right)$ can be regarded as the weighted average backorder cost parameter for LRU i. The Lagrangian dual function is given by

$$
\begin{aligned}
H(\lambda) &= \min_{\substack{\mathbf{S} \in \mathbb{N}_0^{|I|} \\ \psi \in \Psi^{|J|}}} \quad L(\mathbf{S}, \psi, \lambda) \\
&= \min_{\substack{\mathbf{S} \in \mathbb{N}_0^{|I|} \\ \psi \in \Psi^{|J|}}} \quad \sum_{j \in J} L_j(\mathbf{S}^{(j)}, \psi_j, \lambda) - \sum_{k \in K} \lambda_k EBO_k^{obj} \\
&= \sum_{j \in J} \min_{\substack{\mathbf{S} \in \mathbb{N}_0^{|I|} \\ \psi \in \Psi^{|J|}}} \quad L_j(\mathbf{S}^{(j)}, \psi_j, \lambda) - \sum_{k \in K} \lambda_k EBO_k^{obj}
\end{aligned}
$$

The minimization can be moved inside the summation because $L_a(\mathbf{S}^{(a)}, \psi_a, \lambda)$ does not depend on any combination of $\mathbf{S}^{(b)}$ and ψ_b, $\forall a \neq b$ and $a, b \in J$. The Lagrangian dual problem (D) is defined as:

$$
(D) \qquad \max_{\lambda \in \mathbb{R}_+^{|K|}} \qquad H(\lambda)
$$

Solving the dual problem (D) yields a good solution for problem (P), however solving this dual problem is not trivial. Here we restrict ourselves to the observation that problem (D) can be decomposed into a problem that can be solved for each supply source j individually, $\forall j \in J$. Still, in each subproblem per supply source dependencies exist between stock levels $(\mathbf{S}^j)$ and the applied policy to repair or produce parts (ψ_j).

Up to now, we excluded SRUs from the problem formulation. Let X_j be the set of SRUs required by supply source j, $\forall j \in J$, and let $\hat{\mathbf{S}}^{(j)}$ be the corresponding vector of base stock levels. We assume that each SRU $x \in X_j$ is either repaired by supply source j itself, or replenished from external repair shops against a given lead time. We can now replace $\mathbf{S}^{(j)}$ in problem (P) by $(\mathbf{S}^{(j)}, \hat{\mathbf{S}}^{(j)})$, and again the problem decomposes into subproblems that can be solved for each source individually.

This decomposition result is important, because it means that we can consider each repair shop (supply source) individually. In Chapters 4 to 7, we therefore restrict ourselves to repairable parts inventory systems with one internal repair shop.

1.2.2 Interface between repair shops and inventory control

According to Hausman and Scudder (1982), the relevant decisions in a repairable (parts) inventory system concern: (1) the inventory policy (parameters) of the LRUs

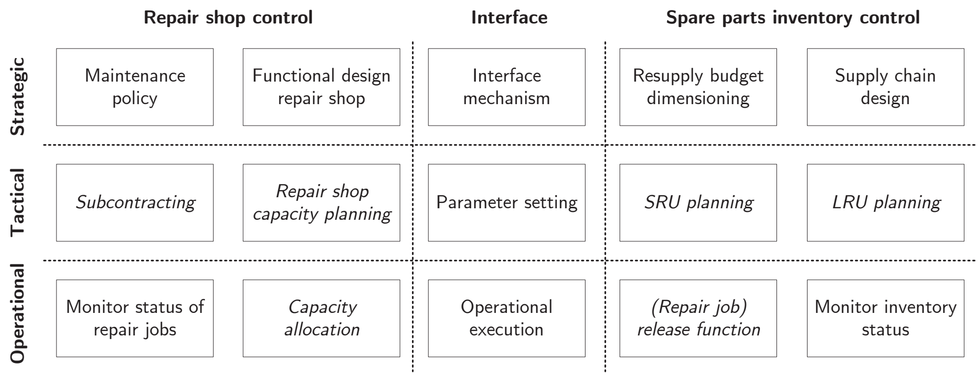

Figure 1.3 *Decisions related to repair shop and inventory control of repairable spare parts*

and SRUs, (2) the repair capacities available to repair failed parts and (3) the scheduling system used to control the flow of work in the repair shop. Figure 1.3 outlines decision functions related to both inventory and repair shop control on the strategic, tactical and operational level. We concentrate on decisions that are directly interrelated; these decision functions are printed in *italics*. Strategic decisions regarding the design and layout of repair shops and the repair process are beyond the scope of this dissertation.

In practice, decisions in repairable parts inventory systems are separated over two departments: an inventory control department, responsible for setting stock levels, and the repair shop, responsible for conducting repairs. Because of the interrelation of the decisions, this requires a certain coordination (set of agreements) between the inventory control department and the repair shop. The agreements between the repair shop and the inventory control department are arranged in an *interface agreement*. This interface agreement consists of decisions at the strategic, tactical and operational level.

At the strategic level, the interface mechanism prescribes how the repair shop and the inventory control department cooperate. For example, the two departments may have agreed that the repair shop has a fixed repair lead time for each part, and that the repair shop is responsible for realizing these lead times. At the tactical level, the (interface) parameters need to be set. In our example, this implies to setting the lead time for each repairable part. At the operational level, the interface agreements are executed. In our example this means that the inventory control department measures the performance of the repair shop, e.g. the fraction of repair jobs in which the actual repair lead time did not exceed the lead time that was agreed upon.

1.3 Research questions & methodology

The main objective of this dissertation is to develop control concepts for MLOs that lead to efficient combinations for availability and costs of spare parts in user networks. These concepts relate to one or more decisions that have to be made by the MLOs. To ensure the practical relevance, we take practice 'as is' and focus on decision functions that are most important and influenceable for MLOs. Case studies are used to extensively describe reality, and to gain insights in these decision functions. By comparing the case study findings with findings from literature reviews, we identify open research topics that are addressed in the following chapters of this dissertation. In this way we also ensure the scientific relevance of this research. We start this research with the following research question:

1. *What are the relevant decision functions in MLOs, and how do they relate to each other?*

Chapter 2 addresses this research question, and is based on Driessen et al. (2015a). We present a normative framework for MLOs in which we decompose all relevant decisions, and group coherent sets of decision functions into processes based on their function. We describe how the decision functions relate to each other and we provide references to state-of-the-art techniques for supporting these decisions. This framework is based on an extensive literature review.

In practice we observe solutions/approaches to decision making in MLOs that we do not understand, i.e. that deviate from the framework. We therefore aim to get a better understanding of decision making in MLOs in practice and its differences with the framework. This raises the following research questions, that give direction to the topics of the remainder of this dissertation:

2a. *Which processes/decision functions in the framework are observed in practice?*

2b. *Which processes/decision functions observed in practice are not included in the framework?*

2c. *Which (interfaces and information flows between) processes/decision functions are filled in differently by the case study companies, and why?*

To answer these research questions, in Chapter 3 we use our framework to study the current design of the spare parts planning and control system in case studies at six prominent companies. First, these case studies reveal that the decisions in repairable parts inventory systems have the highest practical relevance for MLOs in terms of influence on availability and costs of spare parts. Second, in terms of scientific relevance, we find that implementation of our framework has practical implications. We find that decisions are separated over two decision makers, and our framework design may lead to situations in which control and responsibilities are not aligned. This is discussed in more detail in Chapter 3.

The relevant decisions that are in scope of a repairable parts inventory system are the inventory policies of the LRUs and SRUs, the fixed and flexible repair capacities and the scheduling policy of the repair shop (see also Section 1.2.2). Clearly these decisions interrelate, and both repair capacity and inventory can serve as a buffer to meet these spare part availability requirements. Hence we need to find a control concept that optimally balances them. This raises the third research question of this dissertation:

3. *What are ways to design control concepts for repairable parts inventory systems, and which requirements are important to account for in the design process?*

To answer this research question, we review three streams of literature in Chapter 4. We discuss how available literature can be used to design control concepts for repairable parts inventory systems, and what the requirements for application are. Furthermore, we formulate control principles that should be taken into account when designing repairable parts inventory systems. So far, it is however not clear whether (the models from) the reviewed streams of literature can be applied in practice, and in which situations application is beneficial. Therefore, we formulate the fourth research question:

4. *Which models and control principles, identified in the literature review of Chapter 4, are applicable to design control concepts for the repairable parts inventory systems in the case study companies?*

We validate the applicability of the three streams of literature in Chapter 5, by comparing the application requirements and control principles in multiple case studies. From this assessment, we can conclude that there is no repair shop to which the models and control principles can be applied directly, because at least one assumption is not valid. We found that a varying set of repair shops complies to a subset of requirements and control principles per literature stream. This can be explained by the clear differences we observed between repair shops, especially when it comes to the complexity and uncertainty of the materials and capacities required in the repair process. How this impacts the designs of control concepts for repairable parts inventory systems is not clear, and this raises the fifth research question:

5. *To what extent do repair shops differ from each other, and what implications do these differences have on the design of control concepts for repairable parts inventory systems?*

Chapter 6 addresses this research question, and is based on Driessen et al. (2015b). We present a typology of repair shops, based on distinctive levels of complexity and uncertainty with respect to the materials and capacities required in the repair process. We present control concept designs for repairable parts inventory systems, for each stereotype repair shop. We however only qualitatively motivate the underlying design choices; quantifying the impact of the design choices is necessary to test and further improve the designs.

In the designs of Chapter 6, we propose to decouple the inspection and repair phase of the repair process for repair shops with a high material uncertainty. This raises the question whether priority should be given to inspection of failed parts (to reveal which SRUs are required and speed up sourcing of unavailable SRUs) or to completion of the repair of parts for which all required SRUs are available (to increase the availability of LRUs). At the same time, consciously interrupting repair jobs is inefficient; the repair times depend on the time that elapses between inspection and the actual repair of a part. This raises the final research question of this dissertation:

6. *What is the optimal capacity priority policy in repairable parts inventory systems with high material uncertainty?*

In Chapter 7 we set up a simulation model and compare several workload-based capacity priority policies.

1.4 Outline of the dissertation

Figure 1.4 summarizes our research approach and also includes an outline of this dissertation. The chapters in this dissertation have a different scope. In Chapters 2 and 3 we focus on MLOs, as described in Section 1.1. The scope of Chapters 4 to 7 is a repairable parts inventory system, as described in Section 1.2. Chapter 8 concludes this dissertation, and provides some suggestions for further research. Table 1.2 summarizes the aim and methodology of Chapters 2-7.

Table 1.2 Overview of methodology and aim of each chapter.

Chapter	Aim	Methodology
2	Develop a normative decision framework for MLOs	Design research based on a literature review
3	- Validate the completeness and relevance of decision functions in the framework, including their mutual connections - Identify gaps between framework and case studies - Identify most important and influenceable decision functions in MLOs	Empirical research: case studies at 6 MLOs
4	- Identify ways to design control concepts for repairable parts inventory systems, including requirements for application - Formulate principles for control of repairable parts inventory systems	Literature review
5	Validate the applicability of models and control principles, identified in Chapter 4, to design repairable parts inventory systems in practical applications	Empirical research: case studies at 22 repairable parts inventory systems
6	Design control concepts for repairable parts inventory systems, for each of the four repair shop types recognized in practice	Design research, supported by insights from literature (Chapter 4) and case studies (Chapter 5)
7	Validate design choice ‘decouple inspection and repair phase’ for repair shop type II	Quantitative modeling

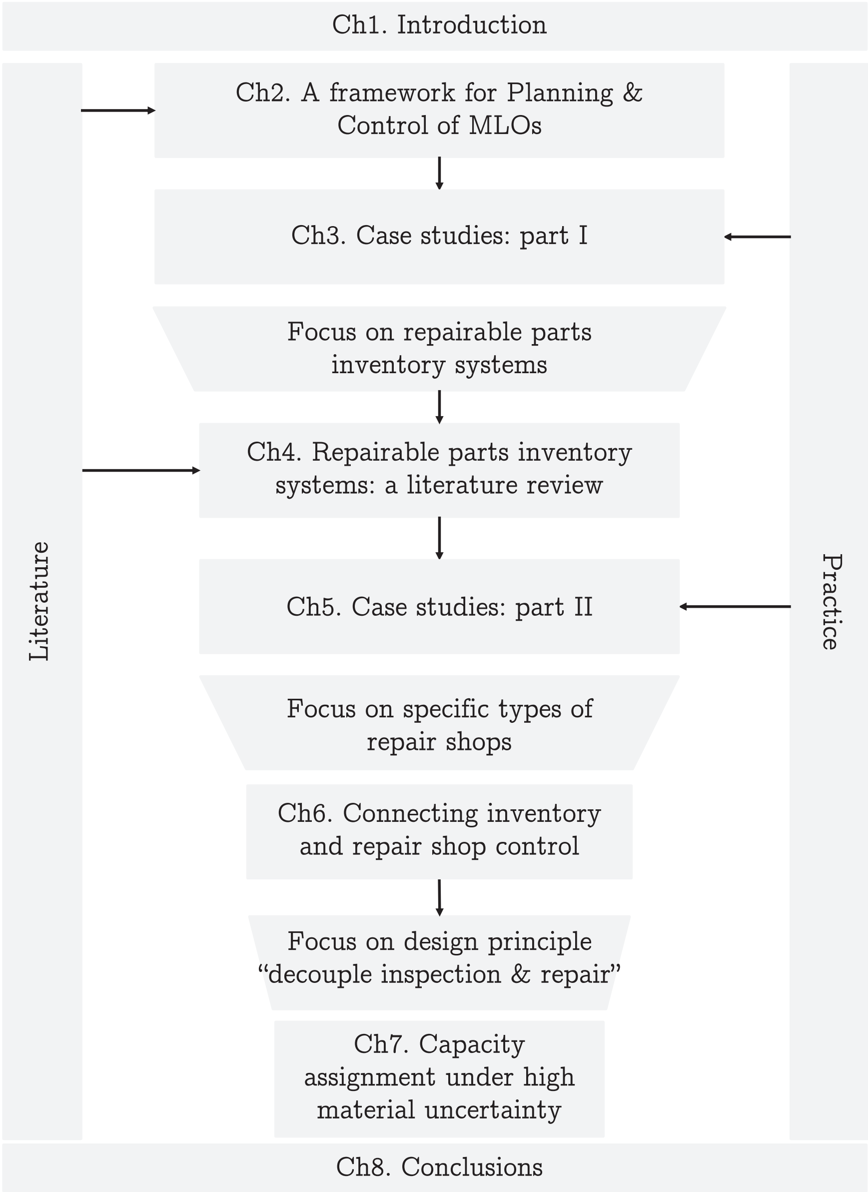

Figure 1.4 *Research approach and outline of the dissertation*

Chapter 2

A framework for planning & control of MLOs

"Tactiek is niet je man uitschakelen of zo. Nee, het is weten waar je mee bezig bent."

Johan Cruijff

2.1 Introduction

In this chapter we focus on the responsibility of an MLO to optimally balance the availability and costs of spare parts. The main contribution of this chapter is the development of a normative framework for planning and control of MLOs, and is based on an extensive review of spare parts literature. In the framework we decompose the main decision functions and describe how they relate to each other. With this decision framework for MLOs we address the first research question of this dissertation:

1. *What are the relevant decision functions in MLOs, and how do they relate to each other?*

To the best of our knowledge, the only other paper in which a framework for spare parts control has been proposed is Cavalieri et al. (2008). Their framework focuses exclusively on the inventory control of spare parts and provides rules of thumb concerning decisions in the framework. We take a broader perspective by incorporating a repair shop and its control and provide references to state-of-the-art techniques for supporting decisions in the framework. The framework itself however does not provide details on any of the decision functions included therein. Rather, it provides a decomposition of all the decisions and how they relate to each other.

The design of our decision framework is based on an extensive review of spare parts literature. The Supply Chain Operations Reference (SCOR) model, developed by the supply chain council and widely accepted as the cross-industry standard for supply chain management, served as a starting point for the design by providing an overview of the main supply chain management decision functions. The unique aspects of decision making in MLOs were gathered from review papers on the management of (repairable) spare parts (Kennedy et al. (2002); Guide Jr. and Srivastava (1997); Basten and Van Houtum (2014)) and an extensive review of spare parts literature.

This chapter is organized as follows. Section 2.2 presents the framework and describes the decisions therein, and Section 2.3 provides relevant references for each part of the decision framework to aid in decision making. Section 2.4 concludes this chapter.

2.2 A decision framework for MLOs

In this section, we present the framework for MLOs. Figure 2.1 presents an overview of processes and decisions in MLOs, including their mutual connections. We distinguish eight different processes, which are numbered one up to eight in Figure 2.1. Within each process, we distinguish different decision levels. Decisions that are not made very frequently, i.e. once a year, are marked ‘S/T’ (strategic/tactical decisions); decisions made regularly, i.e. once a month or quarter, are marked ‘T’ (tactical decisions) and decisions made frequently, i.e. once a day/week, are marked ‘O’ (operational decisions).

An arc illustrates that information, e.g. data or outcomes of decisions, flows from one process to another. This information is needed to make decisions in subsequent processes. We emphasize that there are many feedback loops between the various processes. For readability, these feedback loops are left out of the figure. In Section 2.2.1-2.2.8, we explore each process in Figure 2.1 and describe in more detail what information is passed along the arcs in Figure 2.1. Thus, Figures 2.2-2.9 should be interpreted as more detailed zoom-ins of different parts of Figure 2.1.

The framework we provide will need refinement and alterations for every particular organization, but it serves as a useful starting point in making specific designs of spare part planning and control systems.

2.2.1 Assortment management

Assortment management is concerned with the decision to include a spare part in the assortment and to maintain technical information of the included spare parts. We emphasize that including a part in the assortment does not necessarily mean that this part will also be stocked. The process of managing the assortment can be found in Figure 2.2.

2.2.1.1 Define spare parts assortment

The decision to include (exclude) a part in (from) the assortment is usually taken shortly after procurement (phase out) of a (sub)system and strongly depends on the maintenance policy/program. Parts are excluded from the assortment in case its applicability has expired. There are two options when to include a part in the assortment: before or after the first need for the part.

In case a part is included in the assortment, then there is a possibility that the part is never needed during its lifecycle. Time spent on collecting information and adding the part to the database has been done without any use, which results in unnecessary operational costs.

However, in case a part is not included in the assortment, there are two possible adverse consequences. First, when the part fails and a supplier is still available, the lead time of the part is higher due to data collection and negotiation actions. Second, when the part is needed there may not exist any suppliers for it anymore. In this case, the part may have to be custom made. To do this, in many cases specialized technical information regarding the form, fit and function is needed. If a part is not included in the assortment, this information is not available.

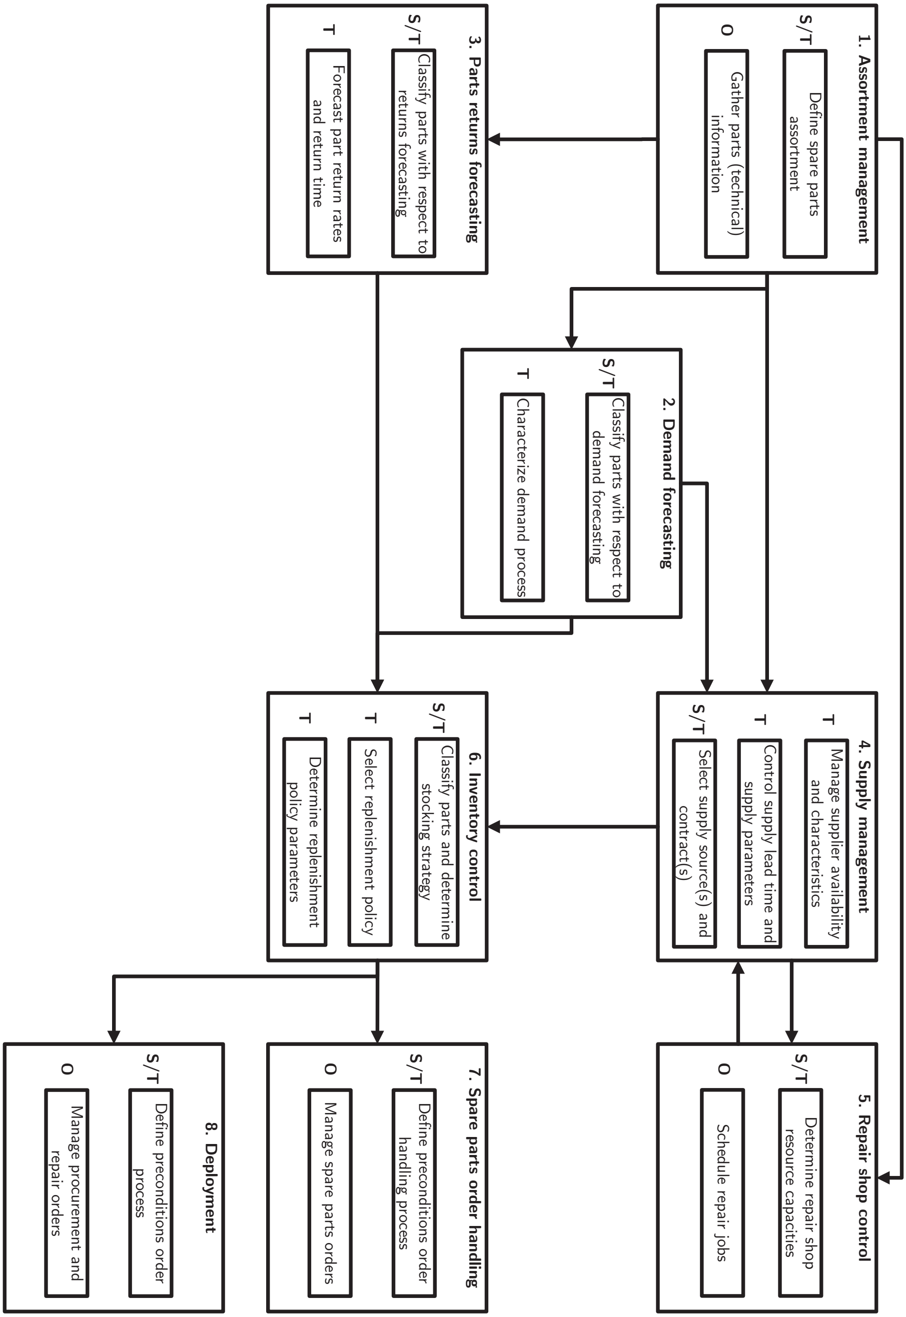

Figure 2.1 *Overview of processes and decisions in maintenance logistics control.*

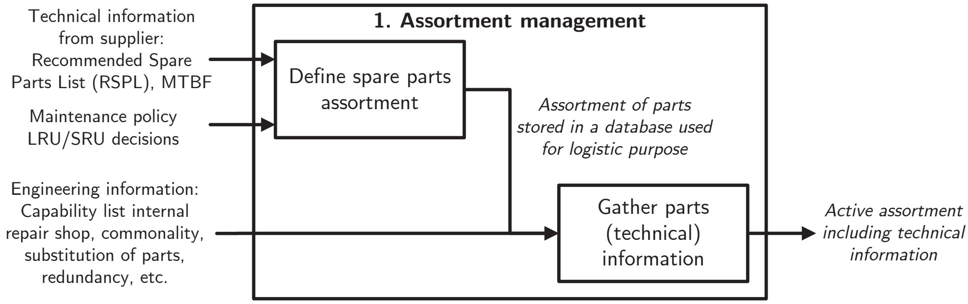

Figure 2.2 *Process of managing a spare parts assortment.*

2.2.1.2 Gather parts (technical) information

Once a part is included in the assortment, (technical) information of the part is gathered and updated when necessary. The MLO needs to decide whether or not to gather and maintain parts technical information that is important for spare parts planning and control: (i) criticality, (ii) redundancy, (iii) commonality, (iv) specificity, (v) substitution, (vi) shelf life, (vii) position in the configuration[1] and (viii) repairability. Additionally technical information regarding form, fit and function may be gathered. We also distinguish so called 'insurance' spare parts.

- Parts criticality is concerned with the consequence of a part failure, that is the type of breakdown and reaction time (Sherbrooke, 2004). Full (partial) system breakdown means that the system is non-operational for (a part of) all assigned usage purposes. Parts that cause full system breakdown are denoted *(partially) critical*. Parts that cause no system breakdown, i.e. the system can be used for all assigned use purposes, are denoted 'non-critical'.

- Parts redundancy is the duplication of system components (parts) with the intention to increase the reliability of the system, usually used to serve as a backup for an operating part that may fail. Information on parts redundancy decreases the number of stocked spare parts as it is known in advance that part failure does not cause immediate system breakdown.

- Parts commonality concerns parts that occur in the configuration of multiple

[1] The configuration is similar to the Bill of Materials. However the configuration changes throughout the lifetime of an asset due to modificative maintenance whereas the bill of materials is a snapshot of the configuration at the time of initial manufacture.

systems that are maintained by the MO. For each system, the MLOs needs to meet a certain service level. Information on parts commonality is needed for customer (system) service differentiation in spare parts planning as well as for the decision where to stock parts, i.e. locally or centrally.

- The specificity of a part concerns the extent to which a part is tailored for and used by a customer. Parts availability at suppliers is usually low, if not zero, for specific parts and hence this might effect the size of the buffer stock needed.
- Parts are substitutional in case different parts have the same form, fit and function. This means that requests for one part can be met by a substitute part. Information on parts substitution is used to prevent stocking parts for which requests can also be met by a substitute part.
- The shelf life of a part is the recommended time period during which products can be stored and the quality of the parts remains acceptable for usage. This information is used to prevent stocking too many parts that are scrapped or revised after the shelf life of the part has expired.
- The configuration is a list of raw materials, sub-components, components, parts and the quantities of each that are currently in a system. Hence this list contains all the SRUs and LRUs in the system that may require maintenance during its use. The position of a part in the configuration is needed to determine at which level parts (SRUs) can be replaced, in order to repair an LRU, and what quantity of each SRU is needed. These different levels in the configuration are also called indenture levels. The initial configuration is usually provided by or available at the OEM and coincides with the bill of materials.
- Parts repairability concerns the identification whether a part is technically repairable and if so, whether or not the internal repair shop has the authorization (from the OEM) and the capability to repair the part. This information is needed to determine the parts supply options.
- Technical information on form, fit and function comes in many forms depending on the technological nature of the part involved. Sometimes this information is of a sensitive nature and the OEM may charge extra for this information and/or requires non-disclosure type contracts.
- ‘Insurance’ spare parts are parts that are very reliable, highly ‘critical’ to system availability and not readily available in case of failure. Often these parts are

far more expensive to procure after the initial buy of the system, compared to buying at the moment of initial system purchase. Because of their high reliability, these spare parts often will not be used during the lifetime of the system. Example of an 'insurance' part is the propeller of a ship.

Parts (technical) information is sometimes provided by the OEM. However, it is also possible that the MLO needs to determine this technical information. All the technical information is used to improve stocking decisions and manage supply risks.

2.2.2 Demand forecasting

Demand forecasting concerns the estimation of demand for parts in the (near) future. Future demand for spare parts is either (partially) planned or unplanned and is characterized in Section 1.1.1. MLOs need to decide whether to use information about planned demand. The demand forecasting process is visualized in Figure 2.3.

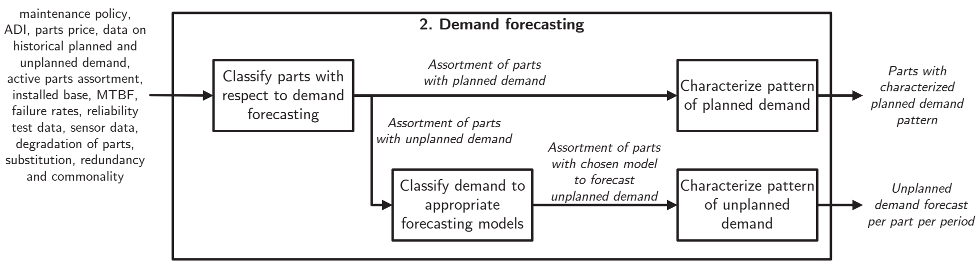

Figure 2.3 *Overview of the demand forecasting process.*

2.2.2.1 Classify parts with respect to demand forecasting

Two types of spare parts are considered: parts for which advance demand information (ADI) is used or not. Using ADI usually decreases the overall forecast error. On the other hand, it is clear that using ADI increases the difficulty, the effort and hence the operational costs to forecast demand (Hariharan and Zipkin, 1995; Benjaafar et al., 2011). In case there is no information available or it is decided not to use it, then all demand is accumulated and one single demand stream is considered. Otherwise, two demand streams (planned and unplanned) are separated.

Within unplanned demand, another classification is made to aid the decision of using a particular forecasting technique. Two factors that determine what methods are appropriate are the interarrival time of demand moments and the variability of demand size. When the time between demand moments is very long, then demand is said to be intermittent. When intermittence is combined with variable demand sizes, demand is said to be lumpy.

2.2.2.2 Characterize demand process

After deciding whether or not to separate demand streams, the demand process needs to be characterized on behalf of the following three purposes: (i) to determine the number of parts to stock, (ii) to determine the repair shop capacity and (iii) to provide the necessary input for updating and characterizing supply contracts.

Planned demand during the delivery lead time is deterministic, since all parts are requested at least early enough in advance. Expected planned demand after this delivery lead time is not known in advance exactly and should hence be forecasted.

Demand for spare parts could also be partially planned in advance. Parts that are needed in about $x\%$ of some types of preventive maintenance inspections are termed *$x\%$-parts*. Combining these probabilities (x) with the information on planned inspections, a forecast for this planned demand stream can be made.

Several methods are applicable to forecast unplanned demand. The first method to forecast unplanned demand is reliability based forecasting. The goal of this method is to forecast parts requirements based on part failure rates, a given installed base and operating conditions. This method determines the failure rate of one part and extrapolates the failure rate to the installed base and varying operating conditions.

The second method to forecast unplanned demand is time series based forecasting. Based on known historic requirements, extrapolations are made using statistical techniques. Examples of well-known time series based forecasting techniques are Moving average, Smoothing methods, Croston's method and bootstrapping (Silver et al., 1998; Croston, 1972; Willemain et al., 2004). The advantage of time series based forecasting is that only historical demand data is needed to forecast demand. Disadvantage is that manual changes to the demand forecasts need to be made in case the installed base changes. The result of characterizing demand is a demand distribution per part per period.

Technical information about substitution is used in forecasting to determine demand for new parts that substitute old parts. Combining demand streams, for different parts that can be met by the same spare part (i.e. substitutes), increases the overall demand forecast reliability. Technical information on commonality of new parts is used to determine how usage in different capital assets affects demand.

2.2.3 Parts returns forecasting

Parts requested by the MOs are sometimes returned in ready-for-use (RFU) condition. In case it is not known which part causes system breakdown, sometimes all parts that may be the cause are requested. After it is found out which part caused the breakdown, unused RFU parts are returned to the original stocking point within an agreed *hand in time*. If the requested part was a repairable, a part is always returned that either (i) needs repair, (ii) is ready-for-use or (iii) is beyond repair and will be scrapped. The MLO needs to account for return rates and hand in times in their planning and control.

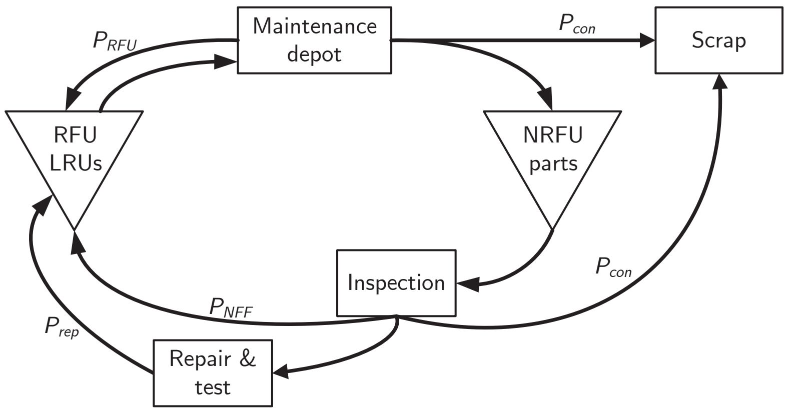

Figure 2.4 *Overview of different part return streams.*

Consider the case of consumables. Here parts are either returned ready-for-use (with probability p_{RFU}) or not at all (with probability p_{con}), see also Figure 2.4. The question is now whether a part *request* should be considered a part *demand* where only part demand influences replenishment decisions. If a procurement order is placed and the part is handed in afterwards, the inventory levels grow unnecessarily.

The case of repairables is different, because repair jobs cannot be released until a failed item is sent to the MLO and often repairable parts have to be tested/certified (again) after their (intended) use, before they are RFU again. Let p_{RFU} denote the probability that a returned part is ready-for-use, p_{rep} denote the probability that a returned part needs repair, let p_{NFF} denote the probability that the inspected part requires no repair ('no fault found') and let p_{con} denote the probability that a returned part will be condemned after inspection (see Figure 2.4). These return fractions are used by inventory control as follows. Requests for parts can be considered as demand and with probability p_{RFU} the lead time is equal to the hand in time, with probability p_{rep} the lead time is the sum of the return lead time, the inspection lead time and the repair and test lead time, with probability p_{NFF} the lead time is the sum of the return lead time and the inspection lead time and with probability p_{con} the lead time is the sum of the hand in time, the inspection lead time and the procurement lead time.

The most straightforward technique to forecast return rates is to use historic return rates, possibly corrected for special events such as unusual accidents. For most parts, this technique is sufficiently accurate. For some parts different failure modes often correspond to different types of returns. Techniques from reliability engineering can be used to estimate these return rates.

2.2.4 Supply management

Supply management concerns the process of ensuring that one or multiple supply sources are available to supply ready-for-use LRUs, as well as SRUs, at any given moment in time with predetermined supplier characteristics, such as lead time and underlying procurement contracts (price structure and order quantities). A process overview of supply management can be found in Figure 2.5.

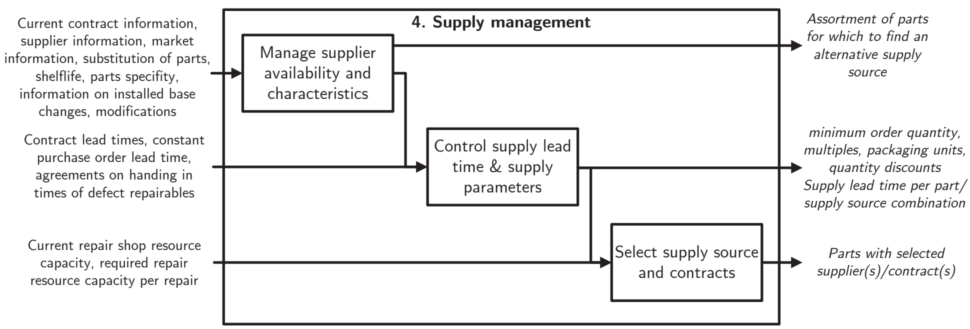

Figure 2.5 *Process of managing the supply of spare parts.*

2.2.4.1 Manage supplier availability and characteristics

The process of managing supplier availability and characteristics within MLOs is concerned with having one or more supply sources available for each spare part in the assortment, including supply characteristics. MLOs have several possible supply types: (i) internal repair shop, (ii) external repair shops, (iii) external suppliers and (iv) re-use of parts, i.e. parts that are known to become available in a short time period caused by phase out of end-of-life systems, possibly used by other companies. Moreover, within each supply type, it is possible to have multiple supply sources.

Each part has either one or more supply sources, one supply source that is known to disappear within a certain time period or no available supply source at all. In the latter case, the MLO needs to find an alternative supply source for all parts that need future resupply. Alternative supply sources are e.g. a new supplier/repair shop, a substitute part or the part needs to become a repairable instead of a consumable, if technically possible.

When the only supply source of a part is known to disappear, the MLO needs to decide whether to search for an alternative supply source or to place a final order at the current supply source. The final order decision concerns the determination of a final order quantity that should cover demand during the time no supply source will be available (Teunter and Klein Haneveld, 1998; Van Kooten and Tan, 2009). The supply availability for these parts is guaranteed through the available inventory. Managing supply availability is also concerned with timely updating and maintaining current contracts with external suppliers.

MLOs also need to gather and maintain information on supply characteristics. Information concerning the following matters is needed to determine the supply lead time (distribution) and to select a (preferred) supply source and contract: (i) contractual or historical repair/new buy price(s) of the part, (ii) quantity discounts (iii) contractual lead and/or repair lead time, (iv) minimum order quantities and (v) multiples.

2.2.4.2 Control supply lead time and supply parameters

The supply lead time consists of: (i) repair or supplier lead time, (ii) procurement time, (iii) picking, transport and storage time of parts and, in case of repairables, (iv) hand in times of failed repairables. For all these components of the supply lead time, agreements are made on planned lead times.

Using planned lead times for internal repair is justified because: (i) MLOs make agreements with the internal repair shop on planned repair lead times and (ii) the repair shop capacity is dimensioned in such a way that internal due dates are met with high reliability. Using planned lead times for external supply is justified because MLOs agree on contractual lead times with their external suppliers.

The supply lead time is determined for each part/supply source combination separately. We distinguish two types of supply lead times: (i) repair lead time and (ii) procurement lead time. For all parts that are known to be 'technically repairable', the MLO needs to determine the procurement lead time, the external repair lead time and the internal repair lead time, in case internal repair is possible. For consumables only the procurement lead time needs to be determined.

2.2.4.3 Select supply source and contracts

The MLO needs to make sure that spare parts can be replenished at any given time. For this purpose, the MLO needs to set up contracts with one or multiple supply sources in a cost efficient way. The decision is based on the following costs incurred while selecting a supply source: (i) setup and variable costs of the repair shop capability and resources, (ii) setup costs of the contract, (iii) procurement or repair costs and (iv) inventory holding costs.

The MLO uses information on supply characteristics and supply lead times to select one or more supply sources out of all possible part/supply source combinations. Important in selecting a supply source is the decision whether to designate a spare part as repairable or consumable. Alfredsson (1997) states: "The task of determining whether an item should be treated as a discardable (consumable) or repairable item is called *level-of-repair-analysis* (LORA). If the item is to be treated as a repairable item, the objective is also to determine where it should be repaired". See also Basten et al. (2009) and MIL (1993) for analogous definitions of LORA.

The LORA is reconsidered each year or in case of substantial changes in the asset base. The outcome of the LORA is used to reconsider the internal repair shop resource capabilities. Note that MLOs need to set up a contract for repair as well as for new buy of repairables.

2.2.5 Repair shop control

The repair shop in the spare parts supply chain functions much like a production unit in a regular supply chain. At the interface with supply management, agreements are made on lead times for the repair of each LRU. Also agreements are made on the load imposed on the repair shop so that these lead times can be realized. For example it is agreed to release no more than y parts for repair during any week.

To comply with these lead time agreements, decisions are made at a tactical and operational level. At the tactical level the capacity of the repair shop is determined and at the operational level, repair jobs are scheduled to meet their due dates. A schematic overview of repair shop control is given in Figure 2.6.

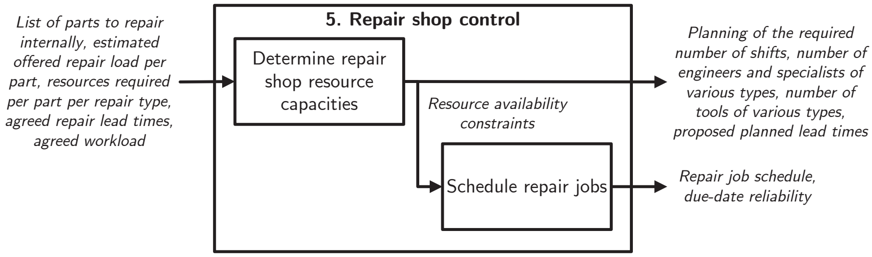

Figure 2.6 *Overview of repair shop control process.*

2.2.5.1 Determine repair shop resource capacities

When a repair job enters the repair shop, the sojourn time in the repair shop consists mostly of waiting time for resources such as specialists, tools and SRUs to become available. The amount of resources that are available in the repair shop determines the waiting times. For control reasons, these resource capacities need to be dimensioned in such a way that most repair jobs are completed within the agreed planned lead times. Within MLOs, internal repair lead time agreements are made (or targets are set) in consultation with the repair shop.

Decisions need to be taken on the amount of engineers and specialists to hire, the number of shifts and the number of tools of various types to acquire. In some instances, these tools are themselves major capital investments. The SRU stocking decision lies beyond the responsibility of the repair shop and is part of the total inventory control decision; see Section 2.2.6 for the reasoning behind this.

The resource capacity dimensioning decisions are based on the estimated repair workload, the repair workload variability (which follows from demand forecasting and parts returns forecasting) and the estimated repair time (and variability) required for an LRU when all resources are available. In making this decision, congestion effects need to be incorporated explicitly.

Since the costs of internal repair are mostly the result of the resources required for repair, the dimensioning decision together with the offered repair load can be combined to estimate the repair lead time and the cost of performing an internal repair. This information is used by supply management to periodically reconsider the LORA decision.

2.2.5.2 Schedule repair jobs

During operations LRUs are released to the repair shop and need to be repaired within the agreed planned lead time. This naturally leads to due-dates for repair jobs. The repair job scheduling function is to schedule the repair jobs subject to the resource constraints which are a consequence of the capacity dimensioning decision. Within these constraints, specific resources are assigned to specific repair jobs for specific periods in time so as to minimize the repair job tardiness. Additionally the repair shop may batch repair jobs to use resources more efficiently by reducing set-up time and costs associated with using certain resources.

2.2.6 Inventory control

The inventory control process is concerned with the decision which spare parts to stock, at which stocking location and in what quantities. Thus, inventory planning is done centrally for all locations (multi-echelon approach). This includes both LRU and SRU inventory locations. The inventory control process is visualized in Figure 2.7.

For control reasons, LRUs required for new projects and modifications are planned separately. We will not discuss the inventory control of these LRUs in detail here.

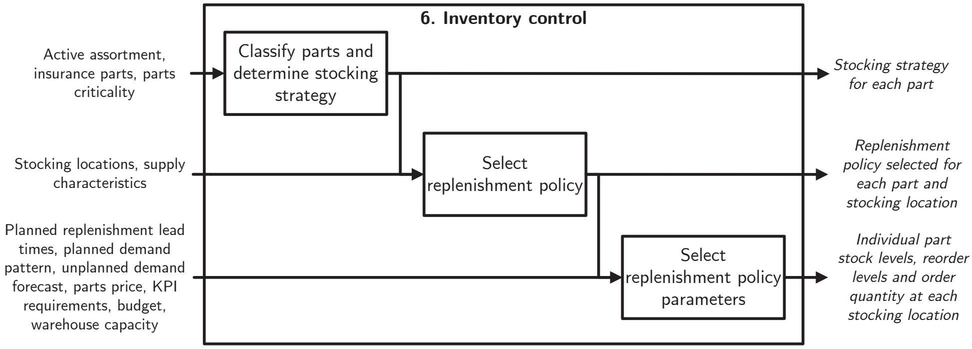

Figure 2.7 *Process overview of controlling inventories.*

2.2.6.1 Classify parts and determine stocking strategy

The MLO has several stocking strategies and classifies the spare parts assortment into different subsets: (i) (partially) critical spare parts and (ii) non-critical spare parts. Insurance parts are a specific subset of critical parts. The decision to stock insurance parts is not based on demand forecasts or on the contribution to a certain service level, but is based on other criteria such as supply availability, failure impact or initial versus future procurement price.

The availability of (partially) critical parts is needed to reduce system downtime. The stocking decision of (partially) critical spare parts depends on the contribution of a part to the overall service level of all (partially) critical parts. The availability of non-critical parts is needed for supporting an efficient flow of system maintenance, non-availability however does not cause immediate system downtime. Separate service level agreements are made for non-critical parts.

2.2.6.2 Select replenishment policy

The MLO is responsible for inventory replenishment of spare parts at all stocking points. For batching reasons the central warehouse replenishes the local stocking points only once during a fixed period (typically a couple of days or one week). This results in a (R, S)-policy for all parts at the local stocking locations (Silver et al., 1998). The length of the review period is set such that internal transport of parts is set up efficiently. In order to reduce system downtime costs, it may be beneficial to use emergency shipments from the central warehouse or lateral transshipments from other local stocking locations to deliver critical parts required at a local stocking point.

The MLO determines the timing and frequency of placing replenishment orders for the central stock based on supply characteristics. Spare parts for which framework-contracts are set up are usually delivered only once during a fixed period. Hence the stock level needs to be reviewed only once during this period, which results in a (R, S)-policy for these parts. The stock level of other parts is reviewed daily, resulting in an (R, s, S) or (R, s, Q)-policy for these parts (Silver et al., 1998).

2.2.6.3 Determine replenishment policy parameters

The MLO uses different methods to determine replenishment policy parameters for non-critical parts and (partially) critical parts. For (partially) critical parts one all encompassing model should be used to aim at a system (multi-item) service level. Optimizing policy parameters to satisfy a system service level is called the *system approach* (e.g. Sherbrooke, 2004; Muckstadt, 2005). In maintenance logistics, this is particularly useful, because MLOs are not interested in the service level of one part but in the multi-item service level instead.

The model should contain the following characteristics: (i) multi-echelon, (ii) multi-item, (iii) multi-indenture structure, (iv) emergency shipments from central depot, (v) lateral transshipments and (vi) multiple service level criteria. Input for this model are demand forecasts and information from supply management (supply lead times, parts prices), parts returns forecasts and information on the current inventory positions and replenishment policies of the spare parts. We note that some organizations also face budget constraints that need to be incorporated in the model.

2.2.7 Spare parts order handling

As discussed in Section 1.1.1, each MO requests the required spare parts. Spare parts are a resource that MOs share, i.e. the same spare parts can be requested by multiple MOs. Spare parts order handling is assigned centrally to the MLO and consists of the following steps: (i) accept, adjust or reject the order, (ii) release spare parts on the order and (iii) handle return order of failed repairable(s). For each of these steps, preconditions need to be defined as well as rules to manage these steps. A process overview of handling spare parts orders can be found in Figure 2.8.

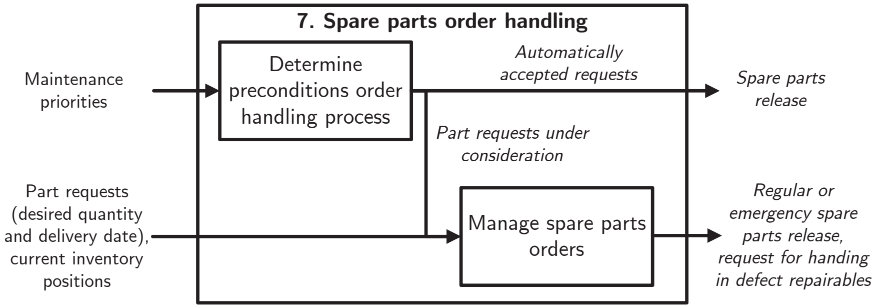

Figure 2.8 *Process of handling spare parts orders.*

2.2.7.1 Determine preconditions order handling process

The first decision in handling spare parts orders is to accept, adjust or reject the order. The advantage of checking spare parts orders is that unrealistic or unusual orders can be adjusted or, in case of incorrect orders, rejected. On the other hand, checking spare part orders is time consuming and increases the operational costs.

When checking spare part orders, the MLO obtains a trigger to contact the MO and adjust the order lead times and/or quantities. In this manner, MOs can reschedule certain tasks of their system maintenance work orders and adjust their spare parts orders based on the new system maintenance schedule. This might decrease system downtime (costs) caused by unavailability of spare parts.

Prioritization amongst spare parts orders while releasing spare parts is not easy in case the available stock is insufficient to meet all demand for that spare part. This is caused by the fact that the required spare parts are (i) part of a set of spare parts needed to start a maintenance task and (ii) are needed to start a different type of maintenance including different levels of criticality. Thus to fill orders, spare parts order handling faces an allocation problem similar to that found in assemble-to-order systems. The optimal solution to this problem is not generally known.

Once spare parts are released on a work order, the return process for failed repairables starts. For this purpose, the MLO creates a return order to hand in the failed repairable by the MO within the agreed hand in time.

2.2.7.2 Manage spare parts orders

Incoming spare parts orders are either automatically accepted or not, based on the preconditions set in the previous section. There might be several good reasons for unusual or unrealistic orders, hence there are no standard rules for accepting, adjusting or rejecting spare parts orders. This task lies with the MLO, who needs to consult with the MO on this.

2.2.8 Deployment

Deployment concerns the process of replenishing spare parts inventories. The deployment process consists of the following steps: (i) define preconditions order process and (ii) manage procurement and repair orders. A process overview of the deployment process can be found in Figure 2.9.

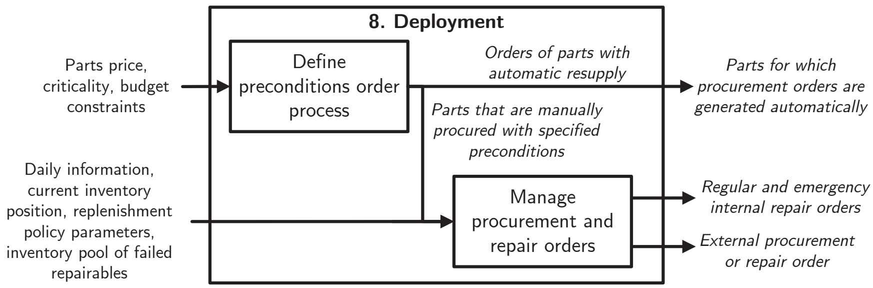

Figure 2.9 *Deployment process.*

2.2.8.1 Define preconditions order process

The replenishment policy parameters set by inventory control implicitly determine when to replenish spare parts inventories and what quantities to repair or procure. Deployment may deviate from this based on new (daily) information not known at the time the replenishment policy parameters were set, or when exceptional repair or procure orders arise from exceptional inventory levels. Deployment then starts a feedback loop to reconsider e.g. the demand forecast or supply lead times that led to this exceptional inventory level. Hence, deployment sets rules for exception management. The MLOs should set a precondition on whether to replenish inventories with or without interference of deployment.

2.2.8.2 Manage procurement and repair orders

The process of managing procurement and repair orders consists of the following steps: (i) procure or repair parts with the right quantity and priority, (ii) check the quality of the received spare parts and (iii) monitor supply lead times. The MLO needs to determine which quantity of each part to order and with what priority, for parts for which the procurement or repair order is checked upon release. The quantity deployment actually orders may deviate from the order quantity set by inventory control, based on newly obtained information. When an order is received, the MLO needs to check the quality of the received parts. When orders do not arrive within the agreed lead time, deployment takes necessary recourse actions.

2.3 Framework related literature

Each decision function in the framework can be supported by several methods and Table 2.1 contains several references that are a good starting point to investigate specific areas of literature in more depth and find models to support decisions that need to be made. The second column of Table 2.1 contains the processes in the framework. The third and fourth column contain the specific topic(s) related to the decision function and the relevant references.

Several review papers have been written on spare part management (Kennedy et al., 2002; Guide Jr. and Srivastava, 1997; Basten and Van Houtum, 2014). These papers provide an enumerative review of the state-of-the-art at the time of writing. In this chapter, we do not attempt to provide an exhaustive review of contributions on these subjects. Rather, we provide relevant references for each part of the decision framework to facilitate decision making.

2.4 Concluding remarks

In this chapter, we have presented a normative framework for planning & control of MLOs. With this framework, we have addressed the first research question:

1. *What are the relevant decision functions in MLOs, and how do they relate to each other?*

Table 2.1 Literature on different decision functions in the framework

	Process	Topic(s)	Literature
1.	Assortment Management	FME(C)A (failure mode effects (criticality) analysis)	Stamatis (1995) Ebeling (1997)
		Criticality, specificity, value	Huiskonen (2001)
		Commonality	Kranenburg and Van Houtum (2007)
		Reliability and quality	Öner et al. (2010)
2.	Demand forecasting	Overview	Altay and Litteral (2011)
		Time series analysis	Box and Jenkins (1970); Chatfield (2004)
		Croston methods	Croston (1972); Teunter and Duncan (2009)
		Life data analysis of equipment	Nelson (1982); Ebeling (1997)
		Prognostics	Heng et al. (2009)
		Linking forecasting to maintenance planning	Wang and Syntetos (2011), Hua et al. (2007)
3.	Parts returns forecasting	Scrap rates/Reliability engineering	Nelson (1982); Ebeling (1997)
4.	Supply management	Level Of Repair Analysis (LORA)	Basten et al. (2009, 2011), Alfredsson (1997); Barros and Riley (2001)
		Contract management	Van Weele (2010)
		Last buy	Van Kooten and Tan (2009), Bradley and Guerrero (2009), Teunter and Klein Haneveld (1998) Teunter and Fortuin (1999)
5.	Repair shop control	Capacity dimensioning and machine repairman models	Iglehart (1965); Borst et al. (2004), Chakravarthy and Agarwal (2003)
		Scheduling/Capacity assignment	Caggiano et al. (2006); Pinedo (2009)
		Overtime usage	Scudder (1985); Scudder and Chua (1987)
		Priority assignment to jobs	Scudder (1986); Guide Jr et al. (2000), Hausman and Scudder (1982)
6.	Inventory control	Review papers and books	Sherbrooke (2004); Muckstadt (2005), Kennedy et al. (2002), Guide Jr. and Srivastava (1997), Basten and Van Houtum (2014), Van Houtum and Kranenburg (2015)
		METRIC-type models (Multi-Echelon Technique for Recoverable Item Control)	Sherbrooke (1968, 1986), Muckstadt (1973); Graves (1985)
		Lateral transshipments	Lee (1987); Paterson et al. (2011), Kranenburg and Van Houtum (2009)
		Emergency procedures	Alfredsson and Verrijdt (1999), Verrijdt et al. (1998)
		Finite repair capacity	Sleptchenko et al. (2002); Caggiano et al. (2006), Díaz and Fu (1997); Adan et al. (2009)
		Empirical research and case studies	Cohen et al. (1997); Bailey and Helms (2007), Wagner and Lindemann (2008)
7.	Spare parts order handling	Allocation policies	Akçay and Xu (2004); Lu et al. (2010)
8.	Deployment literature	Behavioral aspects of planning	Fransoo and Wiers (2006, 2008), Wiers (2009)

To build up the framework, eight relevant decision processes have been identified. In each decision process, several decision functions have been defined. The decision functions in MLOs are related to supply chain management processes (e.g. forecasting and inventory control) and to repair shop control. For each of the decisions in the framework, we have provided literature to assist in decision making.

The framework is based on a literature review, and is especially useful for practitioners. In the next chapter, its completeness, applicability and value is verified (demonstrated) in case studies at six prominent Dutch companies that use and maintain high-value capital assets. With the case studies, we also aim to identify open topics for the remainder of the research.

Chapter 3

Case studies: part I

> *"De waarheid is nooit precies zoals je denkt dat hij zou zijn."*
>
> *Johan Cruijff*

3.1 Introduction

The framework in Chapter 2 serves as a reference model for decision making in MLOs. In the framework, we have grouped *decision functions* (small blocks) into *processes* (large blocks), see Figure 2.1. That is, we have grouped coherent sets of decision functions based on their function. For example, the decisions functions *'classify parts with respect to demand forecasting'* and *'characterize demand process'* are functionally grouped into the process *'demand forecasting'*.

In the framework, we describe what decisions are to be made by MLOs and in what coherence. Further, interfaces and information flows between decision functions are described. Our research has started from the observation that in practice we observe solutions/approaches to decision making in MLOs that we do not understand, i.e. that deviate from the framework in Chapter 2. We therefore aim to get a better understanding of decision making in MLOs in practice and its differences with the framework, and formulate the following research questions:

2a. Which processes/decision functions in the framework are observed in practice?

2b. Which processes/decision functions observed in practice are not included in the framework?

2c. Which (interfaces and information flows between) processes/decision functions are filled in differently by the case study companies, and why?

To answer these questions, we use the framework in in-depth case studies at six MLOs. Case studies are a suitable method to extensively describe and understand the differences between the framework and practice. At the same time they enable us to find new factors or variables that influence the choices for the coherence, interfaces and information flows between decision functions, or to find new decision functions not yet included in the framework. The main findings from the case studies form the basis for the research topics of the next chapters.

This chapter is organized as follows. In Section 3.2 we describe the methodology used to collect the empirical data. Section 3.3 describes, analyzes and interprets the main case study findings, and Section 3.4 concludes this chapter.

3.2 Methodology

Case studies are a suitable method for descriptive and exploratory research on phenomena that are not well understood (Yin (2003)). We have selected the case study companies before we actually started the research (and before we developed the framework of Chapter 2). Prior to the research, we have asked multiple customers (companies) with which Gordian has a relationship whether they are willing to participate in the research. From the companies that expressed their interest, we have selected six companies with the same type of service network, i.e. we have selected companies that both use and maintain high value capital assets. In this section, we introduce the case study companies, we outline the research protocol and the method to collect and analyze the data.

Introduction of the case study companies

The case study companies both use and maintain high value capital assets and operate in different industries. The Dutch Railways (NS) is a passenger railway operator in the Netherlands, and its MLO is specialized in the maintenance, service, cleaning and overhaul of rolling stock. KLM Engineering & Maintenance (KLM E&M) is an aircraft maintenance organization, and carries out maintenance, repairs and modifications on aircraft, engines, and components for their fleet and various other airlines worldwide. Europe Container Terminals (ECT) is a container terminal operator that uses and maintains container handling systems. The Royal Netherlands Airforce (RNLAF), Royal Netherlands Army (RNLA) and Royal Netherlands Navy (RNLN) are armed forces service that contribute to peace and security on a global basis. For this purpose, they use and maintain military equipment such as aircraft and helicopters (RNLAF), tanks and vehicles (RNLA) and frigates and ships (RNLN). Table 3.1 presents general figures about the asset base that the companies use and maintain[1].

Table 3.1 General figures about the asset base of the case study companies

Company	**Scope**	**Number of assets**	**Number of asset types**	**Minimum number per type**	**Maximum number per type**
NS	Carriages	2,876	11	15	868
KLM E&M	Aircraft	211	19	1	41
ECT	Quay cranes	39	7	2	8
RNLAF	Aircraft, helicopters	171	8	4	87
RNLA	Armoured vehicles	1,241	13	4	407
RNLN	Frigates, ships	35	12	1	6

The MLO within NS fits the description of an MLO in Section 1.1. This also holds for RNLAF and RNLA but they have an additional complexity: a yearly budget constraint for replenishment of spare parts. The MLO within RNLN also faces this budget constraint. Another difference at RNLN is that the maintenance depots of the Maintenance Organization (MO) share repair capacity with the repair shops of the MLO, and this impacts decision making related to repair shop control.

KLM E&M maintains airplanes that are used worldwide, and Original Equipment Manufacturers (OEMs) of these airplanes have certified repair shops of airline companies to conduct repairs for other airline companies. KLM E&M has for example contracts with external airline companies on the availability of repairable parts (rotables). For KLM E&M, on the one hand this means that they (can) have a larger repair volume, but also have to serve multiple customers with different service levels. On the other hand, external repair shops can be both supplier and competitor

[1] For each of these equipment types there exist several equipment *subtypes*, but this distinction is left out of the figures. For ECT, figures are restricted to quay cranes only. For RNLA, the figures are restricted to armoured vehicle types only.

of KLM E&M and this may influence repair lead times.

ECT deviates from the other companies because they have no internal repair shops. The assets consist of only limited technology that is specifically designed for ECT, and there are relatively many external suppliers or repair shops available that maintain this technology (and related parts) within relatively short repair lead times. Therefore, it is economically not beneficial to repair parts internally.

Research protocol and interview guide

Multiple case studies in different environments increase the *external validity* (Yin, 2003), i.e. reduce the effects of exogenous characteristics when comparing the framework to practice. We have visited each case study company for 12 to 16 working days and for each of them we have written a final case study report of about 30 pages[2]. These reports are the source for the findings that are presented in Section 3.3. We have prepared a template for the reports and a questionnaire prior to the case studies (see Appendix 3.A.1), which have served as guidelines for conducting the case studies.

Data collection

Our main data source consists of interviews with 8-18 key informants of each company. In total we have conducted 79 interviews (see Appendix 3.A.2 for the list of the positions of the key informants and the discussed topics). The key informants have different positions and work at different levels in the organization (both managerial and operational levels), in order to maximize the diversity among informants and to reveal a variety of perspectives. Examples of the key informant roles include: supply chain or inventory manager, supply chain or inventory planner, repair shop manager, repair shop planner, purchase manager, maintenance manager and maintenance engineer. The number of key informants depends on the size of the company and the extent to which decision making is divided among different positions in the company. We have selected the key informants in consultation with our main point of contact, in general a (supply chain or maintenance) manager with a clear view on the informant's knowledge and role within the company.

With each of the key informants, we have discussed multiple processes and decision functions of the framework. We have selected the discussion topics in consultation with our main point of contact, such that we have discussed each process and decision

[2]These reports are available from the author upon request. We here refer to the following case study reports: Driessen and Arts (2011); Driessen (2011, 2012a,b,c, 2013).

function with at least three key informants of the company, in order to maximize triangulation between them. The interviews were semi-structured (based on the questionnaire) and have lasted between 45 and 90 minutes. Prior to the interviews, a list of questions (a subset of questions from the questionnaire, based on the selected discussion topics) is sent to the key informants so that they could prepare themselves for the meeting. We have sent summaries of the interviews to the key informants (*internal validity*), and have completed them based on their individual feedback. In case of feedback contradictions, we have used source triangulation through additional interviews with other key informants with which we have discussed the same topics (*construct validity*). In addition, we have used results from quantitative analysis and other sources (e.g. internal documents, websites) to further triangulate observations and motivations given by the key informants.

Based on the qualitative and quantitative data from interviews and data analysis, we have made a first version of the case study report. We have sent this version to all key informants for feedback, which we have used to make a second version of the report. We have discussed last feedback in a meeting with all key informants and, based on consensus among them, we have made a final version of the case study report.

Data analysis

Each report consists of six sections. The first two sections contain an introduction and company background. The third section provides a detailed description of the environment in which the company operates, and highlights the company specific characteristics related to the demand and supply of spare parts, and organizational, IT and regulation related aspects. A good understanding of these characteristics is required to carefully interpret the differences between the framework and practice.

The fourth section contains the description of the differences between the framework and practice. For each decision in the framework, we describe how often it is made, based on what information and its coherence with other decision functions. If possible, differences between the framework and practice are clarified by the motivation of the responsible decision maker(s) or by company specific characteristics. The fifth section contains results on some important Key Performance Indicators (KPIs) of Maintenance Logistics Organizations (MLOs), such as stock availability and stock value. These results give an indication of the maturity level of each case study company, and are used to carefully interpret differences between the case study companies. The last section contains the list of key informants that is interviewed.

3.3 Case study findings & interpretation

Table 3.2 gives an overview of the processes and decision functions that are observed in the case study companies (marked with a 'X'), and answers research question 2a. We have found each process, except for parts returns forecasting and repair shop control, in all case study companies. The fact that we have not observed the process *repair shop control* at ECT is explained by the absence of internal repair shops. We have not observed the process *parts returns forecasting* at ECT and RNLA because information is not easily available and the key informants have mentioned that this process only merely impacts the availability and costs of spare parts.

In the remainder of this section, we describe the main case study findings and give answers to research question *2c* for each process and decision function. We furthermore describe the considerations that, besides economic ones, also play a role in decision making. Research question *2b* is answered in Section 3.4.

Assortment management

Both decisions related to assortment management in the framework are found in each of the case study companies. We have observed that most parts in the configuration of an asset is added to the assortment at some point during the asset life cycle rather than at the beginning of it. At RNLAF, RNLA and RNLN, we have found that only 15-25% of the spare parts is added to the assortment within two years after phase-in of the asset type[3]. The explanation for this low value is the observation that OEMs not always supply the Recommended Spare Parts List (RSPL) to the MLO (e.g. OEMs buy off this obligation to protect their intellectual property); if the RSPL is supplied then often it includes only the highest indenture level (critical) parts, and hence non-critical and lower indenture level parts should be identified by the MLOs themselves.

We have observed that especially larger sized (and often repairable) Line Replaceable Units (LRUs) and Shop Replaceable Units (SRUs) are added to the assortment shortly after phase-in of a new asset type. They are (often) more critical and because of their larger size, it is easier to identify them as parts that potentially require maintenance (and hence may require spare parts). The decision to add (smaller sized) SRUs to the assortment is often not taken in the beginning of the life cycle of an asset. In

[3] These numbers are based on the analysis of parts in the configuration of five asset types at these three MLOs.

Table 3.2 Overview of processes and decision functions observed in the case studies.

Process	Decision function	NS	KLM E&M	ECT	RNL AF	RN LA	RN LN
Assortment management	Define spare parts assortment	X	X	X	X	X	X
	Gather parts (technical) information	X	X	X	X	X	X
Demand forecasting	Classify parts with respect to demand forecasting		X	X	X		
	Characterize demand process	X	X	X	X	X	X
Part returns forecasting	Classify parts with respect to returns forecasting	X	X				
	Forecast part return rates and return time	X	X		X		X
Supply management	Manage supplier availability and characteristics	X	X	X	X	X	X
	Control supply lead time and supply parameters	X	X	X	X	X	X
	Select supply source(s) and contract(s)	X	X	X	X	X	X
Repair shop control	Determine repair shop resource capacities	X	X		X	X	X
	Schedule repair jobs	X	X		X	X	X
Inventory control	Classify parts and determine stocking strategy	X	X	X	X	X	X
	Select replenishment policy	X	X	X	X	X	X
	Determine replenishment policy parameters	X	X	X	X	X	X
Spare parts order handling	Define preconditions order handling process		X		X	X	X
	Manage spare parts orders	X	X	X	X	X	X
Deployment	Define preconditions order process			X			X
	Manage procurement and repair orders	X	X	X	X	X	X

this period the asset is under warranty and the OEM is responsible to repair failed parts and to supply the required SRUs. After the warranty period has expired and the decision to maintain parts at internal repair shops is made, SRUs are added to the spare parts assortment. We have observed that often these SRUs are added to the assortment shortly after their first failure, because it is often hard to identify which SRUs may potentially fail, or there is little time available to thoroughly analyze them due to delays in the final decision to internally repair the parts. In the case study at KLM E&M, we have found that 20% of SRU backorders are due to SRUs that were not yet registered in the spare parts assortment.

Two of our case study companies (NS, KLM E&M) distinguish three parts types. Besides *repairables* and *consumables* also the part type *rotables* is used. For the definition of rotables at the NS, we refer to Arts (2013). The definition of rotables at KLM E&M is comparable to our definition of repairables (see Section 1.1.1), but it differs in the fact that rotables can be repaired or maintained infinitely many times after replacement, whereas repairables are condemned after a while. This distinction is not mentioned in the framework of Chapter 2.

Demand forecasting

Demand forecasting is a process we have observed in all case study companies. However, not all of them distinguish parts with and without advance demand information (ADI). This is related to the effort required to gather the ADI and its reliability, compared to its benefits. The ADI is not always reliable or subject to fluctuations: the timing of usage based maintenance depends on the usage of the asset, that may change over time due to changes in e.g. the operating schedule or the (un)availability of other assets. The ADI is mainly available for parts that are replaced based on usage or time; the availability of ADI for parts that are replaced based on their condition is limited, as information from inspections is not registered appropriately and remote monitoring techniques are not yet applied on a large scale.

For a small group of more expensive (repairable) parts under a usage based maintenance concept, such as engines, some companies (RNLAF, RNLN) use the ADI to synchronize the operating schedule and maintenance planning of the assets. By carefully monitoring the remaining number of running hours/cycles of an individual part (e.g. engine) in an asset, they try to adjust the asset operating schedule such that replacement of these parts is synchronized with the availability of spare parts. If it concerns repairable parts, the maintenance planning is sometimes also synchronized with the resource planning of the internal repair shops. In this sense, the demand forecast for spare parts can somehow be influenced by the MLOs, in consultation with the MO, and this aspect is not mentioned in the framework of Chapter 2.

For parts under a condition based maintenance concept, information on planned inspections (to determine *timing* of demand) is sometimes combined with a list of *x%-parts*[4] (to determine *content* of demand) into a demand forecast for spare parts (e.g. RNLAF, NS). This forecast is often compared to a forecast that is made based on time-series, and the best forecast method is selected based on intuition. All companies use time-series to forecast demand for spare parts. This is not surprising as most requests for spare parts are unplanned (see Chapter 1) and other sources to forecast demand, e.g. mean-time-between-failures (MTBF) of parts, are often not available for the MLOs.

All companies distinguish *fast* and *slow* movers; a fast (slow) mover has a relatively high (low) number of historical data points (frequency of demand in the past). According to our key informants, such time-series based forecasting techniques may be suitable for fast movers but not for slow movers, because the forecast is not reliable. In fact, for slow movers, often only implicitly (based on intuition) a forecast is made when inventory managers manually determine the desired replenishment policy parameters (see process inventory control). The observation of the key informants is supported by the quantitative analysis that we conducted. The analysis shows that only 1-5% of all spare parts were used more than 12 times per year (on average over the past two years), and for 40-60% of this group of spare parts the demand coefficient of variation is larger than 1. For such intermittent and lumpy demand patterns, in general time-series based forecasting techniques are less suitable.

The spare parts demand forecasts are made by the department(s) responsible for inventory management of the spare parts. The extent to which these forecasts are shared within the MLO (e.g. internal repair shops and the purchase department) and with external suppliers is limited. Only in three companies the demand forecasts are shared with the internal repair shops in order to forecast the repair shop workload (NS, RNLAF, RNLA). Within KLM E&M, key informants have mentioned that it is more appropriate to forecast workload at the repair shop level because repair shops also face external demand (which is not included in the forecasts of the inventory department). We have observed that the purchase departments make their own forecast to determine which supply sources and contracts should be updated and when, and internal repair

[4] See Section 2.2.2.2

shops make their own forecast to determine the expected workload for each of its resources. These separate forecasts may be based on different assumptions, and hence may lead to future mismatches in the demand and supply of spare parts. Furthermore, although there is a direct link between the demand for LRUs and SRUs, we have not observed that demand forecasts for SRUs are based on the demand forecast of higher indenture level LRUs.

Parts returns forecasting

Parts returns forecasting concerns forecasting the number of ready-for-use (RFU) parts returns, condemnation and no fault found rates. We have observed that none of the case study companies forecast no fault found rates. This is mainly because data that can be used for this purpose is not easily available, and its impact on decision making is limited according to our key informants. All case study companies, except ECT, register condemnation of (repairable) parts in their information system. This information is used to determine the fraction of parts for which a repair job was started, but stopped due to condemnation of the part that was to be repaired. This fraction is used as a forecast for the condemnation rate. For ECT, restricted functionality of their information systems prevent ECT from registering the condemnation of parts.

NS uses time-series based forecast models to forecast the number of returned RFU parts, based on historical data on the number of RFU parts returns. The other case study companies only implicitly forecast RFU parts returns, by subtracting the actual number of returns from the actual demand data that are input for the time-series based forecast models (hence interpreting returns as negative demand). The impact of not explicitly forecasting RFU parts returns is probably limited, as the percentage of RFU parts returns is estimated at 2-5% in our case study companies.

Supply management

Supply management appears to be a time-consuming process in the case study companies, which can be explained by the fact that their assortments comprise tens up to hundreds of thousands spare parts. Moreover, quotations from suppliers on the price and supply lead time of spare parts are often valid only for one year. Especially for parts with a low demand rate, quotations need to be updated frequently between two subsequent resupply orders. Therefore the availability of a supply source is managed only for spare parts that are critical and/or require resupply in the near future. KLM E&M and ECT are the exception; for 80-90% of the assortment resupply is directly possible. This can be explained by the fact that often multiple suppliers are available for the same part, whereas often only one supplier is available for most parts in the other companies.

In contrast to the framework, four case study companies (NS, RNLAF, RNLA, RNLN) have reasons not to assume planned (deterministic) lead times for the external resupply of spare parts. We have observed that their external supply lead times are long and variable, and the percentage of resupply order lines in which the ordered quantity was received on time (i.e. on the requested date) varies between 45 and 70% per company. This observation can be explained by the fact that these companies have little buying power compared to other customers of their suppliers (see also Section 1.1.2), and therefore often only agreements are made on parts price and not on the supply lead time of spare parts. These four companies therefore take into account the supply lead time variability in setting replenishment policy parameters. In contrast, KLM E&M and ECT have more buying power and do make lead time agreements with their suppliers.

At the tactical level, only within KLM E&M agreements are made on repair lead times between inventory control and the internal repair shops. We have not observed a clear differentiation of lead times among the repairable parts. This is motivated by the fact that their repair lead times are aligned with market standards, because they also repair parts for external suppliers. ECT has no internal repair shops, the remaining four companies have a different set of agreements on the interface between inventory control and the repair shops. Therefore at the tactical level no agreements are made on repair lead times.

Supply sources are selected based on criteria such as the price, supply lead time and quality (i.e. failure rate) of the parts. Next to this, legal requirements (e.g. suppliers or external repair shops should be certified by the OEM) and capabilities of the internal repair shops play a role in the selection process. The choice for internal repair depends, besides the aforementioned considerations, also on social factors such as maintaining employment at the internal repair shops, and on the strategic choice of the company to have or maintain their (high) knowledge level about the asset technologies (to be a ‘smart buyer’ of new assets). We have observed that the choice to internally repair (a group of) parts is usually made only when a new asset type enters the fleet, and this choice is only reconsidered on part level.

Repair shop control

In the framework, it is assumed that there is a single repair shop that is responsible for repairing failed parts. In practice, we have observed that all case study companies have multiple repair shops. These repair shops operate independently and have a functional layout. The repair shop capacity is rather inflexible. The repair shops often have long-term contracts with repair engineers that they don't want to stop (for financial as well as social reasons). Recruiting new repair engineers has a long effectuation time due to extensive training programs, and the possibility to hire (lease) service equipment from (to) third parties is limited.

Within KLM E&M and NS, repair jobs are automatically released upon receipt of an non-ready-for-use (NRFU) part. No order acceptance function is observed, that is e.g. based on workload constraints set at the tactical level. The repair shops are responsible for subcontracting repair work, or applying overtime in case of a temporarily high workload. Repair work is scheduled based on earliest due date (KLM E&M) or dynamic priorities (NS).

Within RNLAF, RNLA and RNLN, order acceptance takes place at the operational level. Each repair job is released by deployment, which sets the repair lead time of each repair job in consultation with the repair shop. Deployment is responsible for subcontracting repair work and setting repair job priorities, which are used by the repair shops to schedule their repair work.

Subcontracting repair work usually requires a set-up time that is relatively long (due to e.g. transportation times, negotiations about repair prices) compared to internal repair lead times (e.g. 2-12 months vs. 2-12 weeks). Thus subcontracting is relatively inflexible and not convenient for repair shop capacity management purposes. Overtime is used as a capacity management tool rather than subcontracting.

Inventory control

In the case studies, we have observed that some companies (RNLAF, RNLA, RNLN) have assigned the responsibility to control the SRUs and LRUs to the same department. This choice is made for efficiency reasons, i.e. the same type of work is executed by the same department. KLM E&M and NS have assigned the control of SRUs to the repair shops, in order to align control with responsibilities.

Although most companies (NS, KLM E&M, RNLAF, RNLN) have a distinct (group of) spare parts planner(s) for the parts related to a certain asset type, we have observed that the inventory control parameters are not set using a *system approach*.

All companies do apply an ABC-classification in which service levels are differentiated for parts with different characteristics (price and demand frequency). This enables them to obtain the same effect as a *system approach* in a single-location, single-indenture setting. We have observed that all companies, except NS, manually set the replenishment policy parameters for almost all of their slow- and non-moving parts.

Within KLM E&M the initial stocking decision for repairable parts is sometimes time-phased, i.e. the initial stocking decision is not only made at the beginning of but also at some point during the asset life cycle. This approach is applied to parts with a high MTBF which are replaced based on condition, and its first demand is to be expected only in the long run. It deviates from most theoretical models, that typically assume steady-state demand. It is applicable because the increase of the new buy price after the initial buy is relatively low. Moreover, often KLM E&M has an alternative to source the parts in case needed (e.g. at the second hand market). In the other companies, our key informants have mentioned that new buy prices of repairable parts significantly increase after the initial buy, hence making time-phased stocking less attractive for them.

RNLAF, RNLA and RNLN have a limited yearly budget for the resupply of spare parts. This restriction is not taken into account when setting replenishment policy parameters. In fact, deployment takes this restriction into account when creating resupply orders.

Interface between inventory control and repair shop control

In the framework, we have made two design choices regarding the interface between inventory control and repair shop control. First, we assume that planning and control of SRUs and LRUs is integrated and second, we assume that inventory control makes agreements on the tactical level with the repair shops about planned repair lead times. The repair shops dimension the required repair capacity and apply proper scheduling rules, such that these planned lead times are met with a high probability and repair costs are minimized. Based on the planned repair lead times, inventory control is able to minimize the total inventory costs of SRUs and LRUs, while ensuring a certain availability of the LRUs.

In the case studies we have observed that decisions related to inventory control (planning and control of SRUs and LRUs) and repair shop control (repair capacity dimensioning and scheduling of repair jobs) are separated over two departments (decision makers), each with their own responsibilities and targets. We have observed that the design of the 'interface' in our framework leads to the following two situations, in which control and responsibilities are not aligned. First, to timely meet the planned lead times, the repair shop depends on the availability of the SRUs. This requires that either the SRUs are controlled by the repair shop, or that all SRUs are available when the repair job is released to the repair shop. The latter is difficult we have observed in the case studies, because often an inspection of the part (and hence repair capacity) is required to determine which SRUs are required in the repair process. Second, from the perspective of the inventory department it is sometimes desired to deviate from the planned lead times, by changing repair job priorities, in order to prevent or resolve backorders for LRUs. This may however be inefficient from a repair point of view, for example when repair jobs need to be interrupted.

Each case study company applies a "one-size fits all" approach: the interface between the two departments is the same for all repair shops of the company. We have however observed that the way the decisions are separated over the departments differs among the case study companies, and that different 'interfaces' between the two departments exist. In all case studies, the repair shop is responsible for dimensioning the repair capacity and the inventory department is responsible for LRU inventory control. In contrast, SRU inventory control and scheduling of repair jobs (i.e. priority setting and capacity allocation) are separated differently, and a different interface agreement is applied. Table 3.3 summarizes (the differences between) the three types of 'interfaces' that are found in the case studies. Hereafter, a more detailed description is given.

Table 3.3 Different 'interfaces' found in the case studies.

Company	**Inventory control SRUs**	**Priority setting**	**Capacity allocation**	**Interface**
KLM E&M	Repair shop	Repair shop	Repair shop	Lead time agreements at the tactical level
NS	Repair shop	Inventory manager	Repair shop	Dynamic repair priorities at the tactical level
RNLAF, RNLN, RNLA	Inventory manager	Inventory manager	Repair shop	Lead time agreements at the operational level

Within KLM E&M, at the tactical level agreements are made on the planned lead times by balancing repair (capacity) and inventory holding costs. The planned lead time starts upon receipt of a returned NRFU part and ends when the part enters the RFU stock. To meet these lead times, the repair shop determines the repair job priorities, how to allocate available repair capacity to the repair jobs, whether and when working overtime is applied and when repair jobs are subcontracted. In order to align control with responsibilities, the repair shop plans and controls the SRUs. The repair shop is responsible to meet the planned lead times with a high probability. The inventory department makes, to some extent, use of the possibility to negotiate the repair job priorities with the repair shop, i.e. by changing individual repair job lead times. This means that lead times are decreased for urgent repair jobs while lead times are increased for less urgent repair jobs.

Within NS, at the tactical level dynamic repair priorities for each LRU are set by the inventory department by means of 'min-max levels': LRUs with backorders have the highest repair priority (priority 1), LRUs with an RFU stock level below min-level have second highest priority (priority 2), LRUs with an RFU stock level between min and max level have priority 3, and all other LRUs have priority 4. Based on these repair priorities, the repair shop creates repair jobs for NRFU parts and allocates capacities to the repair jobs. The LRU base stock levels are dimensioned assuming planned lead times of four weeks for all LRUs. The repair shops have targets on the fraction of repair jobs that are executed with priority 1 or 2, as this number indicates whether the repair shop is 'in control' and whether and when overtime should be applied. The repair shop plans and controls the SRUs to align control with responsibilities.

Within RNLAF, RNLA and RNLN, at the operational level lead time agreements are made. The repair priorities are set by the inventory department and negotiation takes place with the repair shop on the due date of each individual repair job. The repair shop is responsible to meet these agreed lead times with a high probability and to do so, the repair shop is responsible for allocation of repair capacity to the repair jobs, and to decide whether and when overtime is applied. The inventory department is responsible for planning and control of SRUs. If repair job due dates cannot be met because of missing SRUs, the repair shop and the inventory department negotiate a new realistic due date for the repair job. Within RNLAF only, at the tactical level agreements are made on the workload offered to the repair shops. Based on this agreement, the repair shops either accept or reject repair jobs, and the inventory department decides on whether the repair job should be subcontracted to external repair shops.

In the case studies at KLM E&M, RNLAF, RNLA and RNLN, we have observed that actual repair lead times strongly vary and repair shops are able to complete only a relatively low fraction of repair jobs within the agreed lead time (60-70%). At the same time, in the same company we have observed that some of the repair shops successfully collaborated with the inventory department (i.e. few interruptions of repair jobs and rescheduling actions occur), whereas we have observed serious lacks of consistency between the inventory department and other repair shops (i.e. many interruptions of repair jobs and rescheduling actions occur). This observation indicates that a "one-size-fits-all" approach is not always effective and, when it comes to the coordination of inventory and repair shop control, shows that refinements of our framework may be necessary in some situations. How to refine the framework is the topic of the next chapters.

Spare parts order handling

Spare parts order handling concerns the acceptance and management of orders for spare parts from the maintenance depots. We have observed that the case study companies do not set preconditions for checking spare parts orders based on economic considerations. Nevertheless, except for KLM E&M and ECT, we have observed that there are other reasons to set preconditions on the acceptance of spare parts orders. Examples include parts with special storage, environmental or safety instructions, parts under technical inspection, and parts that are "nice to take home". Besides these exceptions, spare parts are immediately released according to a 'first-come, first-served' (FCFS) principle.

In case of spare parts shortages, we have found a feedback loop from the MLO to the MO in all case study companies. Negotiation takes place in order to change the maintenance priorities and the due dates of the related spare parts orders. This however occurs occasionally, which can be explained by the fact that parts commonality is low, and often availability of common parts is high. KLM E&M is an exception as it serves not only its own fleet, but also external customers. An internal report shows that FCFS release of spare parts is inefficient; the service level turns out to be 3% lower compared to the best (analyzed) policy to release spare parts to their customers.

When repairable spare parts are released, the MO has the obligation to return the NRFU part within an agreed hand-in time. In most theoretical models, this return time is often assumed to be zero. Except for RNLAF and RNLN, the number of parts that need to be handed in by the MOs, is not clear as appropriate registration is lacking. Hence there is not a real "closed-loop", and this complicates the resupply process for (repairable) spare parts.

Deployment
Each company has a decision support system that creates alerts (replenishment advices) when stock levels are too low, based on the replenishment policy parameters set by inventory control. These alerts are checked by deployment and if required, a resupply order is created. No formal preconditions for automatically processing these alerts, i.e. creating replenishment orders, have been observed in our case study companies. Nevertheless, RNLN has a non-formal policy that states that the Pareto-principle should be applied, i.e. that 80% of the replenishment advices (20% of replenishment value) should be accepted automatically. It is however not clear whether this policy is applied on a large scale within RNLN. We have observed that all companies have authorization levels (based on purchase value) for finally placing replenishment orders at external suppliers.

The replenishment process is different for repairable spare parts, as alerts to replenish the stock of repairable parts are ignored by deployment in all companies. This is motivated by the fact that the number of parts that need to be handed in by the MOs is not always clear in practice, or that this number is not taken into account when creating alerts (RNLAF, RNLN). In fact, repair priorities are set for available NRFU parts only when or shortly before repair capacity is available. This approach prevents a lot of rescheduling of internal repair job due dates, especially when agreed repair lead times are long, due to shifting repair priorities caused by temporary peaks in the demand for spare parts.

In general, deployment relates to the replenishment of RFU stock. However, KLM E&M also has the possibility to loan parts from (or to) other airline companies, which can be seen as a temporarily supply source, and KLM E&M and RNLAF both sell surplus stock to the second hand market. In these cases, deployment also tries to reduce stock levels. The latter is not observed in the other companies. The main motivation given by the key informants is the fact that commonality of parts with other companies is lower.

3.4 Conclusions

In this chapter, we have used the framework of Chapter 2 in six case studies to identify differences in decision making between the framework and practice, by answering the second research question:

2a. Which processes/decision functions in the framework are observed in practice?

2b. Which processes/decision functions observed in practice are not included in the framework?

2c. Which (interfaces and information flows between) processes/decision functions are filled in differently by the case study companies, and why?

From the case studies we have observed the following[5]: (1) each process in the framework, except the process ‘parts returns forecasting’, is found in every case study company and (2) each decision function in the framework is found in at least two of the case study companies, and 13 out of 18 decision functions have been found in all case studies. These observations answer research question *2a*, and demonstrate the applicability of our framework in different environments.

We have not found a generic decision function that is not included in the framework, and this answers research question *2b*. As our case study research includes a large number of interviews (in which we have asked for additional decision functions not yet included in the framework), we may therefore suspect that the framework as presented in Chapter 2 is complete. As such, the framework may also serve as a useful starting point in making specific designs of existing spare part planning and control systems.

The case studies have revealed some interesting findings, and the main finding per process can be found in Table 3.4. The main finding of the case studies relates to the coordination of decisions in the processes supply management, repair shop control and inventory control, and has revealed that some refinements of the framework may be necessary in some situations.

In the framework, we assume that the repair shops are able to meet the agreed lead times with a high probability. This implicitly means that the repair shops are able to carefully control the repair workload and that SRUs do not cause delays of repair jobs (i.e. SRUs are either not required, or are available before the start of a repair job). To meet the agreed lead times turns out to be complicated in practice, and we have found indications that a “one-size-fits-all” control concept for all repair shops is not optimal.

The refinements of the framework relate to the interface between inventory control and repair shop control, and are the topic of the next chapters. In the remainder,

[5]Here we do not account for the fact that the process ‘repair shop control’ and its related two decision functions are not observed at ECT, because ECT does not have internal repair shops.

Table 3.4 Main case study findings of each process.

Process	Main findings
Assortment management	The indenture structure of repairable parts is not always exactly known, and the ‘unknown’ SRUs cause many delays of repairs. In contrast, in the literature it is often assumed that the indenture structure is exactly known.
Demand forecasting	Time series based forecasts are often judged ‘unreliable’ by the key informants, especially for (very) slow moving parts. For these parts, often the demand forecast and inventory control decisions are not made sequentially, but are integrated.
Returns forecasting	The returns forecasting process turned out not be very important for the case study companies. Often forecasts based on averages are sufficient.
Supply management	Many case study companies have reasons not to assume planned (deterministic) lead times for the resupply of spare parts, due to the low on-time delivery performance of both internal repair shops and external suppliers.
Repair shop control	All case study companies (except ECT) have multiple repair shops, instead of one which was assumed in the framework. These repair shops operate independently and have a functional layout. The repair shop capacity is rather inflexible.
Inventory control	Some case study companies have assigned the planning and control of SRUs to the repair shops, in order to align control with responsibilities, whereas it is assigned to the inventory control department in the other companies.
Spare parts order handling	No preconditions based on economic factors for accepting spare parts order were observed. Spare parts are immediately released according to an FCFS-principle.
Deployment	No formal preconditions for automatically processing resupply orders for spare parts were observed.

we concentrate on designing control concepts for repairable parts inventory systems (see Section 1.2). In Chapter 4 we start with a literature review, and compare the literature to additional case studies on repairable parts inventory systems in Chapter 5.

3.A Appendices

3.A.1 Template report & questionnaire

This appendix provides the template that is used for the case study reports. This template has also served as a questionnaire for the interviews with the key informants. Based on the selected discussion topics, we have shared the relevant questions with the key informant prior to the interview. We do not pretend that the list of questions is complete. During the interviews additional questions came up and not all of them (especially company specific questions) are added to the questionnaire. Besides, other information shared by the key informants has been used as input for the case study report.

The case study reports consist of six sections: I-VI. Section I contains a general introduction to the case study. Section II and III of the report provide a characterization of the environment in which the MLO operates. It consists of the following subsections, and includes answers to the following questions (see Table 3.5):

	Topic	Questions
	Background information	What is the size of the company (in FTE)? What customers does the MLO serve? What is the turnover value of the MLO? What is the size of the installed base? Which (maintenance) services does the MLO provide? Which (maintenance) services are subcontracted by the MLO? How is the organizational structure set up? What does the service network look like?
1.	Installed base	Which asset types are distinguished? What is number of assets per type? In which part of the life cycle are the assets currently? What are the main characteristics of the asset base?
2.	Characterization of system maintenance	Which maintenance strategy (or concept) is applied? Who determines the maintenance concept? Which types of maintenance are distinguished? What are drivers for conducting maintenance? Who is responsible for conducting (the various types of) maintenance? What is the ratio between the various maintenance types?
3.	Spare parts typology	Which spare parts types are distinguished? What is the number of spare parts per type? What is the number of parts in the configuration of all assets?

	Topic	Questions
4.	Logistics concept	
4.1	Spare parts demand process	How is demand for spare parts being generated? Who is responsible for creating and handling spare parts orders? How are spare parts requested and how is the request registered? What does the procedure for handling spare parts requests (including issuing of spare parts) look like? How are returns of (non-)ready-for-use parts handled?
4.2	Spare parts supply process	What does the spare parts network look like (in terms of stock locations)? What type of supply sources are distinguished? How is the internal logistics process (i.e. receipt, storage and issue of parts) organized? What transportation modes (and emergency shipments) are distinguished? What are the (general) shipment times of each transportation mode?
4.3	Matching demand and supply of spare parts	How is the demand and supply of spare parts matched by the MLO? Which types of supply chain cooperation exist and what do they look like?
5.	Spare parts demand characteristics	What are characteristics of the demand process? (plot a histogram with demand frequency and demand coefficient of variation)
6.	Spare parts supply characteristics	What are typical supply characteristics of the MLO? What does the supplier base look like (small/large suppliers, geographical)? What is the performance of suppliers in terms of on time delivery? Plot a histogram of the average (repair) lead time and variability
7.	IT-landscape	Which information systems are used to store and analyze maintenance and spare parts related data?
8.	Key Performance Indicators (KPIs)	Which KPIs are measured and what is the current performance? Which targets are set for the KPIs?
9.	Main observations	What are the main differences between the MLO in Chapter 2 and the case study company? How can these differences be explained?

Table 3.5 Content and related questions of sections II and III of the case study report.

Section IV of the report provides a comparison of the framework of Chapter 2 to practice in the case study company. It includes answers to both general questions (related to all processes in the framework) and process specific questions (see Table 3.6). Table 3.6 also includes general questions for the key informants.

	Process	Topic(s) & questions
	General questions (key informant)	What is the key informants position? What is the key informants work experience? What are the key informants tasks and responsibilities? What decisions does the key informant make?
	General questions (process)	Which department is responsible for decision making? What is the number of people (FTE) involved in decision making? What is the education level of the responsible decision makers? How often are decisions made (or reconsidered)? Based on what trade-off is the decision made? How does the asset life cycle influence decision making? Which information is used in decision making, how is the information registered and what is its quality? Which decision support system/models is/are used in decision making?
1.	Assortment Management	What is the origin of the current spare parts assortment, i.e. when were the parts added to the assortment? Based on what trade-off are parts added or removed from the assortment of (active) parts? Which (technical) information about spare parts is available and/or gathered? How is the technical information available (digital, books)? What is the main reason to gather technical information of spare parts? How often is the information updated? What approach is applied to gather/update the information (proactive or reactive)?
2.	Demand forecasting	How is planned and unplanned demand separated and registered? What trade-off determines whether or not to use advance demand information? What methods and information are used to forecast spare parts demand? How are forecasts from multiple methods combined, or is the best method selected? What are the (dis)advantages of the used forecast methods? How are planned and unplanned demand streams combined?
3.	Parts returns forecasting	Which return streams of parts are distinguished? Which return streams are registered? What methods and information are used to forecast part returns? What are the (dis)advantages of the used forecast methods?
4.	Supply management	What does the supplier base look like (small/large suppliers, small/large number of parts, monopolist)? What are the main characteristics of the repair shops (both internal and external)? How can the relationship with the (main) supplier(s) best be described? How can the relationship (and set of agreements) with internal repair shops best be described? Which information is shared with suppliers, and how often? What type of agreements are made with the suppliers?

	Process	Questions
	Supply management	Based on what assumptions regarding the supply of parts are inventory levels set (e.g. planned lead times, include lead time variability)? How (often) is the supplier performance managed by the MLO? How are the various parts of the total supply lead time managed by the MLO? What batching principles are applied (or dictated by the suppliers)?
5.	Repair shop control	What is the main motivation for having internal repair shops? Which internal repair shops can be distinguished, and what distinguishes them from each other? What order acceptance principles are applied in the repair shops? What workload control principles are applied in the repair shops? What capacity flexibility options are used by the internal repair shops? Which factors influence the repair lead times? What batching principles are applied? Which scheduling rules are applied?
6.	Inventory control	Are SRU and LRU inventory control integrated and if not, what is the motivation to separate these decisions? What inventory strategies are used? What ways to differentiate stocking policies between different spare parts are used? What assumptions are made in setting spare parts stock levels? What are motivations to deviate from steady-state inventory models (e.g. shelf life)?
7.	Spare parts order handling	What conditions should be satisfied before spare parts are released on orders? What are the main motivations to set these conditions? Which policy is used to release spare parts to requests? What procedure is followed in case of stockout of spare parts?
8.	Deployment	What conditions should be satisfied before spare parts resupply orders are released? What is the influence/freedom of a planner in adjusting and releasing spare parts resupply orders (in terms of quantity, timing, etc.)? How (often) are spare parts resupply orders followed up by deployment?

Table 3.6 Content and related questions of section IV of the case study report.

Section V presents some figures about the performance of the MLO and section VI concludes with a list of key informants that were interviewed, including their position in the organization and the list of topics that was discussed.

3.A.2 List of key informants and discussion topics

This Appendix gives an overview of the positions of the key informants that were interviewed during the case studies, including the list of topics (processes) that was discussed (see Table 3.7). The numbers in the third column relate to the process numbers in the first column of Table 3.6.

Company	Position	Discussion topics
NS	Supply chain manager	1,2,3,4,5,6,7,8
	Manager inventory control	1,2,3,4,5,6,7,8
	Tactical inventory planner	3,6,7
	Tactical inventory planner	2,3,6,7,8
	Purchase manager	4
	Maintenance engineer	1,6
	Maintenance system engineer	1,2
	Maintenance projects manager	2
	General repair shop manager	3,5
RNLAF	Supply chain manager	1,2,3,4,5,6,7,8
	Tactical inventory planner	1,6
	Tactical inventory planner	2,3,6
	Tactical inventory planner	2,3,6
	Operational inventory planner	1,7,8
	Operational inventory planner	1,7,8
	Purchase manager	4
	Purchase adviser	4,5
	Coordinator internal repair shops	4,5
	Maintenance engineer	2,3,6
	Maintenance planner	7
	Maintenance planner	2,7
	Maintenance projects manager	2,4
	General repair shop manager	5
KLM E&M	Supply chain manager	1,2,3,4,5,6,7,8
	Manager inventory control	1,2,3,4,5,6,7,8
	General repair shop manager	4,5
	General repair shop manager	4,5
	Purchase manager	4,5,7
	Customer order manager	7
	Project manager supply chain	1,2,6
	Project manager supply chain	2
	Repair shop development manager	5,6
	Repair shop development manager	5,6
	Repair shop manager	4,5
	Inventory planner	2,6,8
	Supply chain specialist	2,6,8
	Supply chain specialist	2,6,8

Company	Position	Discussion topics
KLM E&M	Supply chain specialist	2,3,6
	Supply chain engineer	1,2,6,8
RNLA	Supply chain manager	1,2,3,4,5,6,7,8
	General asset manager	2
	Asset manager	2,6
	Asset manager	2,5,6
	Asset manager	2
	Asset manager	1,6
	Manager maintenance engineering	2,6
	Tactical inventory planner	1,2,4,6,8
	Purchase manager	4,8
	General repair shop manager	5
	Manager order intake repair shops	5
	Repair shop manager	5
	Repair shop manager	3,5
	Project manager repair shops	5
	Repair shop planner	5,7
	Repair shop planner	5,7
	Maintenance planner	2,3,7
RNLN	Supply chain manager	1,2,3,4,5,6,7,8
	Supply chain adviser	1,2,3,4,5,6,7,8
	Manager inventory control	1,2,6,7,8
	Manager inventory control	6,7,8
	Manager maintenance engineering	2
	Manager maintenance engineering	2
	Repair shop manager	5,8
	Repair shop manager	5,8
	Repair shop manager	5,8
	Repair shop manager	5,8
	Tactical inventory planner	5,8
	Tactical inventory planner	1,2,3,6,7,8
	Repair shop planner	4,5
	Operational inventory planner	6,7,8
	Operational inventory planner	6,7,8
ECT	General asset manager	4,6
	Manager inventory control	1,2,3,6,7,8
	Operational inventory planner	1,2,3,6,7,8
	Purchase manager	4,8
	Senior Purchaser	4
	Asset manager	2
	Manager maintenance planning	2,3,7
	Maintenance engineer	2

Table 3.7 List of key informants and discussion topics.

Chapter 4

Repairable parts inventory systems: a literature review

"Als je voor elke positie de beste speler kiest, heb je nog geen sterk elftal maar een team dat als los zand uiteen valt."

Johan Cruijff

4.1 Introduction

Repairable parts inventory systems are systems in which a repair shop returns (failed) repairable parts to ready-for-use (RFU) stock. In Section 1.2 we showed, via the Lagrange relaxation method, that we can consider such a system for each internal repair shop individually. The internal repair shop does not share capacity with other supply sources, and SRUs are repaired by the same repair shop or replenished from external suppliers against a given lead time. In the remainder, we concentrate on designing control concepts for repairable parts inventory systems with one internal repair shop. A graphical representation of such systems is given in Figure 4.1[1].

[1]Some flows (no fault found, scrap and subcontracting repair work) are left out of the figure for the sake of clarity.

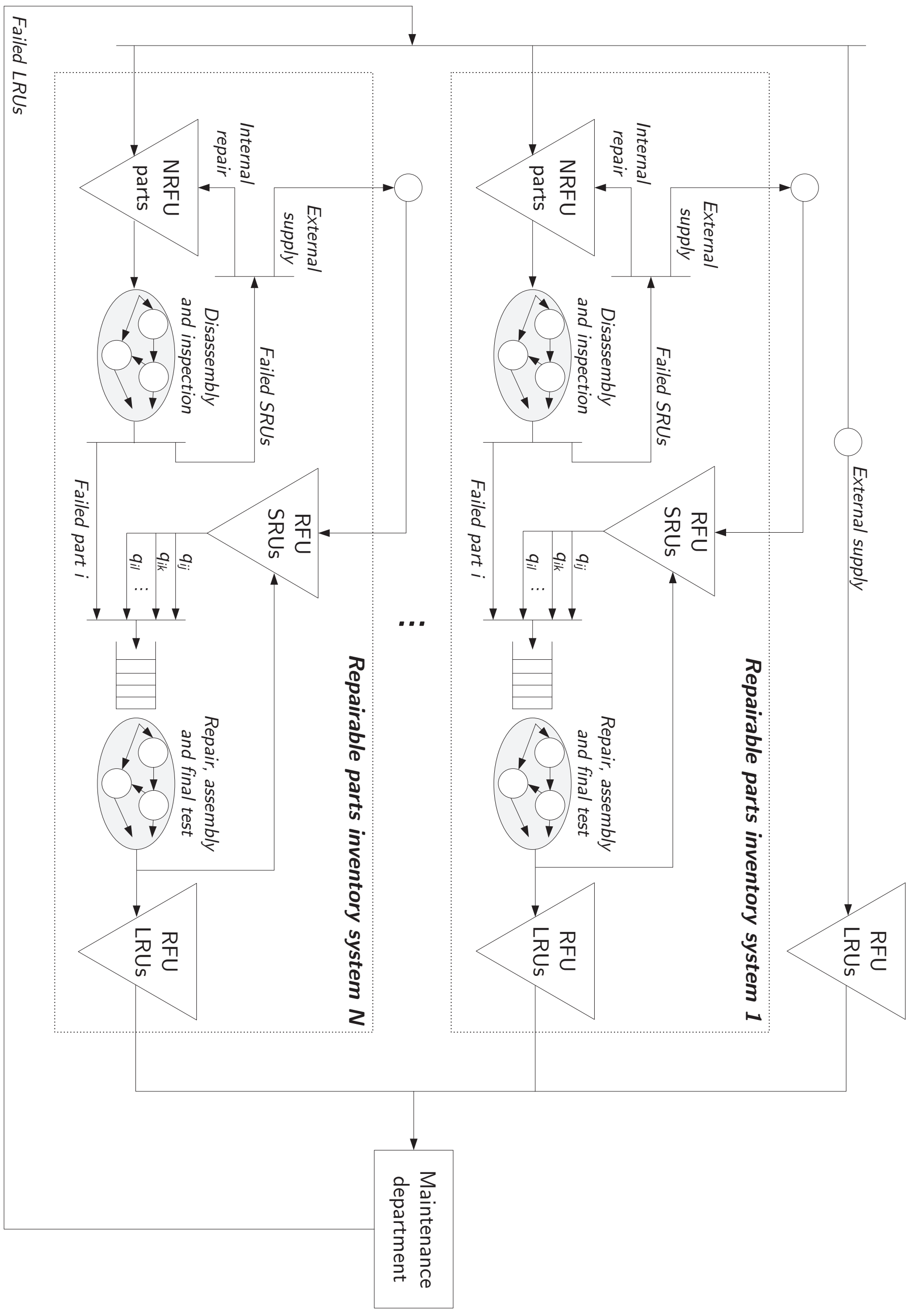

Figure 4.1 *Graphical representation of repairable parts inventory systems.*

Upon failure of an LRU part, a ready-for-use one is requested by the maintenance department. The request is fulfilled immediately by sending a ready-for-use LRU part from the ready-for-use LRU stock if available, and backordered otherwise. Simultaneously, the failed part is returned to the NRFU parts stock location, where it is buffered and from which it is sent into the repair process. The failed part is inspected and disassembled, the set of required SRUs is requested from stock and the failed (repairable) SRUs are returned to the NRFU parts stock location. The repair phase starts when all required SRUs are available, and the failed SRUs are replaced by ready-for-use SRUs. After completion of the repair phase, the part is tested and if it works according to its specifications, it is sent to the RFU parts stock location.

The objective of each repairable parts inventory system is to achieve a certain availability of the high-value capital assets (from a spare parts availability point of view) against acceptable logistics costs. The relevant decisions that are in scope of a repairable parts inventory system are the inventory policies of the LRUs and SRUs, the fixed and flexible repair capacities and the scheduling policy of the repair shop (see also Section 1.2.2). Clearly these decisions interrelate, and both repair capacity and inventory can serve as a buffer to meet the spare part availability requirements. Hence we need to find a control concept that optimally balances them. This raises the third research question of this dissertation:

3. *What are ways to design control concepts for repairable parts inventory systems, and which requirements are important to account for in the design process?*

To answer this research question, we review three streams of literature. An overview of these streams is presented in Section 4.2, and these streams are elaborated in more detail in Sections 4.3-4.5. In Section 4.6 we give an overview of reviewed literature, and discuss related research. Section 4.7 concludes this chapter.

4.2 Overview of literature streams

In repairable parts inventory systems, demand for a repairable part coincides with the return of a NRFU part that needs to be maintained, and a repair job can only start after a demand occurred. This is similar to *pull control systems*, in which production (repair) is triggered by actual demands. In this chapter, we therefore concentrate on literature (streams) related to pull control systems. Two well-known pull control systems are the base stock inventory policy and the Kanban control system (Buzacott and Shanthikumar (1993), Hopp and Spearman (2001)) without batching, i.e. in both systems we apply a one-for-one replenishment policy in which the use of an item to satisfy demand is immediately followed by an action to replenish the open stock position.

The advantage of the base stock inventory policy is that, once a demand for an item arrives, it directly releases a repair job to repair the failed part and hence immediately responds to demand. Disadvantage of a base stock policy is the fact that there is no restriction on the number of outstanding repair jobs, i.e. there is no capacity restriction or workload control mechanism in place. In such a situation, inventory and repair shop control decisions are decomposed. From an inventory control perspective, the repair shop can be regarded as a self-contained unit (or 'black box') that has sufficient flexibility options to always meet certain agreements made with inventory control (e.g. on lead times). In the first stream of literature, we review uncapacitated models for determining inventory (base stock) policy parameters, as well as models that focus on capacity dimensioning with flexibility options. Thereafter, we discuss how and when they can be combined to design control concepts for repairable parts inventory systems.

In case repair capacity and flexibility options are restricted, and resources are highly utilized, integrating capacity and inventory decisions may be worthwhile. The repair process may consist of one or multiple stages, based on the number of resources that are required in the repair process. In the second stream of literature, we review single-stage capacitated inventory models. This stream of literature focuses on inventory models in which capacity (flexibility) is limited, and in which the repair process is modeled as a single stage. The complexity to solve capacitated inventory problems increases in the complexity of the scheduling rule that is used. In a single-stage setting, scheduling rules reduce to priority rules only, and hence priority and capacity allocation decisions coincide. Furthermore, workload control does not add value in a single-stage (single resource), closed-loop setting and is therefore omitted in this stream, and a job is released as soon as capacity becomes available.

The Kanban control system is a way to model a workload constraint in a manufacturing process consisting of multiple stages. In the Kanban control system, the number of outstanding jobs in a (manufacturing) stage is restricted to the number of Kanbans. Once a finished part leaves the stage, the related Kanban card is stored in a Kanban store and available for a new job if needed. Because the number of Kanbans is equal to the base stock level, it is limited in responding to demand in situations with backorders, and especially performs worse when demand or processing times are highly variable. The advantages of both the base stock policy (direct response) and the Kanban control system (i.e. resource constraints and a workload control rule) are combined in *Generalized Kanban Control Systems* (Buzacott (1989)) and *Extended Kanban Control Systems* (Dallery and Liberopoulos (2000)), and these models are the topic of the third literature stream.

The three streams are described in more detail in the following sections, and we discuss how and when these streams of literature can be used to design repairable parts inventory systems. Section 4.6 summarizes the literature review, discusses some of the assumptions generally made in the literature and provides references to literature that relax these assumptions.

4.3 Stream 1: uncapacitated inventory models

In this stream of literature we focus on single-site, multi-item, (multi-indenture) inventory systems for repairable parts. The relevant literature on these types of models starts with the paper by Feeney and Sherbrooke (1966), who formulate the single-item version of this model. This paper has initiated a huge stream of literature on so-called "METRIC-type" models, which have a focus on a multi-echelon setting. Sherbrooke (1968) present a multi-echelon, multi-item inventory model (METRIC), that is extended to a multi-indenture setting (MOD-METRIC) by Muckstadt (1973) and later on improved by Sherbrooke (1986), by means of the VARI-METRIC model, and by Rustenburg et al. (2003). In our setting we can however restrict ourselves to a single-site version of the model, and we refer the reader to Sherbrooke (2004) for a description of a single-site, multi-item inventory system, and to Rustenburg (2000) for a description of a single-site, multi-item, two-indenture inventory system.

In these models, it is assumed that part demands originate from part failures and occur according to stationary Poisson processes. Failed parts are replaced by ready-for-use ones, if possible, otherwise a backorder occurs. Backorders are fulfilled on an FCFS basis. Failed parts are repaired by replacing at most one part at a lower indenture level. The repair lead times are assumed to be independent and identically distributed (i.i.d.) random variables, and can have a general distribution of which only its mean value should be known. A single-site, multi-item, multi-indenture variant of these models can be used to determine the optimal base stock levels of all LRUs and SRUs, given a limited budget for inventory investments, such that the average operational availability of the assets is maximized (i.e. the expected number of backorders for LRUs is minimized). Integrating LRU and SRU stocking decisions significantly reduces total inventory holding costs (Muckstadt (1973) and Rustenburg (2000)).

The assumption of i.i.d. repair lead times is valid if agreed repair lead times are relatively long (how long depends on the repair shop utilization), and repair shops will be able to meet these repair lead times with a high probability. In cases where capacity is flexible and available with little additional costs and lead time, this assumption is also valid. This flexible repair capacity can be arranged by e.g. applying overtime when the shop workload is high or by subcontracting repairs to external repair shops. In such situations, the repair shop can be considered as having 'unlimited capacity' and the aforementioned models are applicable. If realization of agreed repair lead times is however difficult, for example when the repair shop utilization is high, then this assumption leads to significant underestimation of stock levels and service performance (see e.g. Sleptchenko et al. (2002)).

Uncapacitated inventory models are relevant in situations in which inventory and repair shop control are decomposed by means of a 'lead time interface', and in which the repair shop operates as a self-contained unit that is responsible and easily able (because it has sufficient flexibility) to meet these repair lead times. This stream assumes that a regular repair lead time is set for each repairable item, and based on these repair lead times, inventory control is optimized. These repair lead times contain a certain 'slack', to enable the repair shops to realize them with a high probability. This slack should on the one hand be restricted, because it increases the required stock investment, but on the other hand should give sufficient flexibility to the repair shop. The trade-off between the costs of capacity flexibility and inventory should be made at an aggregate level, and the following references address such flexible capacity models.

In an environment with deterministic demand, Bitran et al. (1982) formulate the trade-off between inventories, fixed and flexible (overtime) labor capacity as a Linear Programming problem at an aggregate level. This problem assumes that average manufacturing lead times are given. Recently, Zijm and Schutten (2017) have proposed to determine these lead times by analyzing a production system under workload control, i.e. through a closed queuing network analysis. We hereafter focus on research related to a stochastic (maintenance) environment.

The coordination of initial inventory decisions and the use of flexible manpower has been researched by de Haas (1995). Capacity is modeled as a single operation and (a single type of) manpower is assumed to be the only resource required in the repair process. Simulation is used to estimate repair lead times for a predetermined repair capacity level and flexible manpower policy. The LRU and SRU base stock levels are determined with MOD-METRIC and use the estimated repair lead times as input. By enumerating over several repair capacity levels the total costs of inventory and capacity (flexibility) are evaluated. From the simulated scenarios, the best flexible manpower policy is selected. Related research on application of overtime policies has been done by Scudder (1985) and Scudder and Chua (1987).

Subcontracting repair work is another way to create capacity flexibility. Timely subcontracting is important, as the subcontracting lead time is usually longer than the internal repair lead time. Hence, it requires a much more proactive approach than when applying overtime policies. Keizers et al. (2001) study a repair shop that is responsible for both asset maintenance (preventive and corrective) and component maintenance, and in that sense deviates from our setting. They set a due date target, i.e. the percentage of jobs that has to be finished before the due date, and derive approximations for the makespan distribution. These approximations are used to determine the optimal subcontracting policy. Subcontracting seems especially beneficial for repair shops that face highly volatile demand.

So far, we discussed settings in which inventory and repair shop control make agreements on a regular lead time only. A somewhat different setting is a 'lead time interface' with two lead times. In this setting, a distinction is made between a lead time for *regular* and *expedited* repairs, and the number of expedited repair jobs is restricted to a certain threshold value. Arts et al. (2016) study how expediting repairs of failed LRUs can be used to find the optimal trade-off between stock investments and expediting costs in a single-item model with exponential repair times and fluctuating demand. They propose a (smart enumeration) method to jointly optimize the expediting threshold and base stock level of the LRU. Arts (2017) presents algorithms to determine a near-optimal solution for this trade-off in a multi-item model.

This stream of literature (implicitly) proposes the use of a 'lead time interface', with either one (only regular) or two lead times (regular and expedite). There is no constraint on the workload released to the repair shop. If the repair shop utilization is low, i.e. the repair lead times are more or less equal to the repair times, then a lead time interface with only regular lead times may be applied. If waiting times are substantial and mainly occur in front of the repair shop, and there exist (large) price and demand differences between the repairable parts, then effective prioritization of repair jobs is beneficial. A lead time interface is less applicable in that case, as it would increase the coordination load between the inventory manager and the repair shop on rescheduling repair job completion dates. A lead time interface with a regular and expedite lead time may be beneficial in situations in which the repair lead time contains (high) waiting times (for capacity) in the repair process, i.e. if multiple (highly utilized) resources are required in the repair process.

The single-indenture models are applicable if SRUs cause no delay of repairs. The multi-indenture inventory models do take delays due to out-of-stock SRUs into account. These models however require the assumption that at most one SRU fails upon failure of an LRU and that a failed LRU is inspected upon failure, at no loss of actual repair time. If SRUs cause no delay of the repairs, then a lead time interface is applicable. Integrating SRU and LRU stocking decisions (multi-indenture inventory models) would lead to a situation in which control (of SRUs) and responsibilities are not aligned if SRUs do cause delays of repairs, because timely completion of the repairs depends on the availability of the SRUs. The second assumption may be valid only if actual inspection times are (very) short compared to the actual repair times of the parts, and if waiting times in front of the repair shops are negligible. In that case, the repair lead time can start whenever a part is inspected and all required SRUs are available.

In Table 4.1 we summarize our findings regarding the applicability of the first stream of literature, options that may be valuable when designing repairable parts inventory systems, and modeling aspects that are not yet covered in this literature stream, by comparing the modeling assumptions to case studies.

Table 4.1 Summary of the review of literature stream 1.

Application requirements	**Control principle**
- All required resources have relatively low utilization rates, or capacity is sufficiently flexible in terms of additional lead time and costs - At most one SRU fails upon failure of a next higher indenture level part - Failed parts are inspected upon failure, at no loss of actual repair time - Large problems, due to low computational complexity	Integrating inventory control of LRUs and SRUs significantly reduces inventory holding costs, especially when multiple, expensive SRUs with a low demand rate may be required in the repair process

4.4 Stream 2: single-stage, capacitated inventory models

The second stream of literature focuses on single-stage capacitated inventory models, in which the repair shop has no capacity flexibility options. The repair process consists of a single stage that requires one type of limited capacity only. The models in this stream of literature assume that inventory control and repair priority setting are integrated. This stream of literature has started with deriving approximation methods to solve queuing type models with an exponential server for a single, multi-indentured LRU in a multi-echelon setting, see e.g. Gupta and Albright (1992) and Albright and Gupta (1993). We hereafter focus on the more recent contributions in this literature stream.

Sleptchenko et al. (2003) model a trade-off between repair capacity and inventory in a setting in which the repair shop consists of a single stage. The repair shop is modeled as a multi-class $M/M/k$-queue, in which the number of servers is a decision variable and repair jobs are executed on FCFS basis. Inspection of failed parts takes no time and is performed upon failure (which implies that repair activities may be preempted by inspection and diagnostics activities, but at no loss of actual repair time). This model is applicable for repair shops with identical servers that repair a small number of LRUs, as a limited number of LRUs already requires quite some computation time. Sleptchenko et al. (2003) mention that exponential service times are in general not realistic. They therefore also model the repair shop as an $M/G/c$-queue, but solving this model requires the assumption of at most one LRU. Still, there is room for improvement by applying more sophisticated scheduling rules as an alternative to the FCFS scheduling rule.

Adan et al. (2009) have developed an exact model for simultaneously determining initial stocks and static repair priorities. They consider a multi-item, single-indenture model in which capacity is modeled as an $M/M/1$-queue with multiple priority classes. Each LRU is assigned a priori to each of the priority classes (static priority rules). They show that initial stock investments can be reduced by 40% up to 60% by using only two priority classes, compared to an FCFS service discipline. This model is not restricted to situations with a limited number of LRUs, but again requires the assumptions of a single resource type and exponential repair times. Furthermore, SRUs are not included in the model, or are assumed to cause no delay of the repairs. Also Sleptchenko et al. (2005) developed a model comparable to that of Adan et al.

(2009) which includes SRUs, though this model is applicable in only a limited number of situations (small set of LRUs, maximum of two priority classes).

Static priority policies have the advantage that steady-state repair pipeline distributions can be determined (or approximated) independent of LRU base stock levels and actual on-hand stock levels. Relaxing this assumption makes an exact evaluation hard if multiple LRUs are considered. Tiemessen et al. (2017) study dynamic scheduling rules and use simulation optimization to determine near-optimal LRU base stock levels in a single server, multi-item, single-indenture setting. Chapter 3 of Tiemessen (2014) is an extension of Tiemessen and van Houtum (2013), and again uses simulation optimization to determine near-optimal base stock levels for myopic policies in a similar setting.

The models in this stream of literature assume that there is one decision maker, and repair priority setting and stocking decisions are integrated. These models are applicable in a setting in which the repair process consists of a single stage, and repair jobs are released once repair capacity becomes available. No coordination load between priority setting and capacity allocation is imposed, as these decisions coincide in these situations. Waiting times of the repair jobs will be mainly in front of the repair shops (waiting for capacity), or for the required SRUs. The models in this stream either assume that SRUs cause no delay of repairs, or parts are assumed to be inspected upon failure, at no loss of actual repair time. The latter assumption may be valid only if actual inspection times are (very) short compared to the actual repair times of the parts (and preemption of repair work is not inefficient), or when inspection and repair require dedicated resources, and transfer of information about the cause of the failure of parts takes no time. Table 4.2 summarizes our findings related to the second stream of literature.

Table 4.2 Summary of the review of literature stream 2.

Application requirements	**Control principle**
- Closed-loop repair process that consists of a single stage with restricted capacity (flexibility) - SRUs cause no delay of repairs, or failed parts are inspected upon failure at no loss of actual repair time. - Small problems, as computational complexity quickly increases in the number of decision variables	Integrating scheduling policy and stocking decisions significantly reduces inventory holding costs, especially when large demand and price differences between parts exist

4.5 Stream 3: Generalized Kanban Control Systems

In this section, we first present Kanban systems for multi-stage manufacturing environments. Thereafter, we explain how and when these systems can be applied to repairable parts inventory systems. A framework to analyze multi-stage capacitated manufacturing systems is the Generalized Kanban Control System (GKCS) or the Extended Kanban Control System (EKCS), introduced by Buzacott (1989) (also see Frein et al. (1995)) and Dallery and Liberopoulos (2000) respectively.

The manufacturing system consists of multiple stages in series, and each stage is a manufacturing process on its own (for example a job shop setting) with specific resources. Each stage has an output buffer and resources are not shared among the stages. The workload in each stage is controlled by a fixed set of Kanbans. All output buffers are controlled by a base stock policy and demand that cannot be fulfilled from stock directly, is backordered. The output buffer in the final stage is used to serve customer demand. Both GKCS and EKCS are based on control mechanisms that focus on the linkages between stages (not within a certain stage), and the corresponding decision variables are the number of Kanbans and the base stock level for each stage. See Figure 4.2 for a graphical presentation of the GKCS.

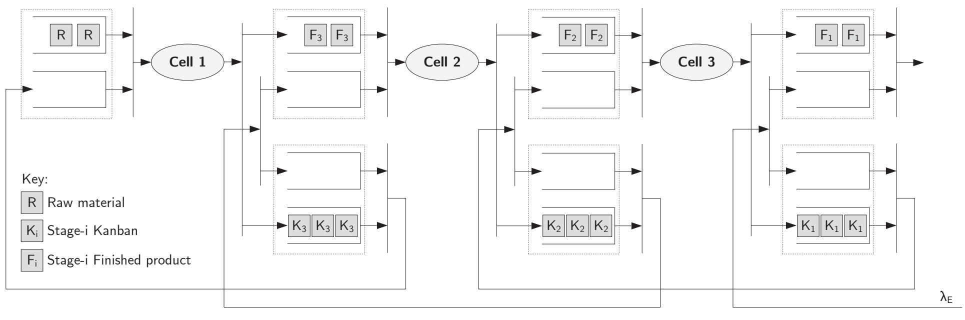

Figure 4.2 *Generalized Kanban Control System (from Zijm (2002)).*

In the GKCS, a production request for stage i starts once a demand for a stage i arrives and sufficient stage i Kanbans are available, and the production request is backordered otherwise. If available, a free Kanban is attached to the production request and together they are transferred into a stage $i-1$ demand. Hence, it directly triggers a production request in the preceding phase. This continues all the way to the first stage, but stops when no Kanbans are available at a certain stage. In contrast, in the EKCS a customer demand is directly transferred into the release of jobs in each of the preceding phases (not subsequently). Hence the EKCS responds quicker to customer demand than the GKCS, but at the expense of higher stock levels. In the sequel, we focus on GKCS only and refer the reader to Dallery and Liberopoulos (2000) for an extensive description of the EKCS and its differences and (dis)advantages with respect to GKCS.

The GKCS framework coordinates capacity and stocking decisions at the tactical level, assuming a FCFS service discipline. The base stock levels and the number of Kanbans (maximum workload) are determined simultaneously, and total repair lead times per product can be estimated. These repair lead times serve as targets for the completion dates of repair jobs at an operational level. Each stage in the repair process operates autonomously and is responsible for meeting the internal lead times for each product. A proper scheduling rule (to prioritize jobs to be released at each of the stages) is applied to meet the completion dates. The GKCS does not include decisions regarding the stocking of raw materials. Only at the first stage raw materials are required, and these are assumed to be always available.

The GKCS has two variants: the single class and multi-class GKCS. The number of classes represent the number of different parts that is produced. Parts from all classes visit all stages in the same order and jobs waiting in front of a stage are served FCFS. Each part has its own base stock level at each stage. In the multi-class GKCS we distinguish three types of workload constraints: (1) each class has a dedicated set of Kanbans, (2) all classes share one set of Kanbans, or (3) a combination of (1) and (2). Multiple combinations are possible because workload constraints can be modeled differently for each class in each of the stages.

Several (approximation) techniques have been developed to analyze GKCS, even for the case of general service times, see e.g. Buzacott and Shanthikumar (1993) and Buitenhek et al. (2000). Buitenhek et al. (2000) use approximate mean value analysis to evaluate system performance measures in a production-to-order system, for a given number of Kanbans, such as the internal lead times per stage, the total manufacturing lead time and customer response times per class.

The GKCS integrates decisions regarding base stock levels and the number of Kanbans at each of the stages (and hence it is assumed that there is no interface between inventory and repair shop control). The GKCS requires the assumptions that stages have dedicated resources and that stages are visited in the same sequence. These stages can be regarded in multiple ways when looking at repairable parts inventory systems.

First, as has been noted by Zijm and Avşar (2003), the stages in the GKCS can be regarded as the indenture levels of repairable parts. Zijm and Avşar (2003) present an approximation method to determine the base stock levels for a single-location, single LRU, two-indenture level model with two SRUs (see Figure 4.3). A workload constraint can be incorporated in this model by combining it with approximate mean value analysis (Buitenhek et al. (2000)), see e.g. Zijm (2002). This extension also allows the use of general service times.

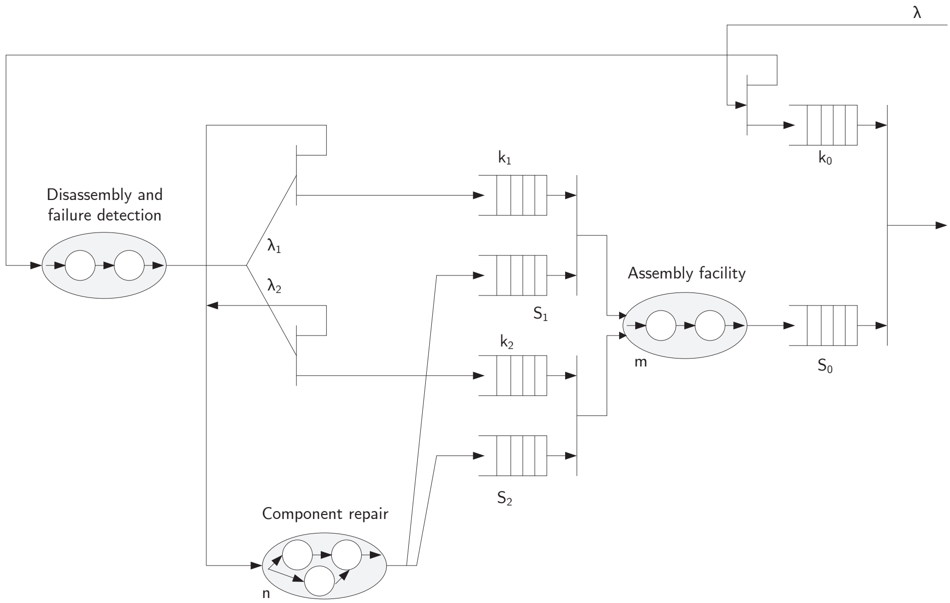

Figure 4.3 *A two-indenture model with a single LRU and two SRUs (from Zijm and Avşar (2003)).*

In this setting the final stage can be referred to as the repair and assembly process of an LRU, and each preceding phase can be seen as the repair and assembly process of a lower indenture level SRU. Both LRUs and SRUs have their own base stock level and number of Kanbans, and repair operations are modeled as $M/M/1$-queues as well as product form queuing networks with a FCFS repair discipline.

Second, the stages can be regarded as the specific resources (e.g. different skill levels of the repair men, test equipment) that are required in the repair process. In that case, the assumption that each stage has dedicated resources is satisfied. It however requires that resources are used in the same sequence for each part and each repair shop (i.e. a fixed repair routing).

Third, the stages can be regarded as the phases (disassembly and inspection, repair and assembly and final test) in the repair process. The assumption that each job visits the stages in the same sequence is satisfied, but then the assumption that each repair phase requires a separate set of resources is needed. Table 4.3 summarizes our findings related to the third stream of literature.

Table 4.3 Summary of the review of literature stream 3.

Application requirements	Control principle
- Multi-indenture setting in which both inspection as well as repair and assembly of parts of each indenture level require separate resources - Single-indenture setting in which each repair job has a fixed routing - Single-indenture setting in which each repair phase requires dedicated resources - Small problems, as computational complexity quickly increases in number of decision variables and priority classes	Integrating inventory and workload control on each stage of the repair process significantly reduces costs, especially when the repair process consists of multiple stages with dedicated resources

4.6 Overview of reviewed literature and relaxations

Table 4.4 provides an overview of the literature we have reviewed. In columns 2-5 we list the four decision functions that are in scope of a repairable parts inventory system. For each stream of literature and each reference therein, we mention how the decision function is modeled. From this overview, we can conclude that a model that includes all decision functions is not yet available in the literature.

The reviewed literature assumes that inventories are controlled by base stock policies (i.e. one-for-one repairs), that a part is repaired by replacement of at most one SRU and that SRU demands are served FCFS. Next, we discuss these assumptions and mention some papers that relax them.

Application of one-for-one repair (or replenishment) is a common assumption in the spare parts literature (Sherbrooke (1968)), because repairable spare parts are usually typified being both expensive and slow moving. For such items economies of scale are in general absent, which justifies the assumption of base stock replenishment policies.

Table 4.4 Literature related to repair shop and inventory control

Stream	Scheduling policy	Repair capacity	LRU base stock levels	SRU base stock levels	References
Uncapacitated inventory models - *Only regular repairs*	Not applicable	Not applicable	Decision variable	Decision variable	Feeney and Sherbrooke (1966), Sherbrooke (1968), Muckstadt (1973), Sherbrooke (1986), Rustenburg (2000) Sherbrooke (2004), Rustenburg et al. (2003)
- *Regular and expedited repairs*	Static/dynamic policies	Decision variable	Predetermined	Predetermined	de Haas (1995), Scudder (1985), Scudder and Chua (1987)
			Not applicable	Not applicable	Keizers et al. (2001)
	Expediting policy	Predetermined	Decision variable	Not applicable	Arts et al. (2016), Arts (2017)
Capacitated, single-stage inventory models	FCFS	Decision variable	Decision variable	Decision variable	Sleptchenko et al. (2003)
	Static policy	Predetermined	Decision variable	Not applicable	Adan et al. (2009)
			Decision variable	Decision variable	Sleptchenko et al. (2005)
	Dynamic/ myopic policy	Predetermined	Decision variable	Not applicable	Tiemessen et al. (2017), Tiemessen and van Houtum (2013), Tiemessen (2014)
Capacitated, multi-stage inventory models	FCFS	Decision variable	Decision variable	Not applicable	Generalized Kanban Control Systems: a.o. Buzacott and Shanthikumar (1993), Frein et al. (1995)
		Predetermined	Decision variable	Decision variable	Zijm and Avşar (2003)

Chua et al. (1993) consider an environment in which batching of repairs is possible and which results in a reduction of the average repair time per item. Their environment is characterized by a high standardization of repair work and high set up times. They conclude that batching is profitable in such environments. Other work related to batching repairs is Scudder (1986).

The assumption that at most one SRU fails upon failure of a next higher indenture level part is highly questionable, see for example Van Jaarsveld et al. (2015). On the other hand, this assumption can be justified if one would include only expensive SRUs in the model and assume that cheap SRUs are always available. If the individual failure rates of expensive SRUs are low (less than 20%, say), then the probability that two expensive SRUs are required simultaneously is small. Hence, the outcome of the models is only moderately affected by this assumption.

Van Jaarsveld and Scheller-Wolf (2015) present alternative rules for allocating SRUs to repair jobs. With the assumption of (at most) one (expensive) SRU failure, we can however justify the assumption of SRU demands being handled FCFS: an SRU that is served from stock directly makes the repair job 'ready to repair', as all required SRUs are available. Only if commonality is at hand, different parts may compete for the same SRU and the assumption of FCFS may be non-optimal. With low failure rates however for both LRUs and SRUs, these situations will rarely occur.

4.7 Concluding remarks

In this chapter, we have studied the literature to answer the third research question of this dissertation:

3. *What are ways to design control concepts for repairable parts inventory systems, and which requirements are important to account for in the design process?*

From the literature review, we have found a number of ways to design control concepts for repairable parts inventory systems. In one stream of literature, inventory and repair shop control are decomposed by means of a lead time agreement, and decisions related to inventory control (determination of base stock levels) and repair shop control (capacity dimensioning and scheduling policy) are modeled and solved independently. The applicability of these designs is limited to situations in which the assumption of i.i.d. repair lead times is justified, i.e. when repair lead times are long compared to the repair shop utilization, or when capacity is flexible.

To a certain extent, it seems possible to integrate decisions related to inventory and repair shop control. In a second stream of literature, a single decision maker is responsible for inventory control and repair priority setting. The applicability of this design and the models that are reviewed is limited to situations in which the repair process consists of a single stage requiring a single type of capacity only. Furthermore, SRUs are assumed to cause no delay of repairs, or inspection is not needed to reveal which SRUs are required.

A third literature stream concentrates on the more advanced Generalized Kanban Control Systems (GKCS). In such systems, decisions regarding base stock levels and the number of Kanbans are fully integrated. Application of the GKCS concept to repairable parts inventory systems is limited to a number of situations. In a single-indenture setting, GKCSs are applicable if each (combination of subsequent) repair phase(s) requires dedicated resources, or when the repair routing is fixed and repair capacity is homogeneous. When a multi-indenture setting is considered, the GKCS concept is applicable if both inspection and repair and assembly of parts at each indenture level require separate resources.

The models of streams 2 and 3 assume that decisions regarding inventory and repair shop control are integrated, and that no interface is required between them. The computational complexity to solve these models quickly increases in the number of decision variables, and this may limit their applicability. This aspect should be taken into account when designing control concepts for repairable parts inventory systems.

We have not found references for some of the control concepts applied in our case study companies. This is partly due to the fact that some of these concepts focus only on operational decision making, while at the tactical level no agreements are made. Nevertheless, in particular the 'min-max concept' currently applied at the NS (see Section 3.3 for a description) is an interesting topic for further research.

From the literature review we have detected principles for control in repairable parts inventory systems. First, we have found indications that SRU and LRU control should be integrated. In the designs of repairable parts inventory systems, inventory and repair shop control are sometimes decomposed by means of a lead time agreement and timely finishing these repair jobs depends on the availability of SRUs. To align control with responsibilities, this requires that SRU and LRU control are decomposed. If and when SRU and LRU control should be integrated is an important choice when designing repairable parts inventory systems. Second, we have found that integrating repair priority setting and LRU stocking decisions substantially reduces the inventory investments. If and when these decisions should be integrated, is an important choice as well, because repair priority setting also influences the efficiency (and costs) of the repair shop. Third, we have found that workload mechanisms are applied at stage level, and hence repair jobs should be released based upon a cap on aggregate flows per resource (or stage).

The three streams of literature we have reviewed are applicable in specific situations. We validate their applicability in Chapter 5, by comparing the application requirements and control principles in multiple case studies.

Chapter 5

Case studies: part II

"Vindt de zwakke plek van je tegenstander en je hebt gewonnen."

Johan Cruijff

5.1 Introduction

In Chapter 4 we have reviewed three streams of literature that, under certain conditions, can be applied to design control concepts for repairable parts inventory systems. In each stream of literature, trade-offs between two or more decision functions in repairable parts inventory systems are modeled and analyzed. As such, the literature has provided insights in the conditions under which it is beneficial to explicitly model a certain decision function, i.e. when a certain decision function significantly influences the performance of repairable parts inventory systems. Based on these insights, we have formulated control principles that should be taken into account when designing repairable parts inventory systems.

So far, it is however not clear whether (the models from) the reviewed streams of literature can be applied in practice, and in which situations application is beneficial. Therefore, we formulate the fourth research question:

4. Which models and control principles, identified in the literature review of Chapter 4, are applicable to design control concepts for the repairable parts inventory systems in the case study companies?

To answer this research question, both quantitative and qualitative data are collected during interviews and plant visits in a multi-case study research design. The aim of the case studies is to extensively describe (the characteristics of each step) in the repair process, i.e. the process of returning a failed and subsequently repaired part to the ready-for-use (RFU) status. In total 22 repairable parts inventory systems from three MLOs are investigated and analyzed. We have assessed the applicability of the three streams of literature, by comparing the application requirements and control principles to the data gathered from the case studies. The findings we present in this chapter form the basis for the designs of the control concepts for repairable parts inventory systems in Chapter 6.

This chapter is organized as follows. In Section 5.2 we describe the methodology used to collect the empirical data. Section 5.3 presents the qualitative and quantitative data we gathered from the case studies, and in Section 5.4 we interpret the main findings from analyzing them. Section 5.5 discusses the research limitations and implications, and Section 5.6 concludes this chapter.

5.2 Methodology

In case study research, it is important to guarantee the consistency of findings arising from several different data collection methods (Yin (2003)). We have selected repair shops in three companies that both use and maintain high value capital assets: Gemeentelijk Vervoersbedrijf Amsterdam (GVB), KLM Engineering & Maintenance (KLM E&M) and Royal Netherlands Army (RNLA). These companies use and maintain assets with a varying level of technological complexity (subway carriages, airplanes, and military equipment, respectively), and also the LRUs have a varying level of technological complexity. The companies operate in different markets. In this way we are able to exclude the effects of exogenous variables (*external validity*).

The case study companies have different motivations for having internal repair shops. For GVB, minimization of repair costs is the main driver. In general they execute only repair work for parts with a relatively high demand rate, i.e. for which it is economically beneficial to set up the internal repair shops. All other parts repair work is subcontracted to external repair shops. KLM E&M services repairable parts (components) for multiple airline companies worldwide. Their focus is on optimizing the profit margin (turnover minus costs) of the repair work, and hence economic factors play a role (e.g. worldwide competition between repair shops). The main motivation for RNLA is to have or maintain a high knowledge level of the asset technologies, and economic factors are less dominant.

In the case study companies, the internal repair shops maintain a wide variety of LRUs and SRUs. The total number of LRUs and SRUs that they maintain varies from 800 up to 5,000 per company. The number of repair shops varies from 3 to 15 per company. In total 22 repair shops were investigated and analyzed.

Starting point for the case studies are the interview summaries from the first set of case studies (see Chapter 3). Based on the interview summaries and insights from plant visits, we have been able to make a first draft description of the entire repair process, i.e. the process of returning a non-ready-for-use (NRFU) repairable part to the RFU status. We have discussed this draft with several repair shop managers and employees of each of the three companies, and we have made a final description based on their feedback (*internal validity*). This final description of the repair process can be found in Section 5.3, and the main point of contact of NS, RNLAF and RNLN[1] have confirmed that the final description reflects their reality (*triangulation*).

For each step in the repair process, we have determined a set of figures needed to assess the applicability of the models and control principles. The use of quantitative figures is preferred. Nevertheless, we have not been able to easily determine some figures (e.g. variation in the repair job routing), or their value is limited. For example, we have estimated the average repair man utilization by dividing the yearly repair shop workload (in working hours) by the average number of working hours per year, summed over all repair men. Based on this calculation, we have found a lower utilization rate than the repair shop manager expected. Due to a lack of information, we have not been able to correct these numbers for factors such as absence through illness, time spent on education and training programs, and time spent on other repair activities conducted by the repair men (e.g. repair work for external customers, or asset maintenance in the case of company 3).

The quantitative data we have gathered is enriched with qualitative data gathered through interviews with repair shop managers and employees. As such, we are better able to interpret the figures we measured. For some steps in the repair process, we have

[1] Please note that ECT does not have internal repair shops and is therefore not consulted.

determined only qualitative figures. These figures are gathered through discussions with repair shop managers and employees, during the tours through the repair shops.

5.3 Empirical results

In this section, we first give an overview of the repair shop layout and the repair process. Next, quantitative and qualitative figures are presented about the demand and economics of the repair shops, and the capacities and materials required in the repair process.

Repair shop layout

Each repair shop is responsible for the repair of a distinct set of parts (both LRUs and SRUs). The inspection and final test of these parts is always conducted by the responsible repair shop, and the repair shop is also responsible for tracking progress of the repair work that is (partly) subcontracted by them to external repair shops. Each repair shop consists of one or multiple work stations, each with its own specific equipment.

The layout of the repair shop concerns the arrangement of machines, tools and equipment used in the repair process, and the layout of the individual work places. In many cases the choice for a certain layout depends on variables such as the demand and economics of the repair shops, but also on the characteristics of the maintenance that is conducted by the repair shops. We have observed three types of repair shop layouts in the case study companies: a product, process and job shop layout.

In a product and process layout the work stations and equipment are located along a single 'production' line, and the repair routing is fixed. The difference between the two is that individual workers perform a variety of tasks in a product layout, whereas in a process layout they only perform a small number of tasks. In a job shop layout, work stations are grouped by the nature of skills and technological processes involved. Not every work station is visited in each repair job, and visits to the work stations (i.e. the repair routing) may vary from one job to another. In Table 5.1, we give an overview of the layout for each repair shop.

A product or process layout is found in repair shops that conduct (high-volume) standardized repair work such as repair shops 2.10 and 2.11. In these repair shops, the repair men capacity is homogeneous. Between the repair shops with a job shop layout, we can see clear differences. In shops with a homogeneous repair capacity, every repair man can solely complete all repair jobs[2]. In shops with a heterogeneous repair men capacity either every repair man can solely complete a subset of repair jobs (based on their skill set), or a variety of skills from multiple repair men is required to complete a repair job. The latter case is for example found in repair shops 2.3, 2.4 and 2.5, where a higher skilled repair man has to conduct the final test of a part.

Among the repair shops a diversity of technologies is found: avionics, mechanics, hydraulics, electronics, etc. The extent to which there exist differences in skill sets between repair men, is found to be correlated with the diversity of products that is repaired. Repair shops that maintain rapidly evolving products, such as avionic and electronic parts, tend to have a higher product diversity than the more 'traditional' products. The product diversity is mainly caused by the fact that parts of both old and new (asset) technologies have to be maintained, requiring a different set of skills.

The repair process

In Figure 5.1, we give a detailed description of the repair process that is observed in all case study companies. This figure contains numbers to which we refer in the remainder of this section. Repairable LRUs and SRUs, that are replaced during line or shop level maintenance respectively, are returned as NRFU items to a central point in the service network (1 and 2), where they are buffered and from which they are sent into the repair process. The return process of NRFU parts mainly depends on the part failure patterns but also on the transportation policy that is applied (e.g. batching of returns from a local warehouse).

Repairable parts are assigned to a repair shop where the repair process starts (with inspection) and ends (with a final test). Parts for which the company does not have the repair capabilities are subcontracted to external repair shops and sometimes repair work can be subcontracted to avoid a temporarily high repair shop workload (3).

The first (possible) phase in the repair process is a pre-inspection phase, in which it is tested whether the part still functions according to its specified requirements. If so, the part is marked 'no fault found' and immediately sent to the RFU stock (4). It is obvious that such pre-inspection is only required if it is not clear whether the part failed (based on visual inspection), which is often the case for electronic and avionic

[2]We emphasize that 'homogeneous' should here be interpreted as: there exist no or only small differences in the skill sets between repair men. Hence there may be a small number of exceptions in which a repair man is not able to solely complete a repair job.

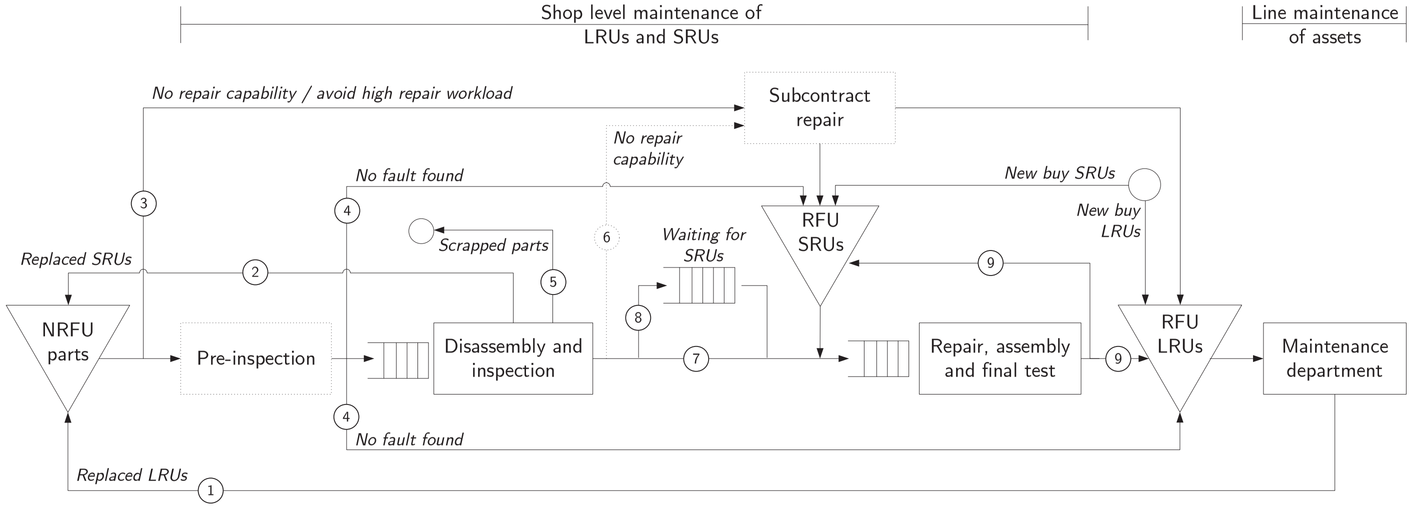

Figure 5.1 *Supply chain of repairable spare parts including the repair process.*

Table 5.1 Qualitative characteristics of the case study repair shops.

Shop	Description	Repair men capacity	Repair shop layout	Product diversity
1.1	Metal workshop	3 (heterogeneous)	Job shop	Low
1.2	Electronics	4 (homogeneous)	Job shop	Low
1.3	Electric (engines)	7 (homogeneous)	Job shop	Low
1.4	Hydraulics/pneumatics	10 (homogeneous)	Job shop	Low
2.1	Mechanic/electronic indicators	12 (heterogeneous)	Job shop	High
2.2	Galley & emergency equipment	17 (homogeneous)	Job shop	Low
2.3	Flight guidance & warning	17 (heterogeneous)	Job shop	Low
2.4	Panel, audio & controllers	11 (heterogeneous)	Job shop	High
2.5	High frequency	10 (heterogeneous)	Job shop	Low
2.6	Electromechanical components	12 (homogeneous)	Job shop	Low
2.7	Fuel & power systems	15 (homogeneous)	Job shop	Low
2.8	Hydraulic power systems	12 (heterogeneous)	Job shop	High
2.9	Engine exterior/metalworkshop	13 (heterogeneous)	Job shop	High
2.10	Brakes	6 (homogeneous)	Product	Low
2.11	Wheels/tires	22 (homogeneous)	Process	Low
2.12	Heat exchangers	4 (heterogeneous)	Job shop	Low
2.13	PDU cargo	6 (heterogeneous)	Job shop	Low
2.14	Water & Waste	4 (homogeneous)	Job shop	Low
2.15	Seats & emergency equipment	6 (homogeneous)	Job shop	Low
3.1	Electronics	(homogeneous)	Job shop	Low
3.2	Hydraulics	(heterogeneous)	Job shop	Low
3.3	Engines/mechanics	(heterogeneous)	Job shop	High

parts. After the pre-inspection phase, it is still not clear what the cause of the failure is.

In the inspection phase, the part is disassembled and the corresponding sub parts are cleaned and inspected to identify the failure mode. It is determined which SRUs failed and should be replaced by ready-for-use SRUs. Failed parts are either scrapped (5) or returned for repair (2). A repairable part is scrapped (condemned) if it is beyond technical or economical repair and the corresponding repair job stops, if it concerns the highest indenture level part.

After the inspection phase it is known which SRUs and which activities (resources) are required in the repair phase, and the repair routing is determined. Sometimes the repair shop does not have the capability to repair the part, and in that case the part is subcontracted to an external repair shop (6). We would like to emphasize that the inspection is not always perfect, that is, sometimes still some more 'inspection' is required during the repair phase. Furthermore, sometimes SRUs are used during the inspection phase in order to identify the failure mode (via 'trial and error').

Parts, for which all required SRUs are available, are queued in front of the repair phase (7). All other parts are put aside by the repair man, waiting for the required SRUs to arrive (8). If inspection and repair of a part are executed by separate repair men, then information about the cause of the failure has to be transferred between them. This is inefficient because it requires extra time and hence increases the repair man utilization. If the repair phase is conducted by the same repair man, then no information transfer is required. Nevertheless, the actual repair time for a part is higher in case the interruption between inspection and repair takes too long (e.g. more than two weeks), because the repair man has to re-analyze some problem details if no proper documentation is available.

In the repair phase the failed SRUs are repaired or replaced by ready-for-use SRUs[3], and all other repair activities are conducted. Upon completion, the part is tested and if it functions according to its specifications, it is send to the RFU stock (9).

Characteristics related to demand and economics of repair shops

Demand for repairable parts is in general unplanned and highly stochastic, and needs to be fulfilled immediately (see also Section 1.1.1). As shown in Table 5.2[4] the volatility in the arrival of repair work, represented by the coefficient of variation (CV) of the number of NRFU part returns per week (cv_j^d), is high for many repair shops. One way to deal with the high volatility in repair volumes is to make use of flexible repair capacity, e.g. by temporarily hiring additional repair capacity, subcontracting repair work to external repair shops or working overtime. This is however not straightforward as we will discuss hereafter.

Internal repair shops often have a high utilization (around 90%) and repair work is executed by repair engineers, who are certified professionals that often have to follow a long training program (of one year, say) before they can do the job. Recruiting such professionals is often difficult. Within this population of professionals, furthermore a distinction can be made between generic professionals and specialists. Especially the unavailability of specialist skills can lead to substantial waiting times in the repair process. Besides certified professionals also expensive, dedicated and highly utilized

[3] The phenomena of *part matching* (failed SRUs are repaired and have to return to the original next higher indenture part for repair), as observed by Guide Jr et al. (2000), was only rarely observed in our case studies.

[4] The definitions of the variables are given in the caption below the table. The shop number consists of the company number (left) and repair shop number (right), separated by a dot.

Table 5.2 Quantitative figures of the case study repair shops.

Shop	Description	n_j	μ_j^d	cv_j^d	r_j^d	q_j^c	μ_j^s	cv_j^s	q_j^i	μ_j^m	μ_j	p_j^0
1.1	Metal workshop	47	22	0.6	20%		4.1	1.2		15.3	5.7	11%
1.2	Electronics	173	59	0.5	35%		1.1	1.9		4.4	0.4	81%
1.3	Electric (engines)	73	32	0.4	30%		4.3	1.2		10.9	3.5	22%
1.4	Hydraulics/pneumatics	61	24	0.4	40%		6.7	1.1		23.6	11.4	18%
2.1	Mechanic/electronic indicators	124	21	0.4	30%	11.1%	3.4	1.1	10.7%	10.5	0.8	78%
2.2	Galley & emergency equipment	102	97	0.2	23%	6.6%	1.9	0.9	9.6%	16.7	2.2	65%
2.3	Flight guidance & warning	73	20	0.3	22%	0.1%	15.1	1.0	2.8%	23.4	3.3	31%
2.4	Panel, audio & controllers	184	60	0.2	9%	1.4%	7.0	1.8	2.4%	13.3	2.3	42%
2.5	High frequency	99	19	0.4	32%	9.8%	5.6	1.1	4.8%	15.1	1.4	56%
2.6	Electromechanical components	126	44	0.3	21%	2.0%	3.9	0.8	11.1%	19.2	3.0	41%
2.7	Fuel & power systems	79	21	0.4	22%	8.9%	6.7	1.5	1.9%	21.7	3.0	61%
2.8	Hydraulic power systems	94	12	0.5	40%	6.5%	11.1	1.1	10.4%	29.2	6.5	41%
2.9	Engine exterior/metalworkshop	52	10	1	25%	0.8%	22.8	2.3	8.0%	31.4	4.9	17%
2.10	Brakes	10	13	0.4	55%	3.2%	10.1	0.7	12.2%	52.3	14.3	20%
2.11	Wheels/tires	21	160	0.2	35%	0.1%	4.8	0.6	8.9%	27.5	0.4	75%
2.12	Heat exchangers	5	3	0.7	40%	0.1%	7.4	0.7	1.1%	6.3	3.4	16%
2.13	PDU cargo	107	36	0.3	45%	1.3%	2.3	2.0	13.7%	21.1	3.6	35%
2.14	Water & Waste	72	27	0.5	18%	8.7%	2.5	0.9	10.9%	31.6	4.1	51%
2.15	Seats & emergency equipment	79	15	0.3	55%	0.2%	8.8	0.6	7.7%	18.9	4.3	26%
3.1	Electronics	49	8	1	45%		8.8	2.1		54.6	7.1	25%
3.2	Hydraulics	56	4	1.3	37%		10.3	2.7		9.5	8.0	24%
3.3	Engines/mechanics	95	21	0.6	69%		9.4	2.0		46.2	14.5	14%

n_j: number of different parts that is repaired by repair shop i (based on data set of 2 years)
μ_j^d: average number of returned NRFU parts per week
cv_j^d: coefficient of variation of the number of returned NRFU parts per week
r_j^d: smallest fraction of parts responsible for (at least) 80% of total number of returned NRFU parts
q_j^c: fraction of unsuccessful repairs (condemnation or scrap rate)
μ_j^s: average service (inspection plus repair) time per part in hours (corrected for batching)
cv_j^s: coefficient of variation of the service time (inspection plus repair) per part in hours (corrected for batching)
q_j^i: fraction of job service time spent on inspection
μ_j^m: average number of (different) next lower indenture level parts, possibly required in a repair job
μ_j: average number of (different) SRUs required per repair job
p_j^0: fraction of repair jobs in which no SRU usage is registered

(test) equipment is used in the repair process, and such equipment is in general not easily available at suppliers.

From this discussion, we can conclude that the applicability to temporarily hire additional repair capacity is limited because of its long effectuation time. From the case studies we observe that subcontracting repair work, to avoid a temporary high workload and resulting long repair lead times, also rarely happens. This is because external repair lead times are in general much longer than internal repair lead times and hence the benefits of subcontracting are minimal. We observe that internal repair shops prefer to work overtime to prevent high workloads and long repair lead times, though we note that this is only possible up to a certain limit per period.

Characteristics of materials and capacities required in the repair process
The inspection phase of a repair job is similar to the 'work preparation' phase in a production process, i.e. only after inspection it is known which materials and capacities (and in which sequence) are required in the repair phase. Table 5.2 contains several quantitative figures. From these figures, we can draw the following conclusions related to the *materials* required in the repair process:

- Inspection takes a significant amount of time (q_j^i), with a weighted average of 7.8% of the total service time over all repair shops and ranging from 1.1% to 13.7% per repair shop.
- The repair shops face different levels of uncertainty when it comes to the SRUs required in the repair process: the average number of (different) SRUs that is required to repair a part (μ_j) ranges from 0.4 to 14.5 (these numbers include also repair jobs in which no SRUs were required), the average number of (different) SRUs potentially required to repair a part (q_j^m) ranges from 4.4 to 54.6[5]. The fraction of repair jobs in which no usage of SRUs was registered also significantly differs[6].
- On average a low number of repairable parts is scrapped in the repair process (q_j^c), with a weighted average of 3.3% over all repair shops and ranging from

[5] These numbers are based on the actual usage of SRUs in our data set. For companies 2 and 3, the usage of *repairable* SRUs is not included because either the company applies a "repair-by-repair"-strategy for repairable SRUs, or their usage is not properly registered.

[6] Usage of SRUs from floor stock (very cheap, fast moving items for universal use that are supplied to the repair shop in bulk) is not registered. These items however constitute only a very small part of the total costs, and repair jobs are in general not interrupted because such items are not available.

0.1% to 11.1% per repair shop.

To conduct the repair work, several types of capacities are required: repair men, machines or equipment, service tools and repair manuals. From discussions with repair shop managers and employees we found that the number of required (specialized) skills and specific, high-value (test) equipment were the main bottlenecks in timely finishing repair jobs. To use the high-value equipment, in general specialized skills are required, and as such these factors are found to be correlated. In the remainder of this section, we therefore focus on the number of specialized skills that is required in the repair process.

Each repair man has a certain set of repair skills (or capabilities), and a variety of skills might be required in the repair process. Due to a lack of data, we have not been able to determine the set of repair skills required for each individual repair job. Therefore we have decided to use qualitative figures, that can be found in Table 5.2. These figures serve as an indication for the number of specialized skills required in the repair process.

First, the number of repair men gives an indication of the size of the repair shops[7]; for large repair shops, we expect that it is easier to have the required skills available than for small repair shops. Second, the heterogeneity of the repair (men) capacity gives an indication of the complexity of the repair work that is executed by the repair shop, and the complexity to have repair men trained (skilled) to do the job. Third, a high product diversity indicates that a variety of skills may be required in the repair process, as mentioned earlier. Based on these indications, repair shops 2.1, 2.4, 2.8, 2.9 and 3.3 require a high number of specialized skills in the repair process.

5.4 Analysis & interpretation

In Chapter 4 we have reviewed three streams of literature literature: uncapacitated inventory models (stream 1), single-stage, capacitated inventory models (stream 2) and multi-stage, capacitated inventory models (stream 3). We have formulated requirements when the literature *could* (application requirements) and when the literature *should* (control principles) be applied in designing control concepts for repairable parts inventory systems. In this section we evaluate the applicability of

[7]The number of repair men available for shop level maintenance cannot exactly be determined for the repair shops of company 3, because these repair men also conduct asset (line) maintenance.

control concepts from the three streams of literature, by comparing the application and control principles to the quantitative and qualitative findings from the previous section.

Literature stream 1: uncapacitated inventory models

Table 5.3 summarizes our observations related to the first stream of literature. The second column contains a list of variables that served as indicators for our observations. From these observations, we can conclude that a varying set of case study repair shops (more or less) complies to each of the application requirements, however no single repair shop complies to all of them.

Table 5.3 Observations from case studies on the applicability of literature stream 1

Application requirement	Application indicators	Observation
All required resources have relatively low utilization rates, or capacity is sufficiently flexible in terms of additional lead time and costs	- Low CV of NRFU part returns (cv_j^d) - Low utilization rate - Short subcontracting lead times - Unrestricted use of overtime	In general, for all shops capacity is rather inflexible due to a high utilization rate, long subcontracting lead times and overtime restrictions per period. The average CV of NRFU part returns per week over all shops is high (0.49). Shops with a relatively low CV are 2.2, 2.3, 2.4, 2.6 and 2.11; these shops are most likely the ones in which overtime can successfully be applied as a capacity management tool, without introducing a high additional lead time and costs.
At most one SRU fails upon failure of a next higher indenture level part	Low average number of SRUs required per repair job (μ_j)	The average number of required SRUs is higher than 2 for all repair shops, except for shops 1.2, 2.1, 2.2, 2.5 and 2.11.
Failed parts are inspected upon failure, at no loss of actual repair time	- Low fraction of job service time spent on inspection (q_j^i) - Low service time per part in hours (μ_j^s)	Inspection takes a significant amount of time. Exceptions are shops 2.3, 2.4, 2.7 and 2.12, and possibly some of the repair shops for which we were not able to determine these figures.

The first observation shows that for most repair shops, it seems worthwhile to integrate decisions related to repair capacity dimensioning and inventory control. This is motivated by the fact that waiting times for capacity can be substantial when the utilization rate and workload volatility are high, and capacity flexibility is low. In that case application of the control principles of literature streams 2 and 3 is beneficial.

The second observation is in line with Van Jaarsveld et al. (2015), who also argue the assumption of at most one SRU failure as the cause of a failure of a next higher indenture level part. Like we mentioned in Section 4.6, this assumption can be justified if only relatively slow moving, expensive SRUs are included in the model (and we implicitly assume that all other SRUs are always available). We are however not able to validate this assumption, due to a lack of accurate data[8]. We would however suspect that the total value of all SRUs in an LRU does not (much) exceed the value of the LRU itself, hence indicating that only a small number of expensive SRUs can fail upon failure of an LRU. Moreover, we suspect that LRUs are designed such that a simultaneous failure of two expensive SRUs only rarely happens.

It is due to the fact mentioned above that we are also not able to assess the control principle of literature stream 1 (i.e. integrating inventory control of LRUs and SRUs significantly reduces inventory holding costs, especially when multiple, expensive SRUs with a low demand rate may be required in the repair process). From the second observation, we can however conclude that this control principle need not be applied to all repair shops.

Literature stream 2: single-stage, capacitated inventory models

Table 5.4 summarizes our observations related to the second stream of literature. From these observations, again we can conclude that a varying set of case study repair shops (more or less) complies to each of the application requirements, however no single repair shop complies to all of them. We emphasize that either the third or the fourth requirement has to be met for this literature stream to be applicable. Therefore, the computational complexity of the problem mostly restricts the applicability of this literature stream.

We can assess the control principle of literature stream 2 (i.e. integrating scheduling policy and stocking decisions significantly reduces inventory holding costs, especially when large demand and price differences between parts exist) by analyzing the value

[8] Company 1 only registers repair (not new buy) prices for the LRUs. In companies 2 and 3, demand registration for repairable SRUs (which are in general the most expensive SRUs) is not done properly, or not at all due to the fact that a "repair-by-repair" strategy is applied for these repairable SRUs.

Table 5.4 Observations from case studies on the applicability of literature stream 2

Application requirement	Application indicators	Observation
Closed-loop setting	Low fraction of unsuccessful repairs (q_j^c)	A low number of repairable parts is scrapped in the repair process, with a weighted average of 3.3% over all repair shops. Exceptions are shops 2.1, 2.2, 2.5, 2.7, 2.8 and 2.14
Repair process consists of a single stage	- Homogeneous repair capacity - Job shop layout	Repair shops 1.2, 1.3, 1.4, 2.2, 2.6, 2.7, 2.14, 2.15 and 3.1 have a job shop layout with a homogeneous repair capacity
SRUs cause no delay of repairs *or:*	- Small set of materials potentially required per repair job (μ_j^m) - Low average number of SRUs required per repair job (μ_j) - High fraction of repair jobs for which no SRUs are required (p_j^0)	Based on the mentioned indicators, the probability that SRUs cause a delay of repairs is relatively small in repair shops 1.1, 1.2, 1.3, 1.4, 2.1, 2.2, 2.5, 2.7, 2.8, 2.11, 1.12, 2.13, 2.14, 3.2
Failed parts are inspected upon failure, at no loss of actual repair time	- Low fraction of job service time spent on inspection (q_j^i) - Low service time per part in hours (μ_j^s)	Inspection takes a significant amount of time. Exceptions are shops 2.3, 2.4, 2.7 and 2.12, and possibly some of the repair shops for which we were not able to determine these figures.
Small problems due to high computational complexity	- Low number of different parts that is repaired (n_j) - Small set of materials potentially required per repair job (μ_j^m)	If both LRUs and SRUs are considered, then this applies to repair shop 2.12. When only LRUs are considered, this also applies to repair shops 1.1, 2.10, 2.11 and 3.1.

of the parameter r_j^d: the smallest fraction of parts responsible for (at least) 80% of the total number of returned NRFU parts. The smaller this fraction, the larger the differences in the individual part demand rates. Applying this control principle seems to be most beneficial for repair shops 2.4, 2.14, 1.1, 2.6 and 2.7[9].

Literature stream 3: Generalized Kanban Control Systems

Literature stream 3 focuses on Generalized Kanban Control Systems (GKCS). As mentioned in Section 4.5, the stages in these models can be interpreted in three ways. We hereafter discuss if and how GKCS can be applied to the repairable parts inventory systems in our case study companies.

[9] Due to a lack of accurate price information, we are not able to assess the price differences between the LRUs.

First, the stages can be regarded as the indenture levels of repairable parts. The final stage of the GKCS is the repair, assembly and test phase of an LRU, and each preceding phase the repair, assembly and final test phase of the lower indenture level parts. The first stage is the disassembly and inspection phase of the LRUs and SRUs. This setting is applicable to repairable parts inventory systems if both inspection and the repair and assembly of parts of each indenture level require separate resources.

In the case study repair shops, this setting applies to repair shops where parts with a high technological complexity (and multiple indenture levels) are maintained by replacing failed SRUs by ready-for-use SRUs from stock. This setting mainly applies to repair shops 2.3, 2.4, 2.7 and 3.3. We observed that the repair men in these shops maintain both LRUs and SRUs, whereas GKCS in this formulation requires that LRUs and SRUs are maintained by separate resources. Hence applicability seems to be limited.

Second, the stages can be regarded as the specific resources that are required in the repair process. In this setting, the repair routing is assumed to be fixed for each part. This setting complies to the case study repair shops with a product or process layout. Application of GKCS to repairable parts inventory systems, in which stages are regarded as specific resources, is hence limited to repair shops 2.10 and 2.11.

Third, the stages can be regarded as the phases in the repair process. To apply GKCS, this requires the assumption that each repair phase has dedicated resources. This in general does not hold for the case study repair shops. Often a single repair man conducts (the greater part of) a repair job. This implies that repair phases have to be combined in order to apply GKCS. For example, the inspection and repair phase can be combined if a separate resource (a higher skilled repair man using specific test equipment) is only required in the final test phase. This setting was found in repair shops 2.3, 2.4 and 2.5. Merging the inspection and repair phase however requires the assumption that SRUs cause no delays of repairs. We refer to Table 5.4 to see for which repair shops this applies.

We assess the control principle of literature stream 3 (i.e. integrating inventory and workload control on each stage of the repair process significantly reduces costs, especially when the repair process consists of multiple stages with dedicated resources) for the setting in which stages are regarded as the phases of the repair process. The assessment is based on the qualitative figures on the number of skills required in the repair process (see Section 5.3). Applying this control principle seems to be most beneficial for repair shops 2.1, 2.4, 2.8, 2.9 and 3.3.

5.5 Research limitations and implications

Clearly, some of the figures and findings from the case studies are correlated to factors such as the size of the company (e.g. repair volumes depend on the number of assets in use/maintenance) and the maintenance strategy that is applied (repair vs. overhaul of a part). The way the repair process is designed also implicitly depends on these factors; e.g. higher repair volumes enable (in general) a more efficient repair process.

Nevertheless, we believe that these case studies provide a good insight in the repair shop (and process) dynamics that can be observed in general. All characteristics related to the uncertainty of the required SRUs are measured at part level, and are aggregated to obtain the necessary statistics on repair shop level. At part level we think it is indeed possible to compare the characteristics of the various shops.

The characteristics related to the capacities required in the repair process are observed at the repair shop level. In fact, the capacity dimensioning decisions (including the variety of available skills) are strategic choices made by the companies, and taken as a given in this case study research. We should be careful in our 'observation' that we see clear differences between repair shops when it comes to the number of skills required in the repair process, as we may observe these differences only as a result of the strategic choices the companies have made.

Nevertheless, we have the following reasons to believe that our observations can be generalized. First, we recall that the three companies have a different motivation for having internal repair shops (see Section 5.2). Second, we can draw the same conclusion from observing only the repair shops of company 2. Third, we triangulated the results related to the capacity characteristics in the three other case study companies, which confirmed that our observations are in line with what is observed in the repair shops of these companies.

5.6 Concluding remarks

In this chapter we have used case studies to answer the fourth research question:

4. *Which models and control principles, identified in the literature review of Chapter 4, are applicable to design control concepts for the repairable parts inventory systems in the case study companies?*

First, we can conclude that there is no repair shop to which the models and control principles can be applied directly, because at least one important assumption is not valid. This however does not mean that the streams of literature cannot be used in designing repairable parts inventory systems. How and when the literature can be used is discussed in Chapter 6.

Second, we can conclude that a varying set of repair shops complies to a subset of requirements and control principles per literature stream. This can be explained by the clear differences we observe between repair shops, especially when it comes to the complexity and uncertainty of the materials and capacities required in the repair process. How this impacts the designs of control concepts for repairable parts inventory systems is not clear, and is discussed in Chapter 6.

The application of the literature is mainly limited by the assumption that inspection takes no time, and repair work is preempted in order to directly inspect a part upon its failure. In that case, the number and the impact of actual repair job interruptions is underestimated and consequently, performance metrics (such as spare parts availability and holding costs) are overestimated, when compared with practice. Other models assume that the inspection and repair phase are conducted by different resources, and do not explicitly model the inefficiency arising from information transfers.

In this chapter, we mentioned several aspects that are not yet fully covered in the literature when it comes to the modeling of the repair process. The following list summarizes these aspects: (1) assumption of i.i.d. repair lead times in uncapacitated inventory models is not always valid, for example if the repair shop utilization is high, (2) inspection of a part takes a significant amount of time and should explicitly be included in the models, (3) in many situations, the inspection and repair phase require the same capacity, (4) repair times are longer in case a repair job is preempted, (5) multiple SRUs may be required in the repair process per failed higher indenture level part, and (6) a repair job in general does not have a fixed routing and is known only after the inspection phase. These aspects will be discussed in the next chapters.

Chapter 6

Connecting inventory and repair shop control

"Hoe simpeler hoe beter, want hoe minder keuze je een speler laat hoe meer kans dat hij het juiste doet."

Johan Cruijff

6.1 Introduction

From the literature review in Chapter 4 we have found three ways to design control concepts for repairable parts inventory systems, each coupled to a specific literature stream. The applicability of each design is limited by the characteristics of a specific situation, and from the case studies in Chapter 5 we have found indications that the design should be tailored to each specific repair shop. This chapter elaborates on the previous chapters and addresses the fifth research question of this dissertation:

5. *To what extent do repair shops differ from each other, and what implications do these differences have on the design of control concepts for repairable parts inventory systems?*

Most authors (e.g. Hausman and Scudder (1982) and Guide Jr et al. (2000)) compare the repair shop setting to the traditional job shop setting as used in typologies for production units (Bertrand and Wortmann, 1992; Bertrand et al., 1990). To typify all repair shops as job shops is too simple. There are reasons to assume that planning and controlling repair shops differs significantly from traditional production units, although it is not clear what distinguishes them from a planning and control point of view. There are well-established typologies for production units (Bertrand and Wortmann, 1992; Bertrand et al., 1990); however, these typologies do not allow for different types of repair shops to be recognized. In this chapter, we present a typology of repair shops which is based on shops having similar characteristics.

This typology can be used to determine which control concept (and control principles) should be applied to which type of repair shop. We present designs for these concepts, this being the second contribution of this chapter. We describe the designs in detail and we (i) map the various decision functions on the interface between the inventory control department and the repair shop, (ii) assign each decision function to either the inventory control department or the repair shop and (iii) refer to literature that can be used to apply these control concepts in practice. The control concepts are based on an extensive design process. We present our design principles and refer to results from Chapters 4 and 5 to support our design choices.

The control concepts may serve as a useful starting point for creating control concepts for repairable parts inventory systems. For companies with existing control concepts, the proposed control concepts have a mirror function. That is, they serve as a reference point to redesign existing control concepts of repairable parts inventory systems. The use of these control concepts as a reference point for designing control concepts of existing repairable parts inventory systems, is explored in focus groups at two companies.

Figure 6.1 gives an outline of the steps towards the designs of the control concepts, including their relation with previous chapters. In Section 6.2 we describe our design process and design choices, and in Section 6.3 we describe the different repair shop types observed in practice. In Section 6.4, we present control concepts for all repair shop types and illustrate how the literature discussed in Chapter 4 can be applied. In Section 6.5 we report on findings from two focus groups, in which we used our control concepts as a reference point for redesigning existing control concepts of repairable parts inventory systems. Section 6.6 concludes this chapter.

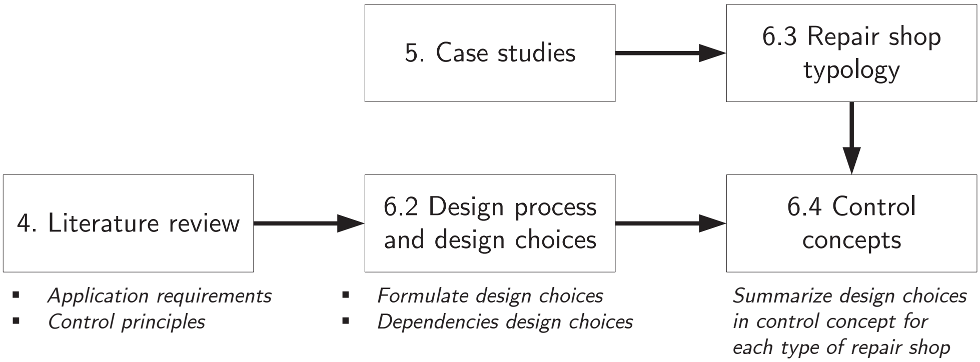

Figure 6.1 *Steps towards the design of the control concepts.*

6.2 Design process and design choices

In general, design starts with specifying the *functional requirements* (objectives) and constraints for the system that is to be designed (see Suh (1990)). The design process involves the search for a system, specified by the *design parameters*, that satisfies the functional requirements. The functional requirements of a repairable parts inventory system can be expressed in terms of *efficiency* and *effectiveness*. The efficiency of a repairable parts inventory system can be expressed in the utilization of the shop's resources, the work-in-process that resides in the shop and the inventory investments in LRUs and SRUs. The effectiveness of the system can be expressed in terms of the minimum or average number of backorders for (repairable) LRUs per asset type. The design parameters of the repairable parts inventory systems are the relevant decision functions (see Section 1.2).

According to Suh (1990), a good design satisfies the following two *design axioms*: (i) the *independence axiom*: maintain the independence of the functional requirements and (ii) the *information axiom*: minimize the information content. Ideally, one might want to optimally solve a model that encompasses all the interrelated decisions in repairable parts inventory systems (i.e. integrated decision making). The design space for repairable parts inventory systems can however be extremely large for real-life problems and many dependencies exist between the decision functions. The first step in the design process is therefore to construct an interface that separates two arbitrary sets of decision functions, that can be modeled and solved independently. This decomposition of decision functions is reasonable if it can be done without introducing a large 'optimality gap' (e.g. a relatively large increase of total costs) and without introducing a high coordination load (and information transfer) between the two decomposed sets of decision functions, or if the computational complexity to solve the design problem is too high.

This decomposition of the design problem requires four design choices. First, with this decomposition into two sets of decision functions an *interface agreement* between them is introduced, and this interface agreement needs to be designed. The other three choices focus on how the decomposition should be designed, and relate to the three control principles identified in Chapter 4. The four design choices are outlined and discussed in Sections 6.2.1-6.2.4, and should in the first place contribute to the objectives of a repairable parts inventory system. Further, we use the following 'guiding principles' in making the design choices:

1. The interface agreement should be part of the design.
2. The coordination between (or integration of) interrelated decision functions should be part of the design.
3. The degree of coordination between individual decision makers should be minimal.
4. The design should be as simple as possible.

6.2.1 Interface between inventory control and repair shop control

In all the case study companies, the decisions in repairable parts inventory systems are divided over two departments: an inventory control department and a repair shop (see Section 1.2.2). In the framework we presented in Chapter 2, the decisions in these systems are related to repair shop control (dimensioning fixed and flexible repair capacities and scheduling repair jobs) and inventory control (setting replenishment policy parameters for SRUs and LRUs), and supply management acts as an *interface*. How to arrange the interface is far from clear. The usual decoupling assumed in the literature, just as in the framework in Chapter 2, is through lead times. However, for efficient control of repair job priorities, repair shop control and inventory control are sometimes assumed to be fully integrated (e.g. Hausman and Scudder (1982)). The design of the interface agreement depends on how the two sets of decisions are decomposed, and hence depends on the next three design choices.

6.2.2 Decoupling of the inspection and the repair phase

In repairable inventory systems, work preparation is part of the repair job (inspection phase). Hence, only after the inspection phase, it is known which resources are required to complete the repair job. This complicates the interface between inventory control and the repair shops regarding the release of repair jobs, as timely completion of repair jobs by the repair shop is influenced by the availability of SRUs.

To simplify the interface between the repair shop and the system that supplies the SRUs, it is beneficial to add a decoupling point between the inspection and the repair phase of repair jobs (as here a large shift in uncertainty occurs). It is however less efficient to interrupt repair jobs after the inspection phase (see Chapter 5). Hence, in designing control concepts for repairable parts inventory systems we need to decide on decoupling the inspection and repair phase, and we need to take the inefficiency of decoupling into account.

This design choice strongly relates to the design of the interface between inventory control and the repair shop. Suppose we decide to decouple the inspection and repair phase of repair jobs, then inspected parts are stocked after the inspection phase and are only released for repair when all the required SRUs are available. In this case, no interface is required between the repair shop and the system that supplies the SRUs. If we do not decouple the inspection and repair phase, then we would require an interface or the SRU stocking decisions are taken at the repair shop level. Hence this design choice to decouple the inspection and repair phase, also strongly relates to the following design choice.

6.2.3 Decoupling of LRU and SRU stocking decisions

Clearly, the stocking decisions of LRUs and SRUs interact and integrating these decisions reduces the stock investments (balance the holding costs of SRUs and LRUs). Following the discussion on the previous design choice, integrating the SRU and LRU stocking decision is only desired in case the inspection and repair phase are decoupled. If the latter is not the case, then the repair shop is responsible for the complete repair (lead time), including the waiting time for SRUs, and SRU stocking decisions are taken at the repair shop level. An alternative is an extra interface between the repair shop and the system that supplies the SRUs, however this is not desirable from a design point of view (minimize information content). Decoupling stocking decisions of LRUs and SRUs seems reasonable in situations in which it is relatively cheap to arrange for a high availability of SRUs, and hence SRUs cause no delays of repair jobs.

6.2.4 Repair job release and resource control

The last design task is to establish a mechanism for releasing jobs to the repair shops and for controlling the flow of repairable parts through the shops. When multiple repair jobs require the same resources, the latter might become bottlenecks and significantly influence the repair lead time. Therefore, we will not assume scheduling to be done in the repair shop at all times; instead we see scheduling as something that can be done either in the repair shop itself or on a higher level (i.e. local vs. central). The design of this mechanism depends on the type of resources required in the repair process.

Within a repair shop, we can distinguish generic and specialized resources, such as specialized test equipment and repair men having different skill levels. The release of repair jobs requiring specialized capacities should be leveled with the specialized repair capacity. However, if specialized repair men only conduct a small percentage of their available time on repair jobs requiring specialized skills, then capacity leveling based on aggregate flows could be sufficient. Repair shops can have resources that require a high-capacity utilization and are inflexible, such as an expensive machine. These type of resources create dependencies between the repair jobs requiring this resource. The allocation of resources should be leveled with the aggregate release of jobs requiring this resource (load leveling per resource), in order to avoid capacity losses and poor performance. Other resources may have a high capacity flexibility or lower utilization levels. For the latter group of resources, aggregate flows of work are sufficient to coordinate the resource requirements.

6.3 Typology for repair shops

This section presents our typology for repair shops. Starting point for setting up this typology is the typology for production units (Bertrand and Wortmann (1992); Bertrand et al. (1990)), which is discussed in §6.3.1. §6.3.2 describes the dimensions of our repair shop typology, including their specific characteristics.

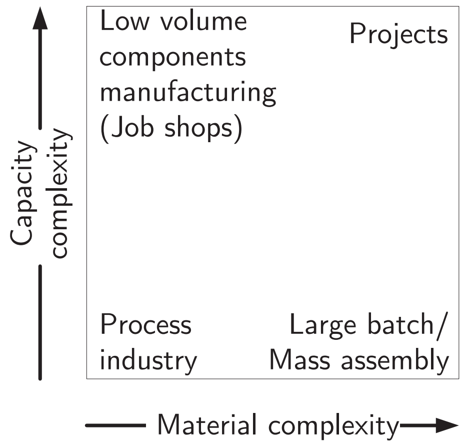

Figure 6.2 *Production unit typology (adapted from Bertrand and Wortmann (1992))*

6.3.1 Production unit typology

As a repair shop is a special kind of production unit, the starting point for a repair shop typology for planning and control purposes is the model presented by Bertrand and Wortmann (1992) for production units. This typology makes a distinction between capacity complexity and material complexity. Uncertainty is added to this typology by considering a project industry shop as much more uncertain than the other types. Figure 6.2 visualizes the typology of production units.

The process industry shop manufactures a limited number of similar product types, in high volumes. This shop typically has a small number of production phases, a low volume flexibility and high changeover costs. Batch sizes and production plans which take setup times into account should be used.

In the mass assembly shop many separate production lines converge at the main assembly line, that has a low volume flexibility. The material coordination is complex due to the assembly nature of the production process. The planning of the production phases is based on flow concepts and are often decoupled from the final assembly line production planning.

In the low volume components manufacturing shop (job shop), products consist of only a small number of materials. The manufacturing of the products requires only a small number of phases but a wide variety of skills though. Key here is to couple the planning of specific capacities to the specific demand for products.

In the project industry products are made on a make-to-order basis. These products require many different technologies and many manufacturing phases. The main issue in project shops is the uncertainty involved in the production process. A rough capacity plan is required, as the many changes in the bill of material of the products require many inefficient rescheduling of a detailed planning.

6.3.2 Repair shop typology: dimensions and characteristics

Several authors compare repair shops to the traditional job shop setting (e.g. Hausman and Scudder (1982) and Guide Jr et al. (2000)). The reason for this is that the number of required materials is usually low, usually specific capacities are required and the shop itself is organized according to functional groups. Routes through the shop fluctuate and there is no fixed sequence of operations, such as in a flow shop. Products are discrete in nature and there are no (semi-)batch characteristics such as linked resources or tanks.

Nevertheless, from the case studies (see Chapter 5) we can see a clear distinction between repair shops with a lower and higher capacity complexity. The *capacity complexity* mainly concerns the requirement of specialized skills of repair men to execute a repair job. This means that the capacity complexity in a repair shop is a useful dimension in a repair shop typology, similar to Bertrand and Wortmann (1992); Bertrand et al. (1990). However, this is not the case for material complexity: material requirements are usually unknown before the inspection phase, lead times for materials are usually long, though only a small number of materials is required in the repair process. Hence this dimension is quite similar for all repair shops and can be typified as moderately complex. This means that when we would use the typology of Bertrand and Wortmann (1992); Bertrand et al. (1990) for repair shops, some categories would never be used. In other words, the typology lacks a descriptive differentiation for repair shops regarding the material dimension.

The typology from Bertrand is based on a dimension related to capacity and a dimension based on material. This seems logical from a planning and control point of view, as planning is the coordination of capacity and material in time. This means that planning complexity must be related to material and capacity complexity. However, as indicated above, the material complexity is not varying to the same extent as it can be the case in 'normal' production units, where shops may produce simple items or

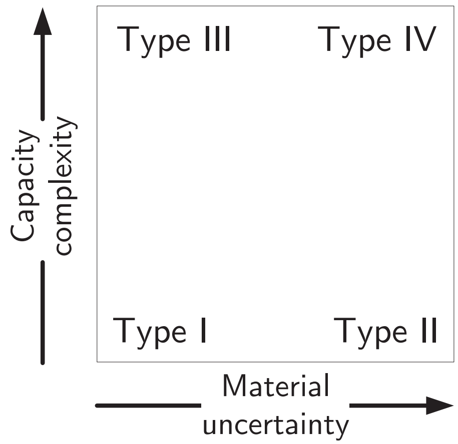

Figure 6.3 *Typology of repair shops for repairable parts.*

complex products such as trucks, or discrete items or process oriented items, such as paint. Instead, the material complexity in repair shops is mostly caused by *material uncertainty*: the extent to which repair jobs for the same item require different sets of materials, i.e. the extent to which it is predictable which materials will be required in the repair process. Contrary to production units, it is not known in advance which materials will be required. The material uncertainty determines whether a repairable item can be repaired in an uninterrupted manner, or whether the repair process has to be paused because some SRU is not available.

This two-dimensional classification results in the typology of Figure 6.3, which shows four basic types of repair shops. Table 6.1 summarizes the characteristics regarding capacity complexity and material uncertainty. We briefly discuss these characteristics hereafter.

In repair shops with a high capacity complexity, the repair job lead time consists to a large extent of waiting time for the required (specialized) resources. Repair shops with a high capacity complexity experience shifting bottlenecks, especially for specialized repair men and high-valued equipment that is required in the repair process, and they often have a specific planning for the bottleneck capacity. Sometimes batching policies can be applied to decrease the average setup times for specific capacities and skills.

In repair shops with a high material uncertainty, repair jobs are often interrupted due

Table 6.1 Characteristics of a low/high capacity complexity and material uncertainty

	Low	**High**
Capacity complexity	Homogenous repair capacity, low number of specialized skills required Narrow product mix, i.e. low diversity in the type of items to repair	Heterogeneous repair capacity, high number of specialized skills required Wide product mix, i.e. high diversity in the type of items to repair
Material uncertainty	Often the same SRUs required for repair jobs of the same LRU Small set of SRUs that has to be kept on stock	The sets of SRUs that are required for repair jobs of the same LRU vary strongly Large set of SRUs that has to be kept on stock

to missing materials. The repair job lead time consists to a large extent of waiting time for materials. Repair jobs for the same LRU may require different sets of SRUs, and in general the number of SRUs that have to kept on stock is high. Moreover, the SRUs can in fact also be repairable parts, that may also be repaired by the same repair shop.

Figure 6.4 gives the classification of the case study repair shops that were introduced in Section 5.2. This classification is based on the empirical results presented in Chapter 5, and we can conclude that all repair shop types can be found in practice. The indicators related to the application requirement "SRUs cause no delay of repair jobs" are used to classify repair shops as having a low material uncertainty (see Table 5.4). Repair shops with a heterogeneous repair capacity and a high product diversity are classified as having a high capacity complexity, because these factors indicate that a high diversity of skills may be required in the repair process[1]. We also classified repair shop 3.2 as having a high capacity complexity, because it both has a relatively high repair time per part and a high coefficient of variation of the repair time. These latter indicators may also point at a high diversity in repair routings and a high number of different required skills.

[1]For the motivation we refer to Section 5.3.

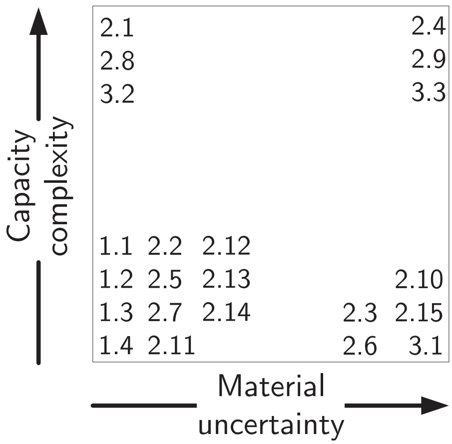

Figure 6.4 *Classification of case study repair shops.*

6.4 Control concepts for repairable parts inventory systems

In this section, we present designs of the control concepts for repairable parts inventory systems for each of the repair shop types presented in the previous section. We take the description of the repairable parts inventory system in Section 1.2 as a reference and, for ease of clarity, we use a somewhat simplified version of the repair process (see Figure 6.5). We assume that replaced parts (LRUs and SRUs) in fact always *failed* and that repair is always possible, that is, both *no fault found* and *condemnation* (or scrap) are excluded[2].

In the remainder of this section, we outline the design of the control concepts for all repair shop types of the typology, presented in Section 6.3. We combine Figures 6.5 and 1.3 and assign the decision functions in Figure 1.3 to the repair shop or to the inventory control department. For each repair shop type, we outline (i) the characteristics and design principles, (ii) the control concept design, and (iii) a review of available literature, including suggestions for further research.

[2]Repair shops with a high 'no fault found' or scrap rate are, from a control perspective, comparable to repair shops with a high material uncertainty. Therefore we do not account for scrap of parts and 'no fault found' in the design of the various control concepts. Nevertheless, we describe how the presented control concepts can be modified to be tailored to these excluded situations.

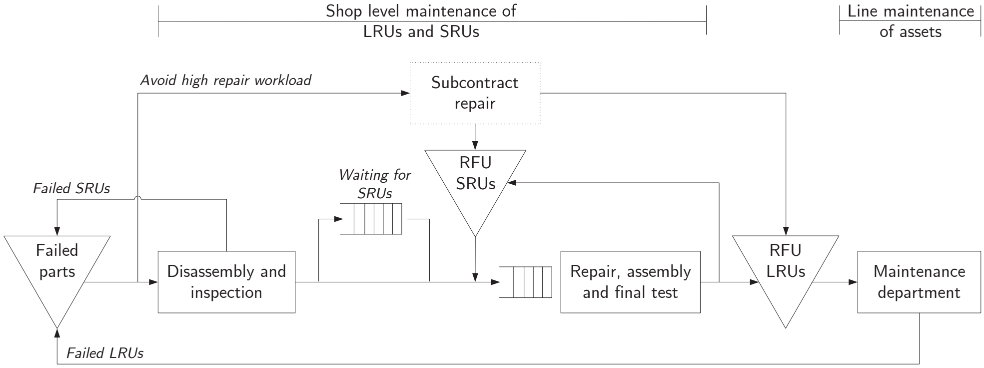

Figure 6.5 *Assumed supply chain of repairable spare parts including the repair process.*

6.4.1 Repair shop type I

6.4.1.1 Characteristics and design principles

Repair shop type I is characterized by a low capacity complexity and a low material uncertainty. Usually this type of repair shop repairs a small number of different parts with a low diversity. The added value of the repairs is typically low because no expensive materials and specialized skills are required, and utilization rates will be high. Typically a relatively small number of parts is responsible for a large portion of the repair work of the repair shop. Examples are the emergency bottles shop or the wheels & tires shop at one of our case study companies.

A low material uncertainty means that no SRUs or often the same set of SRUs is required. The number of SRUs to stock is low and it is easy and relatively cheap to arrange for a high availability of the SRUs. The benefits of integrating SRU and LRU inventory control are low. The repair capacity is homogeneous and the number of specialized skills is low. This means that the repair job release is only influenced by the aggregate release plan and not by specific capacity aspects, and repair jobs can be released as soon as (repair man) capacity becomes available.

The choice of how to set up the interface between the inventory control department and the repair shop depends on how the repair process is organized. Typically this shop type has multiple identical servers (repair men) that are able to conduct all types of repairs, generally without interruptions and with short setup and changeover times. For this reason, the WIP is usually low and repair lead times can be kept short (equal to the effective repair time).

The inventory control department has the best knowledge about which repair jobs to prioritize from an inventory control perspective. Moreover, the inventory control department is best able to balance the costs of regular repairs, overtime and external repairs. Since the capacity complexity is low, the priority from a repair capacity perspective (i.e. to optimize repair capacity utilization) is less important. For these reasons, the inventory control department should be able to set the repair job priorities. In this control concept, the inventory control department 'pushes' the repair jobs into the repair shop.

6.4.1.2 Control concept

We propose to set a target level on the workload (in hours of work) that is cleared by the repair shop each period (e.g. one week), given the constraint that sufficient work is available for the shop. This target level is jointly determined with the fixed and flexible (overtime or external) capacity decisions and the LRU base stock levels at the aggregate level. It is the responsibility of the repair shop to timely clear the shop, i.e. to make sure that the (cumulative) backlog based on the target levels is below a certain threshold.

The repair shop and the inventory control department agree on (i) the maximum amount of cumulative backlog and (ii) the maximum amount of work that is released to the repair shop per period. The process and control concept design of repair shop type I can be found in Figure 6.6.

The low number of repair job interruptions, combined with a restriction on the maximum amount of work released to the repair shop, enables the repair shop to ensure a high probability that the backlog is below the maximum threshold value. We propose to use a constraint (input rule) for the weekly offered workload (in hours). The workload constraint is a trade-off between inventory capacity, fixed and flexible (overtime) repair capacity and subcontracting options.

As the repair shop is responsible for timely clearing the workload, and the benefits of integrating SRU and LRU inventory control are low, it is rather obvious that the repair shop is responsible for the inventory control of the SRUs. In this way, the repair shop has full responsibility for timely clearing the workload.

At the aggregate control level, a queuing type analysis can be used to determine the

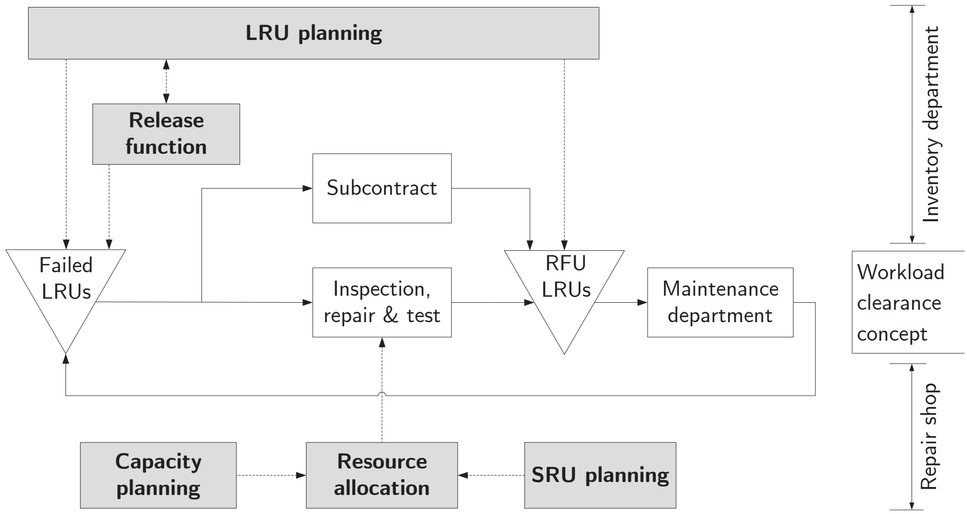

Figure 6.6 *Process and control concept design for repair shop type I*

fixed and flexible repair capacities, the LRU base stock levels and the budget for subcontracting repairs. In this analysis, one should take into account a scheduling rule that accounts for the differences between inventory holding costs of various LRUs. This analysis implicitly also determines the threshold value for subcontracting repair jobs, that is, repair jobs are subcontracted as long as the amount of repair work in the "Failed LRUs" stock exceeds this threshold value.

Detailed scheduling of repair jobs is not required, and the repair jobs can be released by the inventory control department as soon as repair man capacity becomes available. Repair jobs are released based on the run-out times of the LRUs (i.e. dynamic priority rules). Note that in this design, emergency repair jobs do not really affect the repair shop planning. Resource allocation is easy as the repair capacity is homogenous and both repair man capacity and SRUs may be allocated first-come-first-serve to the repair jobs.

Alternatively, sometimes we see that type I repair shops face relatively high setup or changeover times for repair. An example is the emergency bottles shop at one of our case study companies. For this shop, it is important to account for the changeover times between repair of different parts, in order for the repair shop to efficiently utilize its resources. In such cases, it is sometimes preferred to work with static priority rules so that the repair shop itself can determine when to switch the repair from one LRU to another.

6.4.1.3 Literature review and suggestions for further research

The interface between the inventory control department and the repair shop is based on a workload clearance concept, in which the workload constraint is a trade-off between the number of LRUs to stock, the repair shop capacity including possible overtime policies, and subcontracting options. This trade-off is an integrated (long-term) decision made by the inventory control department and the repair shop.

Table 6.2 provides references to literature on decision support models for repair shop type I. We describe how to make/support the decision, how the available literature fits the need for decision support and provide suggestions for further research[3].

6.4.2 Repair shop type II

The control concept design of repair shop type II is not much unlike the control concept of repair shop type I. For this reason, we only describe the main differences between both types, including the implications for the design of the control concept.

6.4.2.1 Characteristics and design principles

Repair shop type II differs from repair shop type I in that it is characterized by a high material uncertainty. The added value of repairs mainly consists of the value of the materials used in the repair process. Repair jobs are conducted one by one; batching of repair jobs of the same parts is more difficult because different SRUs are required in different repair jobs. An example is the electronics repair shop in one of our case study companies.

A high material uncertainty implies that often different sets of SRUs are required for repair jobs of the same LRU and the number of SRUs to stock is high. The number of repair job interruptions is moderate and mainly due to missing SRUs. The benefits of integrating SRU and LRU inventory control can be substantial (Muckstadt (1973)).

Before the inspection phase of a repair job is completed, there is a large uncertainty on which SRUs are required in the repair process. In general, decoupling points are

[3]Research to apply batching policies (of e.g. repairs, return or failed parts) is beyond the scope of this chapter.

Table 6.2 Literature review and suggestions for further research for repair shop type I.

Decision function	**Responsible**	**Suggestions and available decision support**	**Open research topics**
Workload constraint	Inventory control department and repair shop	Trade-off repair and inventory capacity: Sleptchenko et al. (2003).	Extend with the option of working overtime and subcontracting repairs, and use a different scheduling policy.
Workload clearance levels	Inventory control department and repair shop	Set equal to workload constraint.	
Repair shop capacity planning	Repair shop	Follows directly from workload constraint	
LRU planning	Inventory control department	*Flexible repair capacity*: (Variants of) METRIC, e.g. Sherbrooke (1968, 1986) *Inflexible repair capacity*: Adan et al. (2009), Tiemessen et al. (2017), (Arts et al., 2016; Arts, 2017)	Model with expediting and subcontracting repairs
SRU planning	Repair shop	Single-item, single location inventory policies, see e.g. Silver et al. (1998)	
Release function	Inventory control department	Real-time rolling horizon model: Caggiano et al. (2006)	Include option of subcontracting and working overtime
Resource allocation	Repair shop	FCFS allocation of repair capacity. Apply reactive overtime policy, see e.g. Scudder and Chua (1987)	

introduced at points in a repair (or a manufacturing) process where the predictability changes significantly (Bertrand et al. (1990)). For this reason, we decouple the inspection phase and the repair phase of the repair jobs via a stocking point of parts that are "Ready to repair"[4].

[4]This also applies to repair shops with a high 'condemnation' rate, because only after the inspection phase it is known whether the failed part is still repairable, the failed part being one of the resources required to complete the repair job. For repair shops with a high 'no fault found' rate, this may apply to the decoupling of the *pre-inspection* and the *disassembly and inspection* phase, see Figure 5.1.

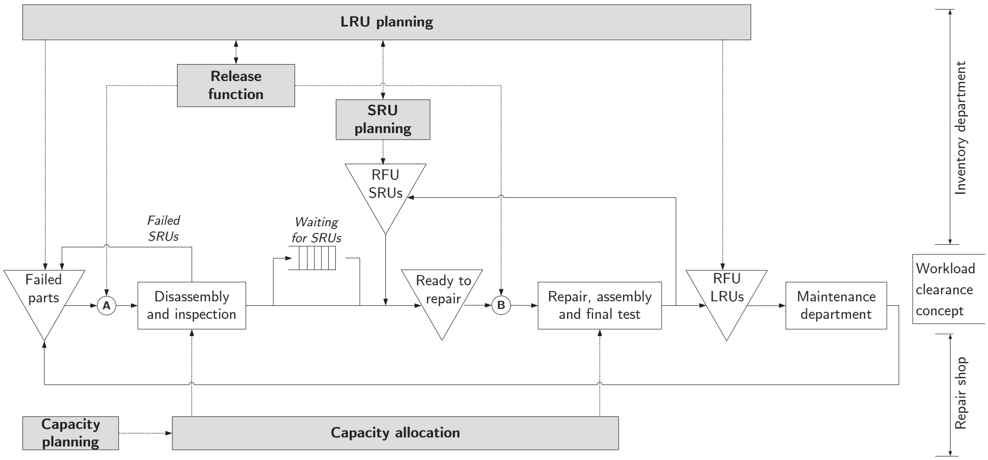

Figure 6.7 *Process and control concept design for repair shop type II*

To minimize the waiting time for SRUs, it is important that the inspection phase of a repair job starts as quickly as possible after the part enters the "Failed LRUs" stock. The waiting time for parts that are ready to be repaired will increase though, if inspection of parts has a higher priority than the repair of parts. For repair shop type II, it is important to be able to find the optimal balance between inspection of parts (to decrease the waiting time of SRUs) and repair of parts (to decrease the waiting time of LRUs). It is the task of the inventory control department to dynamically set the right priorities to inspection or repair of the right parts (based on actual stock levels of both LRUs and SRUs).

6.4.2.2 Control concept

Similar to the design for repair shop type I, we propose to use a target level on the total workload that is cleared by the repair shop each period. In contrast to the design of repair shop type I, the cleared workload may consist of jobs in which only the inspection of a part is conducted. The inspection and repair phase of a repair job are decoupled, in order to be able to dynamically set the right priority to the inspection or the repair of parts. In this control concept, the inventory control department 'pushes' the repair jobs into the repair shop. The control concept design for repair shop type II is visualized in Figure 6.7.

The inventory control department is responsible for the SRU and LRU inventory control. With an integrated approach the inventory control department is able to find the optimal balance between the (average) waiting time for SRUs, that directly influences the waiting time to repair a part and hence inventory levels of the LRUs, and the inventory levels of the SRUs. The inventory control department can influence the time a part spends in the "Failed LRUs" stock and the "Ready to repair" stock, by assigning the desired repair job priorities.

At the aggregate control level, a queuing type analysis can be used to determine the fixed and flexible repair capacities and the LRU and SRU base stock levels. In this analysis, one should take into account a scheduling rule that accounts for the differences in inventory holding costs of each LRU and SRU, the decoupling of the inspection and repair phase of a repair job, and the inefficiency caused by the interruption of repair jobs.

Repair jobs can be released by the inventory control department as soon as repair man capacity becomes available. Repair jobs are released based on the run-out times of the LRUs and availability of the SRUs (i.e. dynamic priority rules). Capacity allocation is easy as the repair capacity is homogenous, however allocating SRUs to repair jobs may be more difficult especially when SRUs are required for the repair of multiple LRUs.

6.4.2.3 Literature review and suggestions for further research

Table 6.3 provides references to literature on decision support models for the decisions in the control concept of repair shop type II, in case it differs from repair shop type I.

Table 6.3 Literature review and suggestions for further research for repair shop type II.

Decision function	Responsible	Suggestions and available decision support	Open research topics
Workload constraint	Inventory control department and repair shop	Trade-off between inventory and repair capacity: Sleptchenko et al. (2003)	Account for priorities on operational level (inspection vs. repair).
LRU & SRU planning	Inventory control department	Trade-off between LRUs and SRUs in MOD-METRIC model: Muckstadt (1973)	Include option that multiple SRUs can fail upon failure of a LRU.
Release function	Inventory control department	Priority rules: Hausman and Scudder (1982)	Priority rules that incorporate distinction between inspection and repair phase.

6.4.3 Repair shop type III

6.4.3.1 Characteristics and design principles

Repair shop type III is characterized by a high capacity complexity and a low material uncertainty. The added value mainly consists of the use of specialized resources in the repair process. The number of repair job interruptions is moderate and mainly due to unavailable (specialized) capacities. For this reason, the WIP in repair shop type III is usually high and a repair man may be working on multiple repair jobs at a time. An example is the mechanics repair shop in one of our case study companies.

The repair capacity is heterogeneous and the number of specialized skills is high. Therefore it is important to effectively utilize the capacity of specialized skills, i.e. repair men with specialized skills should conduct specialized repair steps rather than steps for which only basic skills are required.

Since the material uncertainty is low, like in repair shop type I, the benefits of integrating SRU and LRU inventory control and decoupling the inspection and repair phase are low. Hence the repair shop is responsible for the complete repair job, and it is rather obvious that the repair shop is also responsible for the inventory control of SRUs.

The repair shop has the best knowledge about which repair jobs to prioritize, in order to effectively use the available (specialized) repair capacity. Since the capacity complexity is high, the priority from a repair capacity perspective is very important. Therefore the repair shop is responsible for setting the repair job priorities.

6.4.3.2 Control concept

The inventory control department has the best knowledge about which repair jobs to prioritize from an inventory control perspective. However, the repair shop is responsible for setting the (dynamic) repair job priorities. Therefore we propose that the inventory control department is able to set static repair job priorities. The repair shop is able to dynamically set its own repair job priorities, though respecting these static priorities. The proposed control concept of repair shop type III can be found in Figure 6.8. In this control concept design, the repair shop ‘pulls’ the repair jobs into the repair shop.

We propose to use a control mechanism that decouples the capacity planning of

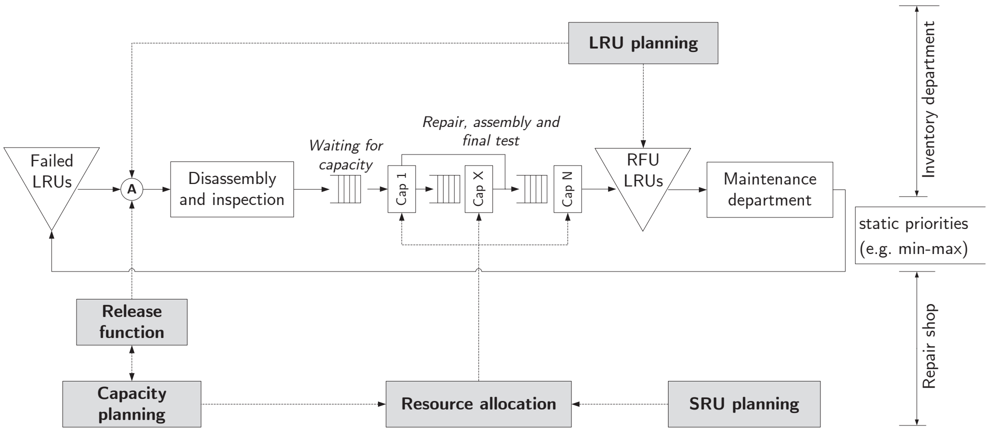

Figure 6.8 *Process and control concept design for repair shop type III*

specialized skills and the LRU planning via these static repair priorities. The inventory control department and the repair shop agree on static priorities, e.g. minimum and maximum stock levels for LRUs, based on which the repair shop has to schedule the repair jobs. With these constraints, the inventory control department is able to determine the number of parts required per LRU. These constraints also ensure that parts are repaired in approximately the right order from an inventory control perspective.

Within the constraints, the repair shop is able to optimize the utilization of the specialists skills, by combining repair steps of different repair jobs. It is important to timely anticipate on shifting bottleneck resources, e.g. by effectively scheduling ahead for the capacity of specialized skills. Proper overtime policies should be applied to make extra use of specialized skills in busy periods and let them idle in quiet periods. For the latter to work out, the static repair priorities should not be too tight.

Repair priorities are in general tight for slow moving LRUs, because the number of parts per LRU is low (typically one). If the repair shop has to repair a relatively high number of slow movers, then an alternative design may be to work with planned lead times instead of static repair priorities (lead time coordination concept). These planned lead times are usually longer than the effective repair time, because of the waiting times for capacities.

In such a concept, the repair shop is responsible for meeting the planned lead times with a high probability. Secondly, the repair shop is responsible to respect a certain repair job release pattern that is agreed upon with the inventory control department, e.g. a minimum amount of repair jobs should be released by the repair shop per period. In this alternative control concept design, the inventory control department should to some extent have the possibility to interrupt the repair shop schedule by expediting repair jobs.

6.4.3.3 Literature review and suggestions for further research

Table 6.4 provides references to literature on decision support models for the decisions in the control concept of repair shop type III, in case it differs from repair shop type I.

Table 6.4 Literature review and suggestions for further research for repair shop type III.

Decision function	Responsible	Suggestions and available decision support	Open research topics
Static repair priorities	Inventory control department and repair shop	Trade-off between extra skills and LRU inventory: Sleptchenko et al. (2017)	Extend to multi-stage setting, relax FCFS-service discipline
		Inventory planning with static priorities: Adan et al. (2009)	Model that integrates repair capacity planning (with a set of specialized skills per repair man) and inventory planning
LRU planning	Inventory control department	Inventory planning with *regular* and *expedite* repair lead time: Arts et al. (2016), Arts (2017)	Extend to multiple capacity types (specializations)

6.4.4 Repair shop type IV

6.4.4.1 Characteristics and design principles

Repair shop type IV is characterized by a high capacity complexity and a high material uncertainty. Usually this type of repair shop repairs a large number of different parts with a large diversity and high value. The added value of the repairs is typically high because expensive materials and specialized skills are required. Utilization rates will be alternating high and low. Examples include the avionics and composites repair

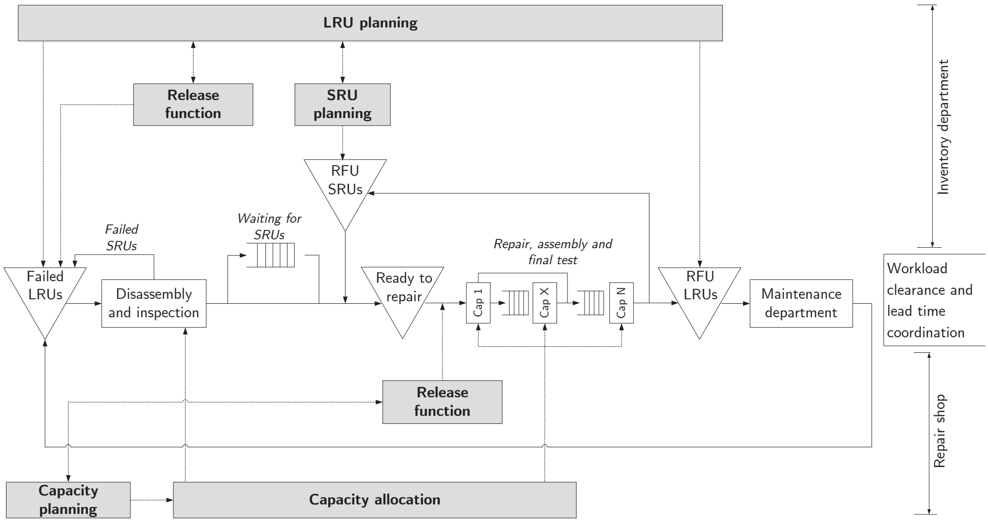

Figure 6.9 *Process and control concept design for repair shop type IV*

shop in one of our case study companies.

The number of repair job interruptions is high and due to unavailable (specialized) repair capacities and/or materials. A high material uncertainty calls for an integrated planning of SRUs and LRUs and the inventory control department to be responsible for repair job release. A high capacity complexity calls for an optimized utilization of the (specialized) repair capacities and hence the repair shop to be responsible for repair job release. The design of the control concept should cope the fact that the repair job release function is a joint responsibility of the inventory control department and the repair shop.

6.4.4.2 Control concept

Like in the design for repair shop type II, we propose to decouple the inspection phase and the repair phase of the repair jobs via a stocking point of parts that are "Ready to repair". This decoupling point enables us to combine the control concept design of repair shop type II and the alternative design (lead time coordination concept) of repair shop type III. Hence the inventory control department and the repair shop need to agree on repair lead times for the repair and test phase of a repair job. The control concept design for repair shop type IV can be found in Figure 6.9.

We propose a workload clearance concept for the inspection of failed parts, that is, the inventory control department and the repair shop agree on: (1) the maximum backlog of inspection work and (2) the maximum amount of work (in hours) related to the inspection of parts, released to the repair shop per period. The inventory control department hence 'pushes' the repair jobs into the repair shop. The workload constraint is only a fraction of the total available repair capacity, as inspection requires only a fraction of the total repair time. Via the workload constraint, the WIP in the repair shop can be kept at an acceptable level.

Furthermore, the inventory control department is able to balance between inspection of parts with a high expected waiting time for SRUs (to decrease SRU waiting time) and parts that should be RFU in a short time period (to decrease LRU waiting time). Like in the design for repair shop type II, the inventory control department is responsible for both SRU and LRU inventory control in order to be able to find the optimal balance between SRU and LRU stock investments.

We propose that the repair shop is able to release ('pull') the repair jobs from the "Ready-to-repair" stock, in order to effectively utilize the available (specialized) repair capacity. Released jobs from the "Ready-to-repair" stock should be completed within the planned lead time with a high probability, and the jobs should be released according to a certain release pattern agreed upon with the inventory control department. We propose to use the lead time coordination concept as discussed in the alternative design of repair shop type III, and hence the inventory control department to some extent has the possibility to expedite repair jobs (that are ready-to-repair).

6.4.4.3 Literature review and suggestions for further research

Table 6.5 provides references to literature on decision support models for the decisions in the control concept of repair shop type IV, in case it differs from other repair shop type III.

6.5 Focus groups

This chapter presents control concepts for repairable parts inventory systems. These control concepts contribute to the following goal: to realize a high availability of high-value capital assets (from a spare parts availability point of view) against acceptable maintenance logistics costs and at an acceptable degree of coordination between

Table 6.5 Literature review and suggestions for further research for repair shop type IV.

Decision function	Responsible	Suggestions and available decision support	Open research topics
Setting lead times & workload clearance levels	Inventory control department and repair shop	Trade-off repair capacity and LRU and SRU stock levels: Zijm and Avşar (2003)	Extend to multiple LRUs, relax FCFS-service discipline

decision makers.

Our hypothesis is that these control concepts serve as a useful reference point for companies to redesign the control concept of existing repairable parts inventory systems. That is, they can be used for a comparison to their current design. This comparison provides insights in how to find an improved balance in the spare parts availability, maintenance logistics costs and degree of coordination.

In this section, we report on focus groups at two companies to verify the hypothesis. This section is organized as follows. In Section 6.5.1, we describe the focus group methodology. In Section 6.5.2 and 6.5.3 we describe the focus groups at KLM E&M and RNLAF in more detail. Section 6.5.4 summarizes the focus group conclusions, and we reflect on the hypothesis.

6.5.1 Methodology

At two companies, KLM E&M and RNLAF, we have organized a separate workshop. Both companies have multiple repair shops and therefore we are able to generate multiple cases per company. Participants of the workshops represent either the inventory control department or the repair shop. The workshops are led by a moderator, who has provided a short introduction to repairable parts inventory systems and the repair shop typology. We have first verified whether all types of repair shops are recognized by the participants. Repair shop types II, III, and IV occur at both case study companies; repair shop type I occurs at KLM E&M but not at RNLAF. The division of the repair shops over the various types can be found in Table 6.6 for KLM E&M and in Table 6.7 for RNLAF. At each company, we have selected one repair shop per repair shop type. Hence, we have studied in depth 4 cases at KLM E&M and 3 cases at RNLAF to test our hypothesis.

Per case, the moderator has outlined the appropriate control concept presented in Section 6.4 while each design choice is discussed in detail. The control concept is compared to the current control concept and next the main differences between the current and the proposed control concept have been identified. Subsequently, advantages and disadvantages of changing from the current to the proposed control concept are discussed. We have stopped the discussion when consensus among the participants was reached. The main findings are summarized by the moderator and the summary is completed based on individual feedback of the participants.

6.5.2 Focus group at KLM Engineering & Maintenance

KLM E&M is an aircraft maintenance organization. KLM E&M carries out maintenance, repairs and modifications on aircraft, engines, and components for their fleet and various other airlines worldwide. KLM E&M has relatively homogenous parts repair streams, because the total volume of parts repair work is relatively high. The parts repair work is distributed among 15 independently operating components repair shops, each responsible for the repair of an individual set of repairable parts.

The participants of the workshop at KLM E&M hold the following positions: one supply chain pool manager, responsible for (LRU) inventory stocking decisions and two managers cell development, together responsible for the coordination of the repair shop decision functions and the design of the repair process in all repair shops within KLM E&M.

One of the participants has presented the current control concept at KLM E&M, which is in fact the same for all repair shops and thus also for all repair shop types. The department responsible for LRU inventory control has made an agreement on a fixed repair lead time for each LRU with each repair shop. Based on the fixed repair lead times, the LRU inventory levels are determined. A target is set on the percentage of repair jobs that has to be completed within the agreed lead time of an LRU. The repair shops are responsible for the inventory control of the SRUs. Hence, the repair shops have full responsibility for achieving the repair lead times.

In Table 6.6 we present the final results of the workshop at KLM E&M. For each repair shop type, we list the number of repair shops classified per type, the repair shop that has served as a case, its characteristics and finally the mapping of the current on the presented control concept.

As can be concluded from Table 6.6, (elements of) the presented control concepts

Table 6.6 Mapping of the control concept in four cases at KLM E&M.

Shop type	Number of shops	Case study	Characteristics	Control concept mapping
I	7	Wheels & tires	High volume repair shop with one-for-one repair and very short lead times.	The current control concept matches with the proposed control concept.
II	4	Toilet repairs	Large number and high diversity of required SRUs with limited availability at suppliers. This leads to many interrupted repair jobs with long and variable lead times.	Integrating the LRU and SRU stocking decision may be useful to improve LRU availability and reduce stock levels. Decoupling inspection and repair is probably not desired for this specific repair shop because it is very inefficient to structurally interrupt the repair jobs in this repair shop.
III	2	Mechanical repairs	Due to long (inefficient) repair set ups, one cannot follow a pure earliest due date rule for the scheduling of the repair jobs, while such a rule would be desirable because of the lead time targets.	Control concept with static repair priorities and repair job priorities set by the repair shop are an alternative to align repair job priority setting and capacity utilization targets.
IV	2	Composites repairs	Because of the high diversity in required SRUs and their limited shelf life, often SRUs are too expensive to stock.	Introduction of a decoupling point between inspection and repair, in combination with dedicated repair job priority setting, will probably result in shorter and more reliable lead times.

provide useful insights for the participants. The main difference in the control concepts is that currently all repair shops are responsible for the inventory control of SRUs, while we propose to integrate the LRU and SRU stocking decisions for repair shop types II and IV. The participants admitted that the current control concepts may result in far from optimal stock levels. The participants have pointed out, however, that the choice for changing the current control concept depends on the expected cost reductions (stock reduction versus increase in coordination).

We have ended the workshop with formulating final conclusions. The participants believe that their current ‘one-size-fits-all’ approach is not optimal, and both the typology and the presented control concepts generate important insights for designing a more differentiated approach.

6.5.3 Focus group at Royal Netherlands Air Force

RNLAF is a high-tech armed forces service that contributes to peace and security on a global basis. For this purpose, it uses and maintains aircraft, helicopters and other weapon systems. At RNLAF, the heterogeneity of parts repair streams is higher than at KLM E&M, because the total volume of parts repair work is small compared to KLM E&M. The total parts repair work is distributed among four independently operating repair shops.

The participants of the workshop at RNLAF hold the following positions: one logistics manager, responsible for the coordination between the inventory control department and the repair shops, the manager of all repair shops, and a planning employee, responsible for the planning of the repair jobs. Within the RNLAF, the inventory control department is responsible for LRU and SRU inventory control and the release of repair jobs to the repair shops. The repair shops are responsible for timely executing the repair jobs. Currently agreements are made on the repair lead time for each individual repair job. This results in a high degree of coordination between the repair shops and the inventory control department.

In Table 6.7, we present the final results of the workshop at RNLAF. For each of the repair shop types II, III, and IV, we again list the number of repair shops per type, the repair shop that served as a case, its characteristics and finally the mapping of the current on our presented control concept. No repair shop was classified as type I.

The main bottleneck in the current control concept is the fact that the repair shops are responsible for meeting targets on the repair lead times, but are not responsible for the availability of the required SRUs. Introduction of a decoupling point between the inspection and repair phase of the repair job is an interesting option, because in that way the repair shops are able to provide the inventory control department with more reliable lead times (after the inspection phase). This however may further increase the already high degree of coordination between the inventory control department and the repair shops. Therefore, the costs reductions and increased coordination should be quantified first before this option can be considered as a serious option.

Together with the participant, we have come to the conclusion that the presented

Table 6.7 Mapping of the control concepts in three cases at RNLAF.

Shop type	**Number of shops**	**Case study**	**Characteristics**	**Control concept mapping**
II	1	Avionic repairs	Many repair job interruptions and low repair lead time reliability due to missing SRUs.	Decoupling inspection and repair might enable the repair shop to provide more reliable completion dates of the repair jobs to the inventory control department.
III	2	Mechanical repairs	Much coordination as reliability of the repair shop depends on SRU availability and repair priorities arranged by the inventory control department.	The high degree of coordination between repair shop and inventory control department can be reduced by assigning the responsibility for SRU inventory control and repair job priorities to the repair shop.
IV	1	Engines and modules repairs	Low lead time reliability and many changes in capacity allocation due to missing SRUs.	Decoupling inspection and repair (lead times) enables the repair shop to create a robust capacity plan and provide more reliable completion dates of the repair jobs.

control concepts give interesting insights for RNLAF to redesign their current control concept, and that a more differentiated approach is required.

6.5.4 Focus group conclusions

In each of the two companies, we see that the same control concept is used for all repair shops. Within the workshops, practitioners agreed that differentiated control concepts are required for the various repair shop types. The typology presented in Section 6.3 can be used to classify the repair shops.

The presented control concepts serve as a reference point. In workshops we have made clear how the companies deviate from the presented control concepts and how they can (re)design their control concepts, in order to decrease the degree of coordination or to decrease the cost level. However, the workshops have also made clear that, besides a qualitative motivation for the design choices, also quantitative support is required for the companies to redesign their control concepts.

6.6 Concluding remarks

In this chapter, we have elaborated on Chapters 4 and 5 and answered the fifth research question:

5. *To what extent do repair shops differ from each other, and what implications do these differences have on the design of control concepts for repairable parts inventory systems?*

From the case studies, we have found that repair shops differ in two dimensions: *capacity complexity* and *material uncertainty*. The value that both dimensions can take on are low and high, and hence we obtain a two-by-two classification and four stereotype repair shops. For each of these repair shop types, designs of the control concepts for repairable parts inventory systems are presented and we have qualitatively motivated our design choices.

For repair shops with a high material uncertainty, we propose to integrate LRU and SRU stocking decisions and to decouple the inspection and repair phase. For repair shops with a low material uncertainty, we propose to decouple the LRU and SRU stocking decisions to enable a simpler interface between the inventory control department and the repair shop (and to align control with responsibilities).

For repair shops with a low capacity complexity, repair priority setting and inventory decisions are integrated. For repair shops with a high capacity complexity, we propose a 'lead time interface' with the option to expedite repair jobs (i.e. a regular and expedited lead time), and repair jobs should be released based upon a cap on aggregate flows per resource.

The use of these control concepts as a reference point for designing control concepts of existing repairable parts inventory systems, is verified in focus groups at two companies. Although these designs are verified in case studies, we can not substantiate the claim that these designs are 'optimal', nor that they are always applicable. Quantifying the impact of the design choices is necessary to further improve and validate the designs. One of our design choices is to decouple the inspection and repair phase for repair shops with a high material uncertainty. The next chapter addresses a quantitative analysis on the impact of decoupling the inspection and repair phase on the total system costs, for repairable parts inventory systems with a repair shop of type II (i.e. with a low capacity complexity and a high material uncertainty).

Chapter 7

Capacity assignment under high material uncertainty

"Elk nadeel heeft zijn voordeel."

Johan Cruijff

7.1 Introduction

Repair shops may have different levels of capacity complexity and material uncertainty. A high material uncertainty implies that the material requirements to repair a failed part are not known beforehand. The inspection phase of a repair job serves as the repair job preparation phase, that is, only after the inspection phase it is known which materials and capacities are required to complete the repair job. In this chapter we restrict ourselves to a repair shop of type II (see Chapter 6), i.e. a repair shop with a low capacity complexity and a high material uncertainty. We assume that interruptions between inspection and actual repair are only caused by unavailability of the required materials (SRUs).

In repairable parts inventory systems, decisions to set the SRU and LRU base stock levels, repair capacities, repair shop scheduling rules and the rules for allocation of SRUs to repair jobs clearly interrelate. For repairable parts inventory systems with high material uncertainty, the decision to set the SRU base stock levels seems especially relevant. We therefore briefly review literature that considers SRU stocking decisions, for a more comprehensive review we refer to Chapter 4.

Main contributions to an integrated approach to set SRU and LRU stock levels are Muckstadt (1973), Sherbrooke (1986) and Rustenburg et al. (2003). The analysis in these contributions requires the assumptions of i.i.d. repair lead times and the fact that at most one SRU fails upon failure of an LRU. Sleptchenko et al. (2002) have showed that if realizing i.i.d. repair lead times is difficult (e.g. high repair shop utilization), then this leads to significant underestimation of stock levels and service performance. Zijm and Avşar (2003) present exact and approximate models to determine optimal base stock levels for a single LRU, two-indenture setting with finite repair capacity. Sleptchenko et al. (2003) present a model that maximizes asset availability, under a budget constraint for SRUs, LRUs and finite repair capacity. Again the assumption of at most one SRU failure upon failure of an LRU is needed to facilitate the analysis. Van Jaarsveld et al. (2015) present a method to minimize the SRU stock investments, while ensuring a minimum probability that all SRUs are available for a random repair job within a certain (LRU specific) time-window. They state and model that LRUs in general consist of multiple SRUs, any of which may be required for the repair of an LRU, and compare their setting to assemble-to-order systems, in which *components* should be interpreted as SRUs and *products* as LRU repair jobs.

Table 7.1 summarizes the model characteristics of the cited references, including the model we present in this chapter. We introduce two new aspects that are currently not covered in the literature. First, inspection and repair both take a significant amount of time and they are executed by the same resources. Second, repair times depend on the time that elapses between inspection and the actual repair of a part. Repair times are higher for interrupted repair jobs, because the inspected part needs to be put aside and the repair man has to go into the problem again, i.e. a part of the inspection has to be redone when the actual repair starts. We take into account that multiple SRUs can be required for the repair of a part, and the set of SRUs that is actually required for repair jobs of the same LRU can vary from one repair job to another.

These new aspects raise the question whether priority should be given to inspection of failed parts (to reveal which SRUs are required and speed up sourcing of unavailable SRUs) or to completion of the repair of parts for which all required SRUs are available

Table 7.1 Model characteristics of cited references and this chapter

References	Modeling capacity	Modeling materials	Method	Limitations
(MOD)METRIC-type models: a.o. Sherbrooke (1968), Muckstadt (1973), Rustenburg et al. (2003)	Not applicable	- Multiple indenture levels - At most one SRU fails on failure of an LRU - Base stock policy parameters of SRUs and LRUs are decision variable	Approximate evaluation	- Repair lead times are i.i.d. variables - At most one SRU required per repair job
Van Jaarsveld et al. (2015)	Not applicable	- Two-indenture levels - Multiple SRU failures upon failure of an LRU - (s,S)-policies for SRUs	Approximate	Inventory policy parameters of LRUs are not included as decision variable
Zijm and Avşar (2003)	Both M/M/1 as well as product form queuing networks with exponential service times	- Two-indenture levels - One SRU fails upon failure of the LRU - Base stock policy parameters of SRUs and the LRU are decision variable	Exact and approximate	- Restricted to one LRU - Inspection as well as repair and assembly of parts of each indenture level require separate resources - Exponential service times are in general not realistic
Sleptchenko et al. (2003)	- M/M/k-queue - M/G/c-queue (but requires assumption of at most 1 LRU)	- Multiple indenture levels - At most one SRU required per repair job - Base stock levels of SRUs and LRU are decision variable	Approximate	- Inspection takes no time and is performed upon failure (i.e. preemption is assumed in order to immediately execute inpections) - Preemption does not affect service times - Exponential service times are in general not realistic - Extension to a setting with many LRUs is limited due to high computation times.
This chapter	- M/G/1-queue - Explicit modeling of inspection time; inspection and repair are executed by the same resource - Repair times depend on the time that elapses between inspection and repair	- Two-indenture levels - Multiple SRU failures possible upon failure of an LRU - Base stock policy parameters of SRUs and LRU are decision variable	Simulation	- Predetermined (possibly non-optimal) SRU base stock levels - Single location model

(to increase the availability of LRUs), and this is the final research question of this dissertation:

6. *What is the optimal capacity priority policy in repairable parts inventory systems with high material uncertainty?*

We model capacity as a single server and we compare (the total costs for) policies that, based on the repair workload in the repair shop, give priority to either inspection or repair of parts. The total costs include inventory holding costs and backordering costs, which clearly depend on the base stock levels of the SRUs and LRUs. One would ideally compare the capacity priority policies under the optimal SRU and LRU base stock levels. However, these optimal base stock levels also depend on the capacity priority policy itself. To the best of our knowledge, models to analytically solve this problem are not known in the literature. Therefore, we take multiple methods to set the SRU base stock levels as given. Based on simulation, the LRU base stock levels are optimized per capacity priority policy. This leads to a fair comparison of capacity priority policies.

This chapter is organized as follows. In section 7.2 we formulate our model. Section 7.3 decribes the options we used to set the SRU base stock levels and section 7.4 discusses the capacity priority policies that we evaluate. In section 7.5 we describe how we set up the simulation experiment and in section 7.6 we present a numerical study. Section 7.7 concludes this chapter and lists some suggestions for further research.

7.2 Model description

We consider a single-location, multi-item, two-indenture inventory model for repairable LRUs and a repair shop responsible for the repair of a fixed and finite set of repairable LRUs, denoted by I. The LRUs are numbered as 1,...,$|I|$. LRUs are used during first level maintenance of an asset, i.e. failed LRUs are replaced by ready-for-use (RFU) LRUs supplied from stock. The inventory of all LRUs is controlled by a base stock policy and the vector of base stock levels is given by $\mathbf{S} = (S_1, ..., S_{|I|})$. The inventory holding costs are denoted by h_i for LRU i. Demands for LRU i occur according to a Poisson process with constant rate λ_i (≥ 0) and the demands are served FCFS. Unfulfilled demand for LRUs is backordered, and a backorder cost b (independent of i) is incurred per part per time unit.

We assume that replaced LRUs actually *failed*, i.e. we assume that there are no 'no fault founds'. The failed LRUs are sent to the stock point "Failed LRUs" and the return time is assumed to be equal to 0. Failed parts are inspected FCFS by the repair shop, and we assume that the repair of a part is always successful (i.e. no condemnation). The inspection phase of a repair job cannot be interrupted (non-preemption) and the average inspection time equals T_i^{Insp} time units for LRU i.

Each repair job may require a different set of SRUs. After the inspection phase, the set of required SRUs is known and demanded. Let $J = \{1, ..., |J|\}$ be the non-empty set of all SRUs. The required quantity of SRU j for the repair of LRU i is equal to 1 with probability $\alpha_{i,j}$ $(0 \leq \alpha_{i,j} \leq 1)$ and equal to 0 with probability $(1 - \alpha_{i,j})$. Consequently, the aggregate demand rate $(\hat{\lambda}_j)$ for SRU j equals:

$$\hat{\lambda}_j = \sum_{i \in I} \lambda_i \cdot \alpha_{i,j}, \quad \forall j \in J.$$

The inventory of all SRUs is controlled by a base stock policy and the vector of base stock levels is given by $\hat{\mathbf{S}} = (\hat{S}_1, ..., \hat{S}_{|J|})$. The inventory holding costs are denoted by $\hat{h}_j$ for SRU j. We assume that all SRUs are consumable[1] and a replenishment order for SRU j is placed directly when a new part of SRU j is demanded (directly after the inspection of a repair job). SRU j has a deterministic replenishment lead time $\hat{t}_j > 0$.

Demands for SRUs are served FCFS and on-hand stock of SRUs, if available, is assigned to the repair job that requires the SRUs. We distinguish assigned and unassigned on-hand SRU stock. Assigned on-hand SRU stock is stock that is already assigned to repair jobs, but are waiting for other SRUs or waiting for repair. Unassigned on-hand stock of SRU j is stock that is available to be assigned to new repair jobs requiring SRU j.

Let $M \subseteq J$ be the set of required SRUs for a random repair job, let $M_a \subseteq M$ be the set of SRUs for which still unassigned on-hand stock is available and let $M_u \subseteq M$ be the set of SRUs for which no unassigned on-hand stock is available. After the inspection phase, for each SRU $j \in M_a$ the assigned on-hand stock is increased with

[1] Note that the assumption of consumable SRUs can also be relaxed to the case where SRUs are repaired by external repair shops with which fixed lead time agreements are made. This assumption excludes repairable SRUs that are repaired by the same repair shop responsible for the repair of the LRU for which the SRU is required. In such cases the capacity priority policy interacts with the replenishment lead times of SRUs, which complicates the evaluation of the policies. Investigation of this specific case is beyond the scope of our study and is a suggestion for further research.

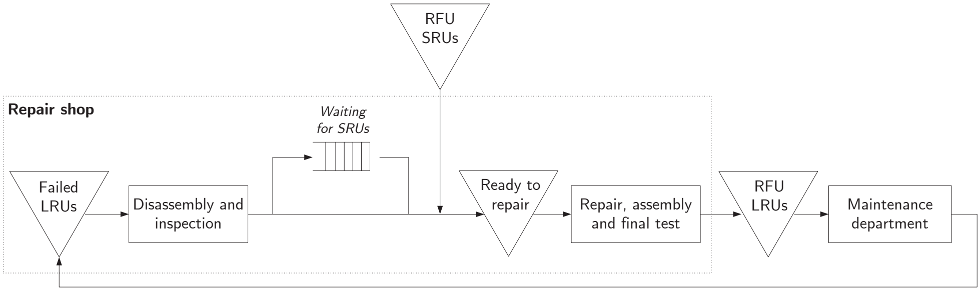

Figure 7.1 *Overview of the general repair process*

one unit and the unassigned on-hand stock is decreased with one unit, but at the same time a replenishment order for one unit is placed. For each SRU $j \in M_u$, the number of backorders is increased with one unit (again, while placing a replenishment order). Once a replenishment order for SRU j arrives, either the unassigned stock of SRU j is increased or the number of backorders of SRU j is decreased, and in the latter case also the assigned on-hand stock of SRU j is increased.

Repair jobs for which all required SRUs are available, are sent to the stock point "Ready to repair". All other repair jobs are queued "waiting for SRUs" until the last required SRU is available from stock. The required SRUs are taken from the on-hand stock when the repair phase of the repair job starts.

The repair phase of a repair job cannot be interrupted (non-preemption) and takes on average T_i^{Rep} time units for LRU i. If the repair phase does not start within T^{Max} (≥ 0) time units after completion of the inspection phase, an additional time $\delta_i \cdot T_i^{insp}$ is required before the actual repair can take place, with $0 \leq \delta_i \leq 1$. This additional time is referred to as 'inefficiency' (a part of the inspection phase has to be redone), and δ_i is referred to as the 'inefficiency factor' for LRU i. After the repair phase, the part is "ready-for-use" again and sent to the ready-for-use LRU stock ("RFU LRUs"). For an overview of the entire process see Figure 7.1.

We assume that the repair shop consists of a single server (repair man) that works with speed 1. The workload, i.e. the sum of inspection and repair time, of an efficiently executed repair job for LRU i has a general distribution with average μ_i and standard deviation σ_i time units. The ratio between inspection and repair is constant for each (efficiently executed) repair job: a fixed percentage of the workload of each (simulated) repair job is required for inspection, denoted by p_i for LRU i. Hence, T_i^{Insp} equals $p_i\mu_i$ and T_i^{Rep} equals $(1 - p_i)\mu_i$.

In case not idle, the server works either on inspection or repair of a single repair job. The selected capacity priority policy, denoted by π, determines when priority is given to either inspection or repair of parts. The server utilization ρ depends on the SRU base stock levels and the capacity priority policy π that is applied. In order to make sure that the system is stable for all situations, we assume that the actual repair time is based on the situation in which all repair jobs are executed 'inefficiently'. The following condition is sufficient to ensure our system is stable:

$$\rho = \frac{\sum_{i \in I} \lambda_i \cdot \left((1+\delta_i) T_i^{Insp} + T_i^{Rep} \right)}{\sum_{i \in I} \lambda_i} < 1$$

We assume that $\pi \in \Pi$, where Π is a class of capacity priority policies that dynamically assign capacity to either inspection or repair of parts, based on information in the repair shop. These capacity priority policies do not take into account information on the ready-for-use LRU stock, and hence the repair pipeline process of LRUs is independent of the base stock levels of the LRUs. This property is also referred to as: "the repair pipeline process is translation invariant" (see e.g. Kohler-Gudum and de Kok (2002)). With this property we are able to determine the optimal LRU base stock levels $\mathbf{S}^*(\hat{\mathbf{S}}, \pi)$, after we obtain the simulation results for a given vector $\hat{\mathbf{S}}$ and capacity priority policy π.

Let $A_i(\hat{\mathbf{S}}, \pi, q)$ be the (steady state) probability that the number of parts of LRU i in the repair shop equals q, $q \in \mathbb{N}_0$, given a vector of SRU base stock levels $\hat{\mathbf{S}}$ and capacity priority policy π. Further, let $B_j(\hat{\mathbf{S}}, \pi, k)$ be the (steady state) probability that the (assigned plus unassigned) on-hand stock of SRU j equals k, $k \in \mathbb{N}_0$. The steady state probabilities of $A_i(\hat{\mathbf{S}}, \pi, q)$ and $B_j(\hat{\mathbf{S}}, \pi, k)$ will be determined with discrete event simulation, and both $\hat{\mathbf{S}}$ and π are input variables for the simulation.

The average total system costs, given a vector of SRU base stock levels $\hat{\mathbf{S}}$, capacity priority policy π and the vector of base stock levels for the LRUs ($\mathbf{S}$), are defined as:

$$C(\hat{\mathbf{S}}, \pi, \mathbf{S}) = \sum_{j \in J} \hat{h}_j \cdot \mathbb{E}[B_j(\hat{\mathbf{S}}, \pi, k)] + \sum_{i \in I} h_i \cdot S_i + \sum_{i \in I} \sum_{q=S_i+1}^{\infty} A_i(\hat{\mathbf{S}}, \pi, q) \cdot (q - S_i) \cdot b \quad (7.1)$$

The inventory holding costs for the SRUs (first part of the cost function) follow directly from the simulation results of the on-hand stock distribution $B_j(\hat{\mathbf{S}}, \pi, k)$. With the repair pipeline distribution $A_i(\hat{\mathbf{S}}, \pi, q)$, we are able to determine the optimal LRU spares vector $\mathbf{S}^*$. The total LRU costs are separable for each LRU i and are equal to:

$$C_i(\hat{\mathbf{S}}, \pi, S_i) = h_i \cdot S_i + \sum_{q=S_i+1}^{\infty} A_i(\hat{\mathbf{S}}, \pi, q) \cdot (q - S_i) \cdot b$$

Its first order difference function is:

$$\begin{aligned}\Delta_i C_i(\hat{\mathbf{S}}, \pi, S_i) &= C_i(\hat{\mathbf{S}}, S_i + 1) - C_i(\hat{\mathbf{S}}, \pi, S_i) \\ &= h_i - b \cdot \sum_{q=S_i+1}^{\infty} A_i(\hat{\mathbf{S}}, \pi, q) \\ &= h_i - b + b \sum_{q=0}^{S_i} A_i(\hat{\mathbf{S}}, \pi, q)\end{aligned}$$

It can be easily seen that $C_i(\hat{\mathbf{S}}, \pi, S_i)$ is decreasing if $\sum_{q=0}^{S_i} A_i(\hat{\mathbf{S}}, \pi, q) \leq \frac{b - h_i}{b}$ and increasing if $\sum_{q=0}^{S_i} A_i(\hat{\mathbf{S}}, \pi, q) \geq \frac{b - h_i}{b}$. Further, $C_i(\hat{\mathbf{S}}, \pi, S_i)$ is convex because:

$$\Delta_i^2 C_i(\hat{\mathbf{S}}, \pi, S_i) = b \cdot A_i(\hat{\mathbf{S}}, \pi, S_i + 1) \geq 0, \quad \forall S_i \in \mathbb{N}_0$$

Hence, an optimal base stock level of LRU i is the smallest value S_i^* that satisfies:

$$\sum_{q=0}^{S_i^*} A_i(\hat{\mathbf{S}}, \pi, q) \geq \frac{b - h_i}{b}$$

With the optimal base stock levels of the LRUs, we are able to determine the total system costs $C_i(\hat{\mathbf{S}}, \pi, S_i^*)$. As mentioned earlier, the vector of SRU base stock levels $\hat{\mathbf{S}}$ and the capacity priority policy π are input variables for the simulation. How they are chosen is explained in sections 7.3 and 7.4, respectively.

7.3 Base stock levels of SRUs

The SRU base stock levels are input variables for the simulation and hence need to be determined upfront. To do so, we need to make an assumption regarding the demand process of the SRUs. While LRU demands occur according to a Poisson process, SRUs clearly do not. This is due to the fact that capacity (inspection) is required to determine the demands for SRUs, and hence demands for SRUs occur a certain (non-deterministic) time after the LRU demand occurred. In the remainder of this section, in which we describe three options to set the SRU base stock levels, we nevertheless assume that the SRU demands do occur according to a Poisson process. The first two options are rather simple, but we would like to mention that the derivation and description of the third option is long.

Option 1: Set all base stock levels of SRUs to 0

In this option, the base stock levels of all SRUs are equal to 0, i.e.

$$\hat{S}_j = 0, \quad \forall j \in J.$$

Option 2: All SRUs have an identical fill rate target γ

In this option, the base stock level $\hat{S}_j$ of SRU j is the smallest value for which the item fill rate of SRU j, denoted by $\beta_j(\hat{S}_j)$, is larger than or equal to γ, i.e.

$$\beta_j(\hat{S}_j) = \sum_{x=0}^{\hat{S}_j - 1} P(X_j = x) \geq \gamma, \quad \forall j \in J,$$

where X_j represents the replenishment pipeline distribution of SRU j, which is assumed to be Poisson distributed with mean $\hat{\lambda}_j \hat{t}_j$.

Option 3: Maximize SRU job completeness, given a constraint on the holding costs for unassigned on-hand SRU stock

In this option, our objective is to maximize the probability that all SRUs for a random repair job are available, given a target value of the holding costs for unassigned on-hand SRU stock C^{obj}. Note that we focus on unassigned on-hand SRU stock only, despite the fact that the cost function in section 7.2 includes both assigned and unassigned on-hand SRU stock. We need to make this assumption because in general (most) inventory models, also the models used in this section, assume that demand is served directly from stock, if available, and hence focus on unassigned stock only.

The probability that all SRUs for a random repair job are available immediately after inspection will be referred to as 'SRU job completeness'. The SRU job completeness for LRU i, given a vector of SRU base stock levels $\hat{\mathbf{S}}$, is denoted by $\beta_i(\hat{\mathbf{S}})$. Let us now define the aggregate SRU job completeness as:

$$\beta(\hat{\mathbf{S}}) = \sum_{i \in I} \frac{\lambda_i}{\lambda} \beta_i(\hat{\mathbf{S}})$$

where $\lambda = \sum_{i \in I} \lambda_i$. The SRU job completeness for a random repair job of LRU i is 1 in case each individual SRU j, $j \in J$, is either (a) not required or (b) required and directly available from stock. In all other cases, the SRU job completeness is 0. Note that the SRU job completeness measure is equivalent to the immediate order fill rate for assemble-to-order systems. For such systems, it is well-known that due to correlation of demand across the components, the net inventories of the components are correlated random variables.

This applies to our problem in the following way. The demand processes of SRUs, that might simultaneously be required for the same repair job, are correlated via the demand process of the parent LRU that requires these SRUs for repair. Hence, the steady-state distributions of the on-hand stock of SRUs are correlated random variables. When demand processes of SRUs are strongly correlated, then one might want to integrally determine the base stock levels of these SRUs.

Recall that $\alpha_{i,j}$ is the probability that SRU j is required for the repair of LRU i and $\beta_j(\hat{S}_j)$ is the fill rate of SRU j. The probability that SRU j, $j \in J$, is not required for a random repair job of LRU i, is equal to $(1 - \alpha_{i,j})$. The probability that SRU j, $j \in J$, is required for a random repair job of LRU i and available from stock, is equal to $\alpha_{i,j}\beta_j(\hat{S}_j)$. For low values of $\alpha_{i,j}$ and a large number of SRUs that might be required for the repair of an LRU, the correlation of the steady-state distributions of the on-hand stock of SRUs will be relatively low. Ignoring the correlations will not strongly influence the choice of the SRU base stock levels and we therefore assume that the steady-state distributions of the on-hand SRU stocks are uncorrelated. We can now write $\beta_i(\hat{\mathbf{S}})$ as:

$$\beta_i(\hat{\mathbf{S}}) = \prod_{j \in J} \Big((1 - \alpha_{i,j}) + \alpha_{i,j}\beta_j(\hat{S}_j) \Big)$$

We can rewrite the SRU job completeness $\beta(\hat{\mathbf{S}})$ as:

$$\beta(\hat{\mathbf{S}}) = \sum_{i \in I} \frac{\lambda_i}{\lambda} \prod_{j \in J} (1 - \alpha_{i,j}(1 - \beta_j(\hat{S}_j)))$$

$$= \sum_{i \in I} \frac{\lambda_i}{\lambda} \left(1 - \sum_{j \in J} \alpha_{i,j}(1 - \beta(\hat{S}_j)) \right.$$

$$\left. + \sum_{j \in J} \sum_{k \in J, k \neq j} \alpha_{i,j} \alpha_{i,k} (1 - \beta_j(\hat{S}_j))(1 - \beta_k(\hat{S}_k)) - ... \right)$$

$$\approx \sum_{i \in I} \frac{\lambda_i}{\lambda} \left(1 - \sum_{j \in J} \alpha_{i,j}(1 - \beta_j(\hat{S}_j)) \right)$$

$$= 1 - \frac{1}{\lambda} \sum_{j \in J} \sum_{i \in I} \lambda_i \alpha_{i,j} (1 - \beta_j(\hat{S}_j))$$

$$= 1 - \sum_{j \in J} \frac{\hat{\lambda}_j}{\lambda} (1 - \beta_j(\hat{S}_j))$$

$$= 1 - \frac{1}{\lambda} \sum_{j \in J} \varphi_j(\hat{S}_j)$$

where $\varphi_j(\hat{S}_j) = \hat{\lambda}_j(1 - \beta_j(\hat{S}_j))$ and where $\sum_{j \in J} \frac{\hat{\lambda}_j}{\lambda}(1 - \beta_j(\hat{S}_j))$ can be regarded as the 'weighted non-fill rate'. We would like to emphasize that the approximation used only holds for small values of $\alpha_{i,j}(1 - \beta_j(\hat{S}_j))$, $\forall i \in I$. That is, for each combination of LRU $i \in I$ and SRU $j \in J$, either the fill rate $\beta_j(\hat{S}_j)$ is high, and/or the probability that SRU j is required for the repair of LRU i is small.

Let $C(\hat{\mathbf{S}})$ denote the total SRU holding costs (calculated over the unassigned on-hand

SRU stock) and let $C_j(\hat{S}_j)$ denote the holding costs for SRU j, $\forall j \in J$, with:

$$C_j(\hat{S}_j) = \hat{h}_j \sum_{x=0}^{\hat{S}_j - 1} (\hat{S}_j - x) P(X_j = x)$$

$$C(\hat{\mathbf{S}}) = \sum_{j \in J} C_j(\hat{S}_j)$$

We can now write our objective as the following problem (P):

$$\begin{array}{lll} (P) & \max & \beta(\hat{\mathbf{S}}) \\ & \text{subject to} & C(\hat{\mathbf{S}}) \leq C^{obj}, \\ & & \hat{\mathbf{S}} \in \mathbf{\Phi}. \end{array}$$

Solving problem (P) is closely related to solving the following multi-objective programming problem (Q):

$$\begin{array}{lll} (Q) & \min & \sum_{j \in J} C_j(\hat{S}_j) \\ & \min & \sum_{j \in J} \varphi_j(\hat{S}_j) \\ & \text{subject to} & \hat{S}_j \in \Phi_j \quad \forall j \in J. \end{array}$$

where $\Phi_j = \mathbb{N}_0$ represents the solution space for $\hat{S}_j$, $\forall j \in J$. For problem (Q), we can derive so-called *efficient solutions*. In short, a solution is *efficient* if there exists no other solution for which the SRU job completeness is higher with the same SRU holding costs, or at least the same SRU job completeness can be attained with lower SRU holding costs. The set of efficient solutions form an *efficient frontier* for the total SRU holding costs and the SRU job completeness (or better: the 'non-fill rates' of SRUs). We can select an appropriate solution for problem (P) from this efficient frontier.

In problem (Q), the objective functions are separable for each SRU j and therefore problem (Q) as a whole is separable. A set of efficient solutions can be generated by a greedy algorithm if we can show that $\varphi_j(\hat{S}_j)$ is decreasing and convex and $C_j(\hat{S}_j)$ is increasing and convex. First, we will show that $\varphi_j(\hat{S}_j)$ is decreasing on $\mathbb{N}_0$:

$$\Delta\varphi_j(\hat{S}_j) = \varphi_j(\hat{S}_j + 1) - \varphi_j(\hat{S}_j) = -\hat{\lambda}_j P(X_j = \hat{S}_j) \leq 0$$

Furthermore,

$$\Delta^2\varphi_j(\hat{S}_j) = -\hat{\lambda}_j\left(P(X_j = \hat{S}_j + 1) - P(X_j = \hat{S}_j)\right)$$

$$= -\hat{\lambda}_j\left(\frac{\hat{\lambda}_j\hat{t}_j}{\hat{S}_j + 1} - 1\right)P(X_j = \hat{S}_j)$$

As $\hat{\lambda}_j \geq 0$ and $P(X_j = \hat{S}_j) \geq 0$, it directly follows that $\Delta_j^2\varphi(\hat{S}_j) \geq 0$ if and only if $\hat{S}_j \geq \hat{\lambda}_j\hat{t}_j - 1$. This means that $\varphi_j(\hat{S}_j)$ is convex on $\Phi'_j = \{\hat{S}_j \in \mathbb{N}_{\nvdash} : \hat{S}_j \geq \hat{\lambda}_j\hat{t}_j - 1\}$, $\forall j \in J$.

Now, we will show that $C_j(\hat{S}_j)$ is decreasing and convex. We first show that $C_j(\hat{S}_j)$ is increasing on $\mathbb{N}_0$:

$$\Delta C_j(\hat{S}_j) = C_j(\hat{S}_j + 1) - C_j(\hat{S}_j)$$
$$= \hat{h}_j\sum_{x=0}^{\hat{S}_j} P(X_j = x) \geq 0$$

Furthermore, we can show that $C_j(\hat{S}_j)$ is convex on $\mathbb{N}_0$:

$$\Delta^2 C_j(\hat{S}_j) = \hat{h}_j P(X_j = \hat{S}_j + 1) \geq 0$$

Hence, $\forall j \in J$, $C_j(\hat{S}_j)$ is increasing and convex on $\mathbb{N}_0$.

We now exclude solutions with $\hat{S}_j < \max\{\lceil\hat{\lambda}_j\hat{t}_j - 1\rceil, 0\}$ from our solution space and discuss the excluded part. The amount $\hat{\lambda}_j\hat{t}_j$ represents the average number of parts in the replenishment pipeline of SRU j. If the average replenishment pipeline stock is smaller than 1 for SRU j, then $\Phi_j = \Phi'_j$ and the solution space remains the same. If the replenishment pipeline stock is larger than 1 (which is in general not the case for slow moving parts), then we exclude solutions with $\hat{S}_j < \lceil\hat{\lambda}_j\hat{t}_j - 1\rceil$. In problem instances with a high target for the SRU job completeness (i.e. a low value of $\sum_{j\in J}\varphi_j(\hat{S}_j)$), these solutions are generally not relevant. Besides this, the average number of parts on stock is low for all values of $\hat{S}_j < \lceil\hat{\lambda}_j\hat{t}_j - 1\rceil$. Consequently, $C_j(\hat{S}_j)$ is also low in these cases and hence we would expect that, in an optimal solution, often the base stock levels of each SRU j $\hat{S}_j$ will be set larger than or equal to $\lceil\hat{\lambda}_j\hat{t}_j - 1\rceil$. For the latter reason, we would expect that reducing the solution space does not significantly influence the optimal solution even if the target for the SRU job completeness is low.

With the adjusted solution space, problem Q now changes to problem Q':

$$
\begin{aligned}
(Q') \qquad & \min && \sum_{j \in J} C_j(\hat{S}_j) \\
& \min && \sum_{j \in J} \varphi_j(\hat{S}_j) \\
& \text{subject to } && \hat{S}_j \in \Phi'_j \quad \forall j \in J.
\end{aligned}
$$

For problem Q', we are now able to construct a set of efficient solutions by a greedy algorithm (see e.g. Fox (1966)), because $C_j(\hat{S}_j)$ in increasing and convex on Φ' and $\varphi_j(\hat{S}_j)$ is decreasing and convex on Φ'. A first efficient solution can be computed by setting $\hat{S}_j = \lceil \hat{\lambda}_j \hat{t}_j - 1 \rceil$, $\forall j \in J$. This solution is efficient (as long as $C(\hat{\mathbf{S}}) \leq C^{\text{obj}}$), because it has the lowest possible total SRU holding costs. Next, we can compute a 'greedy factor', representing the decrease in $\varphi_j(\hat{S}_j)$, equal to $-\Delta\varphi_j(\hat{S}_j)$, relative to the increase in $C_j(\hat{S}_j)$, which is equal to $\Delta C_j(\hat{S}_j)$. The 'greedy factor' τ_j for each SRU $j \in J$ is equal to:

$$
\tau_j = \frac{\hat{\lambda}_j P(X_j = \hat{S}_j)}{\hat{h}_j \sum_{x=0}^{\hat{S}_j} P(X_j = x)}, \quad \hat{S}_j \geq \hat{\lambda}_j \hat{t}_j - 1, \quad j \in J.
$$

The base stock level of the SRU $j^* \in J$ with highest value of τ_j is increased by one unit and a new 'greedy factor' τ_{j^*} is determined based on the new base stock level $\hat{S}_{j^*}$. This procedure is repeated until the total SRU holding costs exceeds its target level C^{obj}. We select the first solution $\hat{\mathbf{S}}$ for which the SRU holding costs exceed C^{obj}. We would like to emphasize that C^{obj} is a target value and should not be interpreted as a constraint, and for this reason we need not necessarily take the second to last solution.

7.4 Capacity priority policies

The second input for our simulation model is the choice of the capacity priority policy. Note that, because of our model assumptions, we are restricted to policies that do not take into account information on the actual ready-for-use LRU stock. In this section we motivate which policies we compare and we explain how each policy gives priority to either inspection or repair of parts.

Giving priority to inspection of failed parts facilitates earlier sourcing of the required SRUs and it potentially reduces the probability that the server becomes idle while there are still parts in the repair shop (but waiting for SRUs), which could reduce the actual server utilization in following periods. On the other hand, it delays the actual repair and hence the availability of ready-for-use LRUs. In addition, separating inspection and actual repair may increase the actual repair time and is hence inefficient from a server utilization point of view. Giving priority to the repair of parts has the opposite effect.

The effects seem to depend on the state of the system. Let us therefore first discuss whether and when inspection or repair should have priority. It is easy to see that situations in which the number of parts in the repair shop is high (either for an individual LRU or for all LRUs), coincide with situations in which the number of LRU backorders is high and/or the ready-for-use LRU stock is low. This is due to the fact the parts circulate in a closed-loop. Hence, we would ideally allocate capacity to repair when the number of parts in repair (or better: the corresponding workload) is high, in order to minimize the (probability of) backorders and the corresponding costs. When the number of parts in repair is low, the probability of backorders will also be low and hence in this situation, giving priority to inspection could be beneficial.

For this reason, we will compare (repair) workload-based capacity priority policies. The repair job workload, i.e. the time required to complete the repair phase of a repair job, is (assumed to be) known after the inspection phase. Let L be the repair workload of all repair jobs that are either in repair or "ready-to-repair", and let $\breve{L}$ (≥ 0) be the workload threshold, based on which will be determined whether priority should be given to inspection or to repair. Note that if either the queue of jobs waiting for inspection or repair is empty, then always priority is given to the first job in the non-empty queue. Otherwise, if $L < \breve{L}$, then priority is given to inspection and if $L \geq \breve{L}$, then priority is given to repair. By varying the values of $\breve{L}$, we can construct a series of capacity priority policies.

Note that we have two extreme policies. In case $\breve{L} = 0$, then priority is always given to repair (provided that "ready-to-repair" parts are available). This policy we will refer to as 'repair priority policy' in the remainder of this chapter. In case $\breve{L} = \infty$, then priority is always given to inspection (provided that parts that need inspection are available). This policy we will refer to as 'inspection priority policy' in the remainder of this chapter. Any hybrid policy, i.e. in which $0 < \breve{L} < \infty$, will be referred to as 'hybrid priority policy'.

To illustrate our capacity priority policies, an example of the Gantt-charts is visualized in Figure 7.2. At the start, three jobs (1,2 and 3) are waiting for inspection. Over time, new jobs (4, 5 and 6) enter the repair shop. In the inspection priority policy, first the inspection of jobs 1, 2 and 3 is executed. Job 3 requires a SRU that is not available from stock and that is ordered at an external supplier. The repair phase of job 3 can only start when the required SRUs are available. Once there are no jobs waiting for inspection left, the repair phase of job 1 starts. The additional time required for repair (i.e. the 'inefficiency') is marked in black. After completion of the repair phase of job 1, it is checked whether there are parts waiting for inspection. As this is not the case in our example, the repair phase of job 2 starts.

In the hybrid priority policy, we see that first jobs 1 and 2 are inspected. The repair workload of these jobs can be found between brackets (0.8 and 1.3 respectively). The repair workload of these jobs together (2.1) exceeds the threshold of 2.0 and hence the repair of job 1 starts (in contrast to what happens under the inspection priority policy). In the repair priority policy, the repair phase (in principal) directly follows the inspection phase. In this example, one of the disadvantages of this policy is illustrated: the repair man is idle after completing the repair phase of job 4 while this is not the case for the inspection priority policy (because the required SRUs were known, ordered and received earlier).

7.5 Simulation

Given a predetermined vector of SRU base stock levels $\hat{\mathbf{S}}$ and a choice of the capacity priority policy π, we are able to simulate the system (performance). The outcome of the simulation are the variables $A_i(\hat{\mathbf{S}}, \pi, q)$ and $B_j(\hat{\mathbf{S}}, \pi, k)$, which are used to determine the total system costs. In section 7.5.1 we describe how we set up the simulation experiment and how we ensure that the simulated values are statistically significant. Section 7.5.2 decribes the simulation state space, (the update of) the related variables and we explain how the capacity priority policies are applied during the simulation.

7.5.1 Set up of simulation experiments

The performance of the system is evaluated via discrete event simulation. We use the same random seed in each simulation run for a fair comparison of the system

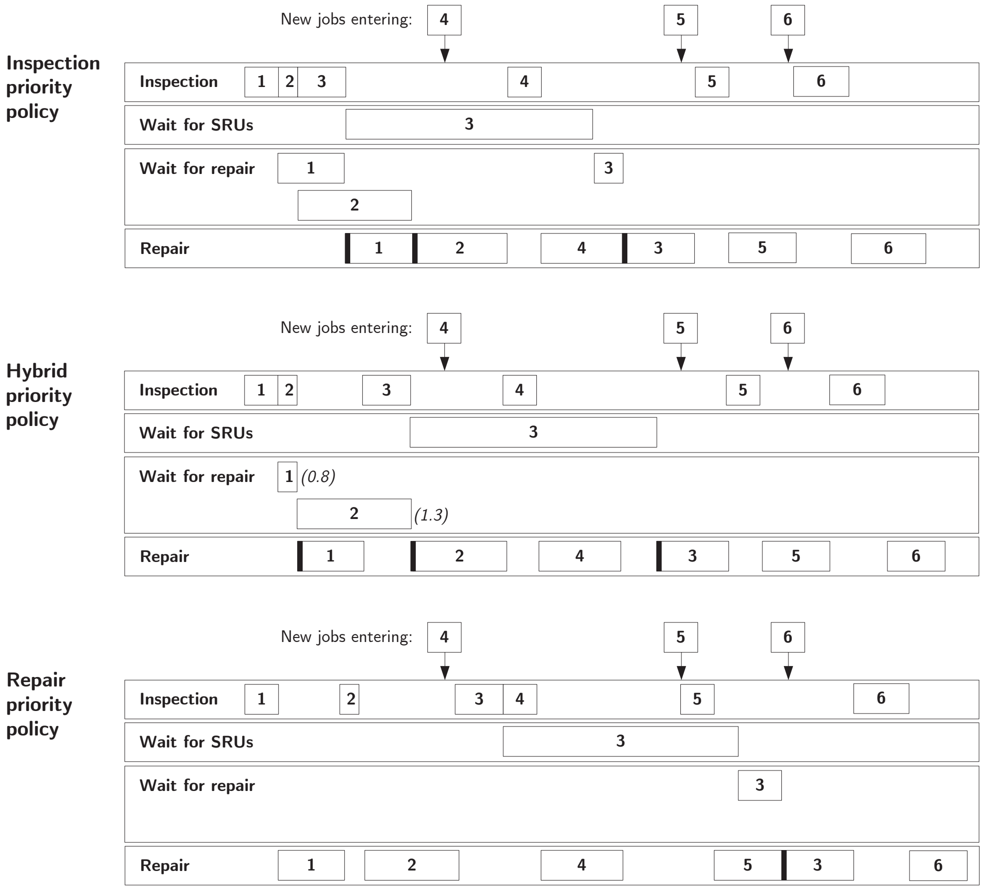

Figure 7.2 *Examples of the Gantt-charts for the three capacity priority policy variants*

performance under different choices of the SRU base stock levels and capacity priority policy. That means that we repeatedly use the same stream of random failures of LRUs, random inspection and repair times and random sets of SRUs required to repair failed parts.

We use the method of non-overlapping batch means (NOBM) to construct a $(1-\epsilon)$%-confidence interval around the steady-state mean total costs; see Steiger and Wilson (2001) for a description of this method. In setting up the simulation experiments, we follow the guidelines in Schmeiser (1982) regarding the batch size and length of the simulation run, in order to avoid simulation biases and to make sure that the batch means are approximately i.i.d. normal variables. We divide the sequence of simulated outputs into a warm-up period of 100,000 demands and 20 batches of initial size 100,000 demands.

We determine the steady-state distribution of the (assigned and unassigned) on-hand stock $B_j(\hat{\mathbf{S}}, \pi, k)$ of SRU j, $\forall j \in J$, for each batch separately. The steady-state distribution $A_i(\hat{\mathbf{S}}, \pi, q)$ of the number of parts of LRU i in the repair shop is determined over all 20 batches together. Based on the steady state distributions $A_i(\hat{\mathbf{S}}, \pi, q)$ we are able to determine an optimal vector of LRU base stock levels $\mathbf{S}$. With equation (7.1), we are able to determine the simulated total costs for each batch separately. The batch mean and standard deviation of the total costs is determined and based on this, a $(1-\epsilon)$%-confidence interval around the steady-state mean total costs is constructed.

The length of the simulation run depends on the desired accuracy of the found solution, and is chosen such that the probability that the simulated total costs are within 1% of its expected value is at least 95%. As long as the desired accuracy is not reached, we double the batch size and combine available batches into new batch results. To avoid long run times, the simulation is always terminated when the batch size has reached 400,000. This stopping criteria was however not needed in the numerical study presented in section 7.6.

7.5.2 State space formulation and variables

During the simulation run the state of the system regularly changes and in this section all possible transition types are described. We describe the state of the system by $\mathbf{x} = (x_1, .., x_{|I|}, \hat{x}_1, .., \hat{x}_{|J|}, T_1, .., T_{|J|}, R_1^1, .., R_{|I|}^1, R_1^2, .., R_{|I|}^2, .., R_1^5, .., R_{|I|}^5, L_{i,w}^v)$, with:

- x_i is the number of parts in the repair shop for LRU i, with $x_i \in \mathbb{N}_0$. This variable is used to instantly update the value of $A_i(\hat{\mathbf{S}}, \pi, q)$.

- $\hat{x}_j$ is the total (assigned plus unassigned) on-hand stock of SRU j, with $\hat{x}_j \in \mathbb{N}_0$. This variable is used to instantly update the value of $B_j(\hat{\mathbf{S}}, \pi, k)$.

- $T_j = (\tau_{j,1}, \tau_{j,2}, ..., \tau_{j,u_j})$, with $u_j \in \mathbb{N}_0$ is the number of parts in the replenishment pipeline of SRU j, where $\tau_{j,u}$ is the remaining lead time for the u-th replenishment order for SRU j, $\tau_{j,u} \in \mathbb{R}$. The case $u_j = 0$ corresponds to the situation where there are no outstanding replenishment orders for SRU j, hence $\tau_{j,0} = 0$.

- $R_i^v = (r_{i,1}^v, r_{i,2}^v, ..., r_{i,w_{i,v}}^v)$, with $w_{i,v} \in \mathbb{N}_0$, is the list of repair jobs for LRU i that is in status v where $r_{i,w}^v$ is the w-th repair job for LRU i that reached status v. The possible status of a repair job includes $v \in V$, with $V = \{1, 2, 3, 4, 5\}$ and where status v corresponds to: (1) waiting for inspection in the "Failed LRUs" stock, (2) in the inspection phase, (3) waiting for SRUs, (4) waiting for repair in the "Ready-to-repair" stock and (5) in the repair phase. The case $w_{i,v} = 0$ corresponds to the situation where there are no repair jobs for LRU i with status v.
- $L_{i,w}^v$ is the repair workload, i.e. the time required to complete the repair phase, for the w-th repair job of LRU i with status v. These values are used to instantly update the value of the actual repair workload L.

In the following sections, we describe the updates of the state space variables and application of the capacity priority policy.

7.5.2.1 Update of state space variables

We have the following state space variable update types:

- Type 1: failure of LRU i. Hence the number of parts in the repair shop is increased by one, i.e. x_i is updated to $x_i + 1$, and the number of repair jobs for LRU i is also increased by one, i.e. $w_{i,1}$ is updated to $w_{i,1} + 1$ and $r_{i,w_{i,1}+1}^1$ is added to the list of repair jobs R_i^1 that are waiting for inspection in the "Failed LRUs stock".
- Type 2: start of inspection phase for a repair job for a part of LRU i. The repair jobs in the list R_i^1 are updated according to $r_{i,w}^1 = r_{i,w+1}^1, \forall w \in \{1, ..., w_{i,1} - 1\}$. That is, the first repair job $r_{i,1}^1$ is updated to $r_{i,1}^2$ and the index w of the other repair jobs in list R_i^1 is decreased by one.
- Type 3: completion of inspection phase of repair job $r_{i,1}^2$ for LRU i. Let $M \subseteq J$ be the set of SRUs required to repair the failed part. For all SRUs $j \in M$ with $u_j < \hat{S}_j$, the assigned on-hand stock is increased with one unit and the unassigned on-hand stock is decreased with one unit. If all required SRUs are assigned, i.e. if $u_j < \hat{S}_j$, $\forall j \in M$, then $r_{i,1}^2$ is updated to $r_{i,w_{i,4}+1}^4$. In all other cases, $r_{i,1}^2$ is updated to $r_{i,w_{i,3}+1}^3$. For all SRUs $j \in M$ a replenishment order is created, i.e. $\tau_{j,u} = \hat{t}_j$ with $u = u_j + 1$ is added to the list of outstanding replenishment orders T_j and u_j is updated to $u_j = u_j + 1$.

- Type 4: arrival of replenishment order of SRU j. The total (assigned plus unassigned) on-hand stock of SRU j, $\hat{x}_j$, is updated to $\hat{x}_j + 1$. If $u_j > \hat{S}_j$ then the SRU is assigned to the first repair job waiting $r^3_{i,z}$ for SRU j, and if all required SRUs for repair job $r^3_{i,z}$ are assigned then $r^3_{i,z}$ is removed from list R^2_i and added to list R^4_i. The value $\tau_{j,1}$ is removed from the list T_j and the other values in the list are updated to $\tau_{j,w} = \tau_{j,w+1}$, $\forall w \in \{1,..,u_j - 1\}$.

- Type 5: start of repair phase for a repair job for a part of LRU i. The repair jobs in the list R^4_i are updated according to $r^4_{i,w} = r^4_{i,w+1}$, $\forall w \in \{1,...,w_{i,4} - 1\}$, that is, the first repair job $r^4_{i,1}$ is updated to $r^5_{i,1}$ and the index w of the other repair jobs in list R^4_i is decreased by one. Furthermore, the total on-hand stock $\hat{x}_j$ of SRU j is updated to $\hat{x}_j - 1, \forall j \in M$, with $M \subseteq J$ the set of required SRUs to complete the repair job.

- Type 6: completion of repair phase of repair job $r^5_{i,1}$ for LRU i. Repair job $r^5_{i,1}$ is removed from list R^5_i. The number of parts in the repair shop for LRU i, x_i, is updated to $x_i - 1$. The values of $L^5_{i,w}$ are updated to $L^5_{i,w+1}$, $\forall w \in \{1,..,w_{i,5} - 1\}$, and the index w is decreased by one.

- Type 7: Update of the required workload for repair job $r^z_{i,y}$, where $z \in \{3,4\}$, because the maximum postponement time T^{Max} has expired and an additional time is required to complete the repair job. The repair workload $L^z_{i,y}$ of job $r^z_{i,y}$ is updated.

7.5.2.2 Capacity priority

Depending on the selected capacity priority policy and the state of the system, priority is given to either inspection or repair of a part. Priority needs to be given only when the server completes the inspection or repair phase of a repair job (note that we assume non-preemption for both the inspection and repair of parts), or when a repair job enters one of the stock points "Failed LRUs" or "Ready-to-repair" and the server is idle at that time. In case the stock point "Failed LRUs" ("Ready-to-repair") is empty and the server is idle, then priority is always given to the repair (inspection) of parts, provided that the stock point "Ready-to-repair" ("Failed LRUs") is not empty.

When both the stock points "Failed LRUs" and "Ready-to-repair" are not empty, then we use the following rule for setting priority. Let $r^1_{\eta,1}$ be the repair job for a part of LRU $\eta \in I$ that entered the "Failed LRUs" stock the first, out of all parts in the "Failed LRUs" stock, and let $r^4_{\theta,1}$ be the repair job for a part of LRU $\theta \in I$ that entered the “Ready-to-repair” stock the first, out of all repair jobs that are “Ready-to-repair”. The total repair workload of parts in repair plus the number of parts in the “Ready-to-repair” stock, denoted by L, is now equal to:

$$L = \sum_{i \in I} \sum_{v \in \{4,5\}} \sum_{w \leq w_{i,v}} L^{v}_{i,w}.$$

If L is larger than or equal to the threshold value $\breve{L}$, then priority is given to the repair job $r^{4}_{\theta,1}$. If L is less than the threshold value $\breve{L}$, then priority is given to repair job $r^{1}_{\eta,1}$.

7.6 Numerical study

In this section, we present the results of our simulation studies for a single base case under the options for determining the SRU base stock levels as described in section 7.3. In section 7.6.1 we first outline our base case. We have constructed a base case that represents a real-life repair shop environment with high material uncertainty. In section 7.6.2 we present the results of our simulation studies for this base case.

7.6.1 Base case description

We consider a case in which a single repair man is responsible for the repair of 20 LRUs. The LRUs vary in demand rate (5 levels) and price (4 levels) and each combination occurs once, see table 7.2 for an overview. We have numbered the LRUs A1, ..., A4, ..., E1, ..., E4 for later references. The letter relates to the demand class (A=very fast moving, E=very slow moving) and the number relates to the price class (1=cheap item, 4=expensive item). The utilization of the repair man is at most 90%, which is the case if all repair jobs are interrupted after inspection.

The total repair job workload (inspection plus repair time) of each LRU is Gamma distributed, with an average of one time unit and standard deviation 0.5. The inspection time of a repair job always equals 10% of the simulated total repair job workload. The repair time equals 90% of the simulated total workload.

For repair jobs that are interrupted for more than T^{Max} units after inspection, the repair time is increased with the inefficiency factor multiplied by the inspection time of that specific repair job. In general, the ‘inefficiency’ is higher for parts that

Table 7.2 LRU characteristics

Demand class	Demand rate per time unit	Price 40.000	160.000	280.000	400.000	Inefficiency factor (%)
Very fast moving	0.1126	A1	A2	A3	A4	15%
Fast moving	0.0563	B1	B2	B3	B4	20%
Medium moving	0.0282	C1	C2	C3	C4	25%
Slow moving	0.0141	D1	D2	D3	D4	30%
Very slow moving	0.0093	E1	E2	E3	E4	50%

Table 7.3 SRU class characteristics for all LRUs

SRU class	Number of SRUs	Price (% of LRU price)
1	10	0.25%
2	10	0.5%
3	10	1.0%
4	10	1.5%
5	5	2.0%
6	1	5.0%
7	1	10.0%
8	1	15%
9	1	20%
10	1	25%

are repaired less often, because the part is more unknown to the repair man. The inefficiency factors hence increase when LRU demand rates decrease and can be found in table 7.2; the average inefficiency factor (weighted for the various demand rates) is equal to 20%.

Each LRU contains 50 SRUs and each SRU has a probability of 20% that it is required during a repair job (i.e. $\alpha_{i,j} = 0.2 \quad \forall i \in I$ and $\forall j \in J$). We assume that each SRU is required for at most one LRU (no commonality), hence we have $50 \cdot 20 = 1000$ SRUs. The prices of the SRUs are a fixed percentage of the LRU price for which the SRUs are required. Table 7.3 provides an overview of the SRU classes we have, including the number of SRUs in that class and their price (related to the price of the LRU). Each SRU in a certain class has the same characteristics.

All SRUs are consumable and have a deterministic replenishment lead time, varying from 5 to 100 time units. The replenishment lead times are higher for SRUs with a higher price and/or lower demand rate. The weighted average lead time over all

Table 7.4 SRU replenishment lead times per SRU class per LRU (in time units)

	SRU class									
LRU	1	2	3	4	5	6	7	8	9	10
A1	5	10	15	20	25	30	35	45	55	65
A2	10	15	20	25	30	35	40	50	60	70
A3	15	20	25	30	35	40	45	55	65	75
A4	20	25	30	35	40	45	50	60	70	80
B1	10	15	20	25	30	35	40	50	60	70
B2	15	20	25	30	35	40	45	55	65	75
B3	20	25	30	35	40	45	50	60	70	80
B4	25	30	35	40	45	50	55	65	75	85
C1	15	20	25	30	35	40	45	55	65	75
C2	20	25	30	35	40	45	50	60	70	80
C3	25	30	35	40	45	50	55	65	75	85
C4	30	35	40	45	50	55	60	70	80	90
D1	20	25	30	35	40	45	50	60	70	80
D2	25	30	35	40	45	50	55	65	75	85
D3	30	35	40	45	50	55	60	70	80	90
D4	35	40	45	50	55	60	65	75	85	95
E1	25	30	35	40	45	50	55	65	75	85
E2	30	35	40	45	50	55	60	70	80	90
E3	35	40	45	50	55	60	65	75	85	95
E4	40	45	50	55	60	65	70	80	90	100

SRUs equals around 30 time units. The replenishment lead times range from 5 to 65 time units for SRUs related to LRU A1 and range from 40 to 100 time units for SRUs related to LRU E4. See table 7.4 for a complete overview of all SRU replenishment lead times.

The holding costs h_i and $\hat{h}_j$ are equal to 0.1% of the part price per time unit. This corresponds to an annual holding costs rate of 12.5% (assuming 125 time units is equivalent with one year). The backorder costs b are equal to 5,500 per time unit for each LRU.

7.6.2 Scenario results

In this section, we present the results we have obtained under the different options for determining the SRU base stock levels. In section 7.6.2.1 we present the results for the option where all SRUs have an identical fill rate target γ, for $\gamma = \{0\%, 50\%, 75\%, 90\%, 95\%\}$. Note that the scenario with fill rate target 0% is equivalent to the option in which all SRU base stock levels are 0. The holding costs over the unassigned on-hand SRU stock based on the scenarios, in which all SRUs have an identical fill rate target, serve as a target for the holding costs over the unassigned on-hand SRU stock in the scenarios in which we optimize $\beta(\hat{\mathbf{S}})$. We study two options: the 'inefficiency' either is obtained directly (see sections 7.6.2.1 and 7.6.2.2), i.e. $T^{Max} = 0$, or only after a certain period (section 7.6.2.3), in this specific case T^{Max} equals 5.0 time units. In the latter case the 'inefficiency' is absent, if the repair phase starts within 5.0 time units after the inspection phase was completed.

For each scenario, we present the total system costs per time unit for various values of the workload threshold $\breve{L}$. From this, one can determine the optimal workload threshold value for each scenario. Further, we present more detailed results for each scenario under the optimal workload threshold value. The 'Total SRU holding costs' here represent the average holding costs per time calculated over the assigned and unassigned on-hand SRU stock. Most of the other characteristics are intuitively clear, some of them are explained hereafter:

1. Time repair shop empty: percentage of time that the repair shop is empty (i.e. $c_i = 0, \quad \forall \quad i \in I$).
2. Average waiting time SRUs job complete: average time from completion of the inspection phase of a random repair job until the moment that all required SRUs are available from stock.
3. Average time waiting capacity: average time that a random repair job is waiting for capacity (in the buffer "Failed LRUs" or in the buffer "Ready-to-repair")
4. Average time in service: average time that a random repair job is in service (inspection plus repair) by the repair man.

7.6.2.1 Scenario results if all SRUs have an identical fill rate target γ

Table 7.5 shows the total system costs for the scenarios in which all SRUs have an identical fill rate target. For each workload threshold value, it can be seen that the total costs decrease when increasing the fill rates to 90%, but increase when the fill rates are increased to 95%. Under low individual fill rate targets, the inspection priority policy ($\breve{L} = \infty$) is the best choice. For moderately high fill rate targets, a hybrid priority policy is the best choice and under high fill rates (90% en 95%) the repair priority policy is the best choice. An individual fill rate target of 90% leads to the optimal total system costs of 21,058.79 per time unit; the difference in total system costs under this scenario but with different workload threshold values is at most 4.15%.

Table 7.5 Total costs per time unit if all SRUs have an identical fill rate target γ

$\breve{L}$	$\gamma = 0\%$ (0)	$\gamma = 50\%$ (1)	$\gamma = 75\%$ (2)	$\gamma = 90\%$ (3)	$\gamma = 95\%$ (4)
0.0	55,742.89	26,318.37	22,433.93	**21,058.79**	**21,288.10**
1.0	54,965.74	26,108.42	22,476.21	21,228.97	21,433.41
2.0	53,112.54	25,257.55	22,285.28	21,526.94	21,876.41
3.0	51,205.35	24,581.50	22,100.51	21,570.12	22,017.17
4.0	49,119.39	24,106.39	21,986.73	21,582.43	22,060.77
5.0	46,456.32	23,813.60	21,923.57	21,601.22	22,100.80
6.0	43,091.67	23,621.02	21,891.22	21,623.53	22,137.44
7.0	39,254.73	23,515.44	21,882.95	21,652.68	22,174.80
8.0	35,845.16	23,449.86	**21,880.86**	21,682.65	22,212.37
9.0	33,450.33	23,424.54	21,893.58	21,710.10	22,244.59
10.0	31,760.20	**23,408.14**	21,902.55	21,734.63	22,275.90
11.0	30,827.27	23,413.57	21,916.15	21,757.75	22,303.67
12.0	30,357.40	23,414.19	21,928.40	21,780.53	22,328.90
∞	**29,661.31**	23,475.62	22,011.89	21,931.75	22,493.75

Table 7.6 shows the detailed results for the same scenarios under the optimal value of the workload threshold, i.e. the workload threshold that minimizes the total system costs. From this table, we see that increasing the individual target fill rates for all SRUs positively influences the repair job lead times in two ways. Firstly, the average waiting time until all SRUs for a random job are 'complete', decreases. Secondly, the number of inefficient repairs decreases (especially under the repair priority policy), which results in a lower actual server utilization and a lower waiting time for capacity.

7.6.2.2 Scenario results if SRU job completeness is maximized

Table 7.7 shows the total system costs per time unit for the scenarios in which we optimize the SRU job completeness, given a constraint on the holding costs over the unassigned on-hand SRU stock. The relationship between the scenarios in section 7.6.2.1 and the scenarios in this section, are marked by the numbers between parentheses in the table headers.

In all scenarios, in which we maximize the SRU job completeness, we see that the repair priority policy is always the best choice. This is in line with our findings in section 7.6.2.1, in which we saw that the repair priority policy is the best choice if

Table 7.6 Detailed results for optimal capacity priority policy per scenario

Results per scenario	$\gamma = 0\%$ (0)	$\gamma = 50\%$ (1)	$\gamma = 75\%$ (2)	$\gamma = 90\%$ (3)	$\gamma = 95\%$ (4)
Optimal workload threshold	∞	10.0	8.0	0.0	0.0
Cost deviation w.r.t. repair priority policy	87.93%	12.43%	2.53%	0.00%	0.00%
Cost deviation w.r.t. inspection priority policy	0.00%	0.29%	0.60%	4.15%	5.66%
Total costs per time unit	29,661.31	23,408.14	21,880.86	21,058.79	21,288.10
Total LRU holding costs	24,320.00	14,080.00	11,280.00	8,040.00	6,960.00
Total backorder costs	4,589.84	4,109.86	3,918.39	3,175.78	2,678.61
Total SRU holding costs	751.47	5,218.27	6,682.47	9,843.02	11,649.49
Total assigned SRU holding costs	751.47	806.08	548.25	150.19	74.73
Total unassigned SRU holding costs	0.00	4,412.20	6,134.23	9,692.83	11,574.76
Target unassigned SRU holding costs (C^{obj})	0.00	4,408.69	6,131.27	9,690.07	11,571.96
Time repair shop empty	0.0000%	0.0000%	0.0313%	1.1287%	3.7290%
Aggregate SRU fill rate	0.00%	73.22%	87.33%	95.96%	97.82%
SRU Job completeness	0.00%	22.93%	44.30%	74.85%	84.63%
Percentage repair directly after inspection	0.00%	5.08%	11.29%	74.85%	84.63%
Average waiting time SRUs job complete	53.041	17.069	9.062	3.066	1.483
Average time waiting capacity	6.008	5.676	5.605	5.342	4.984
Average time in service	1.020	1.019	1.018	1.005	1.003
Average repair job lead time	60.068	23.763	15.685	9.413	7.470

Table 7.7 Total costs per time unit if SRU job completeness is maximized

$\breve{L}$	Maximize $\beta(\hat{\mathbf{S}})$ (0)	Maximize $\beta(\hat{\mathbf{S}})$ (1)	Maximize $\beta(\hat{\mathbf{S}})$ (2)	Maximize $\beta(\hat{\mathbf{S}})$ (3)	Maximize $\beta(\hat{\mathbf{S}})$ (4)
0.0	55,742.89	**24,575.68**	**21,931.89**	**20,328.76**	**20,835.42**
1.0	54,965.74	24,711.71	22,061.67	20,396.63	20,877.44
2.0	53,112.54	24,842.96	22,380.23	20,813.10	21,228.35
3.0	51,205.35	24,815.21	22,431.84	21,032.41	21,520.89
4.0	49,119.39	24,789.18	22,444.73	21,083.36	21,576.43
5.0	46,456.32	24,771.91	22,460.92	21,121.93	21,614.80
6.0	43,091.67	24,764.54	22,476.26	21,159.05	21,651.06
7.0	39,254.73	24,759.11	22,491.79	21,196.12	21,684.30
8.0	35,845.16	24,756.85	22,510.54	21,228.39	21,711.72
9.0	33,450.33	24,758.97	22,528.20	21,256.60	21,739.12
10.0	31,760.20	24,765.53	22,545.14	21,284.35	21,765.85
11.0	30,827.27	24,774.55	22,561.02	21,309.02	21,789.00
12.0	30,357.40	24,782.96	22,575.59	21,332.24	21,810.94
∞	**29,661.31**	24,874.56	22,699.29	21,503.51	21,977.17

aggregate SRU fill rates are high. In the scenarios in which we optimize the SRU job completeness, the aggregate fill rates range from 94.91% to 99.63% (see Table 7.8).

By maximizing the SRU job completeness, we see that the total system costs decrease for scenarios (3) and (4): from 21,058.79 to 20,328.76 and from 21,288.10 to 20,835.42, respectively. From the evaluated scenarios, we see that scenario (3), in which we maximize the SRU job completeness, has the lowest total system costs. In this scenario, a repair priority policy is the best choice.

For relatively low investments in SRU stocks, like in scenarios (1) and (2), we see that the total system costs increase when we maximize the SRU job completeness. This can be explained by two factors. First, we see that different capacity priority policies (i.e. a different optimal workload threshold) are the best choice in these scenarios. If we compare the total system costs under a repair priority policy ($\breve{L} = 0.0$) in for example scenario (2), then we see that the total system costs decrease from 22,433.93 to 21,931.89 instead.

Second, in scenario (1) and (2) under maximization of the SRU job completeness, we see that especially expensive SRUs, that are (potentially) required for the repair of expensive, slow moving LRUs (D4, E3 and E4) are not being stocked. This causes high waiting times for these specific LRUs and consequently higher (holding plus backordering) costs for LRUs. We don't see this phenomena in scenarios (3) and (4).

Table 7.8 Detailed results for optimal capacity priority policy per scenario

Results per scenario	Maximize $\beta(\hat{\mathbf{S}})$ (0)	Maximize $\beta(\hat{\mathbf{S}})$ (1)	Maximize $\beta(\hat{\mathbf{S}})$ (2)	Maximize $\beta(\hat{\mathbf{S}})$ (3)	Maximize $\beta(\hat{\mathbf{S}})$ (4)
Optimal workload threshold	∞	0.0	0.0	0.0	0.0
Cost deviation w.r.t. repair priority policy	87.93%	0.00%	0.00%	0.00%	0.00%
Cost deviation w.r.t. inspection priority policy	0.00%	1.22%	3.50%	5.78%	5.48%
Total costs per time unit	29,661.31	24,575.68	21,931.89	20,328.76	20,835.42
Total LRU holding costs	24,320.00	16,440.00	11,760.00	8,000.00	6,400.00
Total backorder costs	4,589.84	3,166.62	3,679.68	2,485.13	2,817.19
Total SRU holding costs	751.47	4,969.06	6,492.21	9,843.63	11,618.23
Total assigned SRU holding costs	751.47	557.38	324.60	94.67	39.25
Total unassigned SRU holding costs	0.00	4,411.68	6,167.61	9,748.96	11,578.98
Target unassigned SRU holding costs (C^{obj})	0.00	4,408.69	6,131.27	9,690.07	11,571.96
Time repair shop empty	0.0000%	0.0000%	0.0149%	2.7776%	6.2960%
Aggregate SRU fill rate	0.00%	94.91%	97.29%	99.21%	99.63%
SRU Job completeness	0.00%	66.52%	80.28%	93.32%	96.75%
Percentage repair directly after inspection	0.00%	66.52%	80.28%	93.32%	96.75%
Average waiting time SRUs job complete	53.041	15.414	8.385	1.824	0.781
Average time waiting capacity	6.008	6.230	5.417	4.847	4.751
Average time in service	1.020	1.007	1.005	1.002	1.001
Average repair job lead time	60.068	22.652	14.807	7.673	6.533

In these scenarios all SRU base stock levels are at least equal to 1 and, because the largest part of individual SRU waiting times is decreased by increasing the base stock level from 0 to 1, extremely long repair lead times for expensive slow moving LRUs (like D4, E3 and E4) do not occur often. This significantly reduces the LRU costs.

Inspection and hybrid priority policies tend to work well if many SRU base stock levels are equal to zero. We observed this phenomenon in scenarios 1 and 2, in which the total costs increased when the SRU job completeness was maximized (and the number of SRUs with base stock level 0 increased). Repair jobs that require SRUs with base stock level 0, directly benefit themselves from earlier ordering of these SRUs. Demands for SRUs with base stock level greater than 0, are fulfilled from replenishment orders that are generated by demands from previous repair jobs. Hence, a repair job, that only requires SRUs with base stock levels greater than 0, does not directly benefit from an early inspection.

Table 7.9 Total costs per time unit if inefficiency only drops-in after 5 time units

$\breve{L}$	Maximize $\beta(\hat{\mathbf{S}})$ (0)	Maximize $\beta(\hat{\mathbf{S}})$ (1)	Maximize $\beta(\hat{\mathbf{S}})$ (2)	Maximize $\beta(\hat{\mathbf{S}})$ (3)	Maximize $\beta(\hat{\mathbf{S}})$ (4)
0	55,742.89	24,563.30	**21,921.67**	**20,321.08**	**20,830.51**
1	54,965.74	24,554.74	21,924.34	20,326.97	20,833.96
2	53,112.54	24,495.84	21,929.91	20,377.35	20,880.48
3	51,205.35	24,451.57	21,935.75	20,430.08	20,943.26
4	49,119.39	**24,435.07**	21,956.14	20,474.84	20,988.37
5	46,456.32	24,574.50	22,186.48	20,792.40	21,303.05
6	43,091.67	24,676.95	22,370.09	21,055.97	21,567.51
7	39,254.73	24,678.39	22,394.42	21,106.68	21,611.83
8	35,845.16	24,678.73	22,413.99	21,139.17	21,642.30
9	33,450.33	24,681.44	22,434.87	21,167.06	21,669.65
10	31,760.20	24,686.84	22,450.18	21,194.70	21,696.19
11	30,827.27	24,696.68	22,466.48	21,219.62	21,719.51
12	30,357.40	24,705.81	22,483.23	21,242.57	21,740.81
∞	**29,661.31**	24,796.75	22,607.43	21,407.46	21,902.17

7.6.2.3 Scenario results if inefficiency only drops-in after 5 time units

Table 7.9 shows the total system costs for the scenarios in which we optimize the SRU job completeness, given a constraint on the holding costs over the unassigned on-hand SRU stock, and in which the inefficiency is only obtained after $T^{Max} = 5.0$ time units. For scenario (1), we now see that a hybrid priority policy (with $\breve{L} = 4.0$)

is the best choice. In this situation, the total system costs are 0.5% lower than in the repair priority policy. However, we still see that, under the evaluated scenarios, scenario 3 has the lowest total system costs. In this scenario repair priority is the best choice.

7.7 Conclusions & suggestions for further research

In this chapter, we have studied the control concept for a repairable parts inventory system with a repair shop of type II as proposed in Chapter 6. In this control concept, we have proposed to decouple the inspection and repair phase of a repair job. From the literature it is not clear whether priority should be given to either inspection of failed parts (to reveal which SRUs are required and speed up sourcing of unavailable SRUs) or to completion of the repair of parts for which all required SRUs are available (to increase the availability of LRUs), and therefore we formulated the following research question:

6. *What is the optimal capacity priority policy in repairable parts inventory systems with high material uncertainty?*

In this chapter, we have evaluated workload-based capacity priority policies in a single base case for several scenarios to set the SRU base stock levels. This base case was constructed such that the repair shop faces a high material uncertainty when conducting repairs. From this evaluation study, we can draw the following conclusions:

1. In a setting with a high aggregate SRU fill rate, a smooth repair process is enabled (i.e. low number of repair job interruptions because of missing SRUs) and repair priority policies outperform inspection and hybrid priority policies. This also holds for situations in which the inefficiency is obtained only after a certain time period.

2. In a setting with a low to moderately high aggregate SRU fill rate, inspection and hybrid priority policies outperform repair priority policies.

From the scenarios we have evaluated, the scenario with a high aggregate SRU fill rate combined with a repair priority policy leads to the lowest total costs. To come to these conclusions, we made some simplifying assumptions. These assumptions possibly influence the optimal capacity priority policy, and they can be relaxed in further research.

First, the analysis can be extended to cases with condemnation and ‘no fault found’ rates larger than 0. We suspect that these parameters have significant impact on the optimal priority policy, as the repair jobs for condemned parts or parts with ‘no fault found’ are directly transferred into a replenishment order for an LRU or sent to the ready-for-use LRU stock. In both cases, the time required until the part is RFU again is shorter in a inspection/hybrid priority policy than in a repair priority policy. It is however not clear whether this leads to an extra buffer of parts in the repair process (after the pre-inspection phase, see also Section 5.3), or whether buffers can be combined. This is a suggestion for further research.

Another extension could be to incorporate priority rules for executing repair jobs (instead of the FCFS service discipline). Our model is, with some simple extensions, also applicable for static repair priority policies. For example, (at least) two priority classes could be created, of which one class consists of the LRUs for which the expected SRU job completeness is relatively low. Note that the assignment to priority classes can be done upfront, as we also determine the SRU base stock levels upfront. In a new, to be defined capacity priority policy one could give priority to inspection of the parts that have a relatively high probability of being interrupted because of missing SRUs. Another extension could be the use of dynamic (or myopic) priority policies, but note that this requires a different approach.

In our model, we assume that capacity is fixed and a possible extension is the application of overtime policies. How this impacts our model and the choice for a capacity priority policy is not obvious. It is likely that overtime policies are more effective when smart priority rules are applied. That is, it is not effective to work overtime on repair jobs on an FCFS basis, rather priority should be given to the repair of LRUs with backorders or with the shortest inventory run-out times.

We have evaluated three options to set the SRU base stock levels, however we did not use an ‘integrated’ approach to set the LRU and SRU base stock levels (like we proposed in the design of repair shop type II). The capacity priority policies under the *optimal* SRU base stock levels have not been evaluated. These optimal SRU base stock levels also depend on the capacity priority policy itself, and a model to integrally solve this problem is not available in the literature. One possibility to find (close to) optimal base stock levels is via simulation based optimization techniques. A drawback of this approach is that computation times may quickly rise in the number of SRUs, and this approach thus may become infeasible. We moreover believe that a different method does not impact the choice for a certain capacity priority policy in our base case, as a high aggregate SRU fill rate will probably be the result of any method, and in that case the repair priority policy outperforms the other capacity priority policies.

In our analysis, we assumed that SRUs are assigned to repair jobs on a FCFS basis. Evaluation of an alternative SRU assignment policy is another suggestion for further research, and such policy will likely increase the SRU job completeness. We therefore suspect that this is mainly in favor of the repair priority policy.

Chapter 8

Conclusions & suggestions for further research

"Wij hebben het uitstekend gedaan, simpelweg omdat we niet beter konden."

Johan Cruijff

This research has focused on concepts and models that help MLOs to improve the balance between the availability and costs of spare parts in user networks. A normative decision framework for MLOs is presented in Chapter 2, which is used to compare the framework to practice in six companies that collaborated in this research (Chapter 3). We have concluded that the study of repairable parts inventory systems is highly relevant from both a scientific and practical point of view. In Chapter 4 we have reviewed some important literature streams, and in Chapter 5 we have identified gaps with practice when it comes to designing control concepts for repairable parts inventory systems. In Chapter 6 we have presented designs of control concepts for repairable parts inventory systems, for each of four types of repair shops found in practice. In Chapter 7, we have verified the design of one of the repair shop types in a quantitative study. This concluding chapter summarizes the results and insights that are obtained from this research, and addresses suggestions for further research.

8.1 A framework for planning & control of MLOs

The main objective of this dissertation has been to develop control concepts for MLOs that lead to efficient combinations for availability and costs of spare parts in user networks, by developing concepts and models that have both a high practical as well as scientific relevance. These concepts and models relate to one or more decisions that have to be made by MLOs. To find the most relevant decision functions, we have started our research with the following research question:

1. *What are the relevant decision functions in MLOs, and how do they relate to each other?*

The relevant decision functions in MLOs are related to supply chain management processes (e.g. forecasting and inventory control) and to repair shop control. In the decision framework that is presented in Chapter 2, we cluster all relevant decision functions, describe how they relate to each other and we have provided references to state-of-the-art techniques for supporting these decisions. With this framework, we have addressed our first research question.

The framework is based on a literature review, and is especially useful for practitioners. In Chapter 3 its completeness, applicability and value is verified (demonstrated) in case studies at six prominent Dutch companies that use and maintain high value capital assets. With the case studies, we have also aimed to identify open topics for the remainder of this research. As such, we have formulated the following research questions:

2a. *Which processes/decision functions in the framework are observed in practice?*

2b. *Which processes/decision functions observed in practice are not included in the framework?*

2c. *Which (interfaces and information flows between) processes/decision functions are filled in differently by the case study companies, and why?*

Each decision function is found in at least two of the case study companies, and 13 out of 18 decision functions are found in all case studies. We have not found a generic decision function that is not included in the framework. Comparing the framework to case studies has revealed that decisions related to repair shop (capacity dimensioning and scheduling of repair work) and inventory control (setting SRU and LRU stock levels) are implemented differently in the various companies, and are most relevant for MLOs in terms of availability and costs of spare parts. These decisions relate to decisions in *repairable parts inventory systems*, and clearly these decisions interrelate. A control concept is required to optimally balance the decisions in these systems.

In the framework *repair shop control* and *inventory control* are separate processes, and *supply management* acts as an interface. We have made two design choices regarding this interface. First, we have assumed that planning and control of SRUs and LRUs is integrated and second, we have assumed that inventory control makes agreements on the tactical level with the repair shops about planned repair lead times. Application of our framework has turned out to have practical implications: decisions are in practice separated over two decision makers, and situations may exist in which control and responsibilities are not aligned. Moreover, repair shop control and inventory control are sometimes assumed to be fully integrated, to enable efficient control of repair job priorities (e.g. Hausman and Scudder (1982)). How to set up a control concept for repairable parts inventory systems is hence also relevant from a scientific point of view, and this is the motivation for research questions 3-5.

8.2 Designs for repairable parts inventory systems

Chapters 4-6 focus on designing control concepts for repairable parts inventory systems. In Chapter 4 we have reviewed three streams of literature: uncapacitated inventory models (stream 1), single-stage capacitated inventory models (stream 2) and multi-stage capacitated inventory models (stream 3), and we have raised the following research question:

3. *What are ways to design control concepts for repairable parts inventory systems, and which requirements are important to account for in the design process?*

In the first stream of literature, repair shop control and inventory control decisions are decomposed by means of a 'lead time'-interface with either one (regular) or two (regular and emergency) lead times, like in the framework of Chapter 2, and these decisions are solved independently. This stream of literature is applicable if the repair shop is able to meet these lead times with a high probability, which may be difficult when the repair shop utilization is high and capacity flexibility options are limited (or expensive), and either inspection is performed upon failure of a part, or SRUs do not cause a delay of repairs.

Literature streams 2 and 3 assume that inventory and repair shop control decisions are integrated and no interface is required between them. Application of the models from these streams is limited in two ways. First, the models are limited to situations in which SRUs do not cause delays of repairs (stream 2), or to situations in which both inspection and repair and assembly of parts at each indenture level require separate resources (stream 3). Second, the computational complexity to solve these models quickly increases in the number of decision variables, and this limits their practical applicability.

From the three literature streams we have reviewed, we have found three important principles to account for when designing control concepts for repairable parts inventory systems. From the first stream (uncapacitated inventory models), we found that integrating SRU and LRU stocking decisions is beneficial to significantly reduce inventory holding costs, especially when SRUs are expensive and have low demand rates. From the second stream (single-stage capacitated inventory models), we have found that integrating repair priority setting and LRU stocking decisions significantly reduces inventory investments, especially when there exist large demand and price differences between the LRUs. From the third stream (Generalized Kanban Control Systems), we have found that repair job release (workload control) should be based on aggregate flows per specialty or resource (i.e. stages), especially when multiple (highly utilized) resources are required in the repair process.

The three streams of literature we have reviewed are applicable in specific situations. So far, we do not know whether the models and control principles identified in the literature review can be applied to the repairable parts inventory systems found in the case study companies. Therefore, we have raised the following research question:

4. *Which models and control principles, identified in the literature review of Chapter 4, are applicable to design control concepts for the repairable parts inventory systems in the case study companies?*

We have collected quantitative and qualitative during interviews and plant visits in multiple case studies, and we have concluded that some important modeling aspects are not yet covered in the literature. First, the inspection phase of a repair job takes a significant amount of time from a repair man and is similar to the 'work preparation' phase in a production process, i.e. only after this phase the repair routing and required materials are known. Second, we have found that the repair routing and the required set of materials strongly vary from one repair job to another, and multiple SRUs may be required to repair a failed part. Third, in many situations the inspection and repair phase require the same capacity, and the repair man has to go into the problem again (i.e. a part of the inspection has to be redone) when the repair phase is not conducted shortly after the inspection phase. These aspects deserve attention in further research.

From the mapping of the models and control principles to the data we gathered from the case studies, we can conclude that there is no repair shop to which the models and control principles can be applied directly, because at least one assumption is not valid. We have found that a varying set of repair shops complies to a subset of requirements and control principles per literature stream. This can be explained by the clear differences we have observed between repair shops, especially when it comes to the complexity and uncertainty of the materials and capacities required in the repair process. How this impacts the designs of control concepts for repairable parts inventory systems is not clear, and therefore we have raised the fifth research question:

5. *To what extent do repair shops differ from each other, and what implications do these differences have on the design of control concepts for repairable parts inventory systems?*

From the case studies, we have found that repair shops differ in two dimensions: *capacity complexity* and *material uncertainty*. Both dimensions can take the values *low* and *high*, which results in a 2x2 typology of repair shops. This typology was inspired by the typology for production units (Bertrand and Wortmann, 1992; Bertrand et al., 1990). The capacity complexity concerns the requirement of specialized skills of repair men to execute a repair job. The material uncertainty concerns the extent to which repair jobs for the same item require different sets of materials, i.e. the extent to which it is predictable which materials will be required in the repair process.

For each of the four repair shop types, we have formulated design choices and have given a qualitative motivation for these choices. For repair shops with a high material uncertainty, we have proposed to integrate LRU and SRU stocking decisions and to decouple the inspection and repair phase. For repair shops with a low capacity complexity, repair priority setting and inventory decisions are integrated, and are assigned to the inventory department. For repair shops with a high capacity complexity, we have proposed a ‘lead time interface’ with the option to expedite repair jobs (i.e. a regular and expedited lead time). Based on these design choices, refinements of the framework in Chapter 2 are made explicit by presenting designs for the control concept of repairable parts inventory systems in Chapter 6.

The use of these control concepts as a reference point for designing control concepts of existing repairable parts inventory systems is verified in workshops at two companies. Although these designs are verified in case studies, we can not substantiate the claim that these designs are ‘optimal’, nor that they are always applicable. Further refinement of the designs may be necessary for each specific repair shop in practice, and quantifying the impact of the design choices is necessary to test and further improve the designs. In Chapter 7 we have evaluated the design of a repairable parts inventory system with a repair shop of type II, i.e. a repair shop with a low capacity complexity and a high material uncertainty, but evaluating the design of the other repair shop types deserves attention in further research.

8.3 Capacity assignment under high material uncertainty

In the design for repair shop type II in Chapter 6, we have proposed to decouple the inspection and repair phase with a stock point for parts that are ‘ready-to-repair’. This has raised the question whether priority should be given to inspection of failed parts (to reveal which SRUs are required and speed up sourcing of unavailable SRUs) or to completion of the repair of parts for which all required SRUs are available (to increase the availability of LRUs). At the same time, consciously interrupting repair jobs is inefficient; the repair times depend on the time that elapses between inspection and the actual repair of a part. This has raised the final research question of this dissertation:

6. *What is the optimal capacity priority policy in repairable parts inventory systems with high material uncertainty?*

A repair priority policy combined with a high aggregate SRU fill rate of around 99%, that is set using a ‘system approach’, leads to the lowest integral costs for the base case and the scenarios that we evaluated. The repair priority policy also outperforms all other policies in all scenarios in which the aggregate SRU fill rate is moderately high to high (of at least 95%, say), also when the inefficiency of interrupting repair jobs is obtained only after a certain time period. Hence, one may argue that the decoupling of the inspection and repair phase, and consciously assigning capacity to either inspection and repair of failed parts, is not beneficial. However, there are still issues that deserve further research.

Inspection/hybrid policies tend to work well if many SRU base stock levels are 0, or when the aggregate SRU fill rate is relatively low. These situations are conceivable in real-life situations, because some of our implicit assumptions are not valid. For example, the configuration of LRUs (i.e. the lower indenture level parts) is not always (completely) known and this might lead to more SRUs having a base stock level of 0. Furthermore, if SRU failure rates are not known exactly this could lead to a lower aggregate SRU fill rate, when actual failure rates exceed expected failure rates. In practice we for example observe aggregate SRU fill rates of around 90-95% for repair shops with a high material uncertainty. In such cases, inspection or hybrid priority policies may sometimes outperform repair priority policies. Hence, how these assumptions impact the choice for a certain capacity priority policy is not clear and deserves further research.

8.4 Further research and future perspectives

The suggestions for further research in the previous sections have focused solely on the research we executed, and mainly relax the assumptions that we have made in developing the concepts and models. In this section, we put our conclusions and suggestions for further research in perspective to developments that are ongoing in the field of (or fields related to) service logistics. We briefly discuss how these developments impact the concepts, models and results that are presented in this dissertation, and how these developments might give interesting directions for further research. We conclude with a reflection on how we organized our research.

Recent developments in condition based maintenance and (remote) condition monitoring, combined with developments in the area of data science (e.g. data mining, machine learning), enable maintenance organizations to predict the need for maintenance more carefully and in real time. As such, also the need for and timing of the LRU replacements can be predicted with more certainty. This *advance demand information* (ADI) should enable repair shops to effectively prioritize repair jobs based on possible future stockout occasions (known from this ADI). This is however difficult in a standard 'lead time interface' (with only a regular repair lead time), but it may be possible in a lead time interface with a regular and an expedited lead time. Furthermore, it enables repair shops to better predict the workload that is offered to the shop, and consequently this should improve their ability to efficiently plan repair work, and to effectively apply proactive overtime policies or emergency repairs. This is an interesting direction for further research, and may also lead to some modifications of the control concept designs presented in Chapter 6.

In the near future, such techniques possibly also reveal which SRUs need to be replaced in the repair process. These developments would make the inspection phase shorter or even unnecessary, and that would also make some of the results of this research out of date. Although first results from these techniques are promising and show an increased potential, application on a large scale will probably take at least five up to ten years. Next to that, one may wonder whether using these techniques for all parts is cost effective. Investments can be high and the monitoring systems itself can fail and need to be maintained as well. Research in this area currently focuses on monitoring the condition of LRUs, and only to a less extent on the SRUs. Therefore we believe that this development has only limited impact on the need for continuation of the research directions mentioned in the previous chapter.

Additive manufacturing (AM) is another recent development, and enables MLOs to manufacture parts on demand within a short time period (a production time of one or two days, say, should be possible in the near future). However, printing multi-indentured LRUs turns out to be complex and might even be impossible, at least in the near future. SRUs are less complex to print than LRUs, and this may create opportunities. For example, combining an inspection/hybrid based capacity priority policy with AM may be beneficial, as these policies will probably work well when lead times for SRUs are short and SRU base stock levels are equal to 0. However, printed parts are possibly less reliable (i.e. have higher failure rates). Furthermore, one should take into account possible congestion effects if too many SRUs need to be printed instantaneously by the same printer, possibly meaning that not all SRUs can be supplied within a very short lead time and SRUs should be stocked upfront anyhow. Moreover, although AM is evolving, it is questionable whether application of AM for electronic parts on a large scale can be expected in the near future. With the growing interest and potential we nevertheless expect that the number of third parties offering AM-services will increase and this development may eventually take away these constraints. We believe that the need for further research on decoupling the inspection and repair phase increases when application of AM evolves.

Another development is that asset life cycles increase, not only because assets become more reliable but also because the return on investment is relatively lower due to increased research and development costs for new assets. At the same time life cycles for spare parts decrease due to rapid technological developments, especially for electronic parts. This increases the risk that suppliers suddenly stop supplying these parts. A high material uncertainty seems to be common in electronics or avionics repair shops (repair shop types II and IV), and hence these shop types face significantly higher supply risks for the required SRUs. This may have a negative impact on the actual aggregate SRU fill rate that is realized by MLOs. This tends to be in favor of inspection/hybrid capacity priority policies, and hence these developments increase the need for further research on decoupling the inspection and repair phase.

Nevertheless, the question arises whether all mentioned developments are profitable for MLOs. OEMs may use the application of condition monitoring techniques as a marketing instrument to expand their influence and maintenance activities (both asset maintenance and maintenance of repairable parts) in user networks (*servitization*). This will limit the role of MLOs in user networks, and internal repair shops will face a lower workload or may even disappear. Repair work will be conducted by external repair shops, with whom MLOs arrange for example parts availability contracts. On the other hand, in some cases these technological developments also offer opportunities for MLOs to expand their services and move upwards in the service supply chain. For example, KLM E&M provides component (part maintenance) services for several other airline companies.

The control concepts we have designed in Chapter 6 are specific for repairable parts inventory systems with an *internal* repair shop. The concepts and models that are developed will probably also be applicable for repairable parts inventory systems with an *external* repair shop, though we note that we should drop our assumption of a single stock location for ready-for-use LRUs when the repair shop serves a global market. The MLOs in the case study companies often have agreed on a 'lead time interface' (with only a regular lead time) with their external repair shops. From the case studies, we observed that the fraction of repair jobs that is executed within the agreed repair lead times is around 60%-70% and hence there seems to be room for improvement. The designs in Chapter 6 may be a useful starting point for designing alternative control concepts. We however note that the control concepts discussed in Chapter 6 may require frequent coordination (on e.g. shifting repair priorities) between the MLO and the external repair shops, and external repair shops also conduct repair work for other companies, limiting their ability to reschedule repair work. Recent developments like *internet of things* and *control tower technology* may speed up the implementation of control concepts for repairable parts inventory systems in a setting with external repair shops.

Last but not least, we reflect on how we have organized our research. In this research we have closely collaborated with industry. This is not only fun (not 'another day at the office') but it has also provided an inspiring environment to generate new research ideas. Our idea to decouple the inspection and repair phase, for example has come up during discussions on the work floor. This research is partly based on an *empirical* approach, and the empirical information gathering has been useful to substantiate claims and choices in designing control concepts for repairable parts inventory systems. As such, new research directions are defined based on realistic assumptions and validated concepts. We hope that this dissertation is a starting point for more research on repairable parts inventory systems.

Bibliography

Adan, I.J.B.F., A. Sleptchenko, and G.J. Van Houtum. Reducing costs of spare parts supply systems via static priorities. *Asia-Pacific Journal of Operational Research*, 26(4):559–585, 2009.

Akçay, Y. and S.H. Xu. Joint inventory replenishment and component allocation optimization in an assemble-to-order system. *Management Science*, 50(1):99–116, 2004.

Albright, S.C. and A. Gupta. Steady-state approximation of a multi-echelon, multi-indentured repairable-item inventory system with a single repair facility. *Naval Research Logistics*, 40:479–493, 1993.

Alfredsson, P. Optimization of multi-echelon repairable item inventory systems with simulataneous location of repair facilities. *European Journal of Operational Research*, 99:584–595, 1997.

Alfredsson, P. and J. Verrijdt. Modeling emergency supply flexibility in a two-echelon inventory system. *Management Science*, 45(10):1416–1431, 1999.

Altay, N. and L.A. Litteral, editors. *Service Parts Management: Demand Forecasting and Inventory Control*. Springer, 2011.

Arts, J.J. *Spare Parts Planning and Control for Maintenance Operations*. PhD thesis, Eindhoven University of Technology, School of Industrial Engineering, 2013.

Arts, J.J. A multi-item approach to repairable stocking and expediting in a fluctuating demand environment. *European Journal of Operational Research*, 256:102–115, 2017.

Arts, J.J., R.J.I. Basten, and G.J. van Houtum. Repairable stocking and expediting in a fluctuating demand environment: Optimal policy and heuristics. *Operations*

Research, 64(6):1285–1301, 2016.

Arts, J.J., R.J.I. Basten, and G.J. Van Houtum. Maintenance service logistics. In Zijm, W.H.M., M. Klumpp, S.S. Heragu, and A. Regattieri, editors, *Operations, Logistics and Supply Chain Management*. Springer Publishing Company, New York - Berlin, to appear in 2018.

Bailey, G.J. and M.M. Helms. Mro inventory reduction - challenges and management: a case study of the tennessee valley authority. *Production Planning & Control*, 18 (3):261–270, 2007.

Barros, L.L. and M. Riley. A combinatorial approach to level of repair analysis. *European Journal of Operational Research*, 129(2):242–251, 2001.

Basten, R.J.I. and G.J. Van Houtum. System-oriented inventory models for spare parts. *Surveys in Operations Research and Management Science*, 19(19):34–55, 2014.

Basten, R.J.I., J.M.J. Schutten, and M.C. van der Heijden. An efficient model formulation for level of repair analysis. *Annals of Operations Research*, 172(1): 119–142, 2009.

Basten, R.J.I., M.C. van der Heijden, and J.M.J. Schutten. A minimum cost flow model for level of repair analysis. *International Journal of Production Economics*, 133(1):233–242, 2011.

Benjaafar, S., W. L. Cooper, and S. Mardan. Production-inventory systems with imperfect advance demand information and updating. *Naval Research Logistics*, 58 (2):88–106, 2011.

Bertrand, J.W.M. and J.C. Wortmann. Information systems for production planning and control: developments in perspective. *Production Planning & Control*, 3(3): 280–289, 1992.

Bertrand, J.W.M., J.C. Wortman, and J. Wijngaard. *Production control : a structural and design oriented approach*. Elsevier, 1990.

Bitran, G.R., Haas. E.A., and Hax A.C. Hierarchical production planning: a two stage system. *Operations Research*, 30(2):232–251, 1982.

Borst, S., A. Mandelbaum, and M.I. Reiman. Dimensioning large call centers. *Operations Research*, 52(1):17–34, 2004.

Box, G.E.P. and G.M. Jenkins. *Time series analysis, forecasting and control*. San

Francisc0: Holden-Day, 1970.

Bradley, J.R. and H.H. Guerrero. Lifetime buy decisions with multiple obsolete parts. *Production and Operations Management*, 18(1):114–126, 2009.

Buitenhek, R., G.J. Van Houtum, and W.H.M. Zijm. AMVA-based solution procedures for open queueing networks with population constraints. *Annals of Operations Research*, 93:15–40, 2000.

Buzacott, J.A. Queueing models of Kanban and MRP controlled production systems. *Engineering Costs and Production Economics*, 17:3–20, 1989.

Buzacott, J.A. and J.G. Shanthikumar. *Stochastic models of manufacturing systems*, volume 4. Englewood Cliffs, NJ: Prentice Hall, 1993.

Caggiano, K.E., J.A. Muckstadt, and J.A. Rappold. Integrated real-time capacity and inventory allocation for repairable service parts in a two-echelon supply system. *Manufacturing & Service Operations Management*, 8(3):292–319, 2006.

Cavalieri, S., M. Garetti, M. Macchi, and R. Pinto. A decision-making framework for managing maintenance spare parts. *Production Planning & Control*, 19(4):379–396, 2008.

Chakravarthy, S.R. and A. Agarwal. Analysis of a machine repair problem with an unreliable server and phase type repairs and services. *Naval Research Logistics*, 50 (5):462–480, 2003.

Chatfield, C. *The analysis of time series: an introduction.* Chapman & Hall/CRC, 2004.

Chua, R.C.H., G.D. Scudder, and A.V. Hill. Batching policies for a repair shop with limited spares and finite capacity. *European Journal of Operational Research*, 66: 135–147, 1993.

Cohen, M.A., Y. Zheng, and V. Agrawal. Service parts logistics: a benchmark analysis. *IIE Transactions*, 29:627–639, 1997.

Croston, J.D. Forecasting and stock control for intermittent demands. *Operational Research Quarterly*, 23(3):289–303, 1972.

Dallery, Y. and G. Liberopoulos. Extended Kanban control system: Combining Kanban and base stock. *IIE Transactions*, 32(4):369–386, 2000.

de Haas, H.F.M. *The coordination of initial stock and flexible manpower in repairable* *item systems*. PhD thesis, Eindhoven University of Technology, School of Industrial Engineering, 1995.

Díaz, A. and M.C. Fu. Models for multi-echelon repairable item inventory systems with limited repair capacity. *European Journal of Operational Research*, 97:480–492, 1997.

Driessen, M.A. Logistiek Centrum Woensdrecht - Case studie, available upon request from the author, 2011.

Driessen, M.A. KLM Engineering & Maintenance - Case study, available upon request from the author, 2012a.

Driessen, M.A. Defensie Bedrijf Grondgebonden Systemen - Case studie, available upon request from the author, 2012b.

Driessen, M.A. Marinebedrijf - Case studie, available upon request from the author, 2012c.

Driessen, M.A. Europe Container Terminals - Case studie, available upon request from the author, 2013.

Driessen, M.A. and J.J. Arts. Nedtrain - Case study, available upon request from the authors, 2011.

Driessen, M.A., J.J. Arts, G.J. van Houtum, W.D. Rustenburg, and B. Huisman. Maintenance spare parts planning and control: a framework for control and agenda for future research. *Production Planning & Control: The Management of Operations*, 26(5):407–426, 2015a.

Driessen, M.A., J.W. Rustenburg, G.J. Van Houtum, and V.C.S. Wiers. Connecting inventory and repair shop control for repairable items. In Zijm, H., M. Klumpp, U. Clausen, and M. Hompel, editors, *Logistics and Supply Chain Innovation. Lecture Notes in Logistics*, pages 199–221. Springer, Cham, 2015b.

Ebeling, C.E. *Introduction to reliability and maintainablity engineering*. McGraw-Hill, 1997.

Feeney, G.J. and C.C. Sherbrooke. The (s 1, s) inventory policy under compound poisson demand. *Management Science*, 12:391–411, 1966.

Fox, B. Discrete optimization via marginal analysis. *Management Science*, 13(3): 210–216, 1966.

Fransoo, J.C. and V.C.S. Wiers. Action variety of planners: Cognitive load and requisite variety. *Journal of Operations Management*, 24:813–821, 2006.

Fransoo, J.C. and V.C.S. Wiers. An empirical investigation of the neglect of mrp information by production planners. *Production Planning & Control*, 19(8):781–787, 2008.

Frein, Y., M. Di Mascolo, and Y. Dallery. On the design of Generalized Kanban control systems. *International Journal of Operations & Production Management*, 15(9):158–184, 1995.

Graves, S.C. A multi-echelon inventory model for a repairables item with one-for-one replenishment. *Management Science*, 31(10):1247–1256, 1985.

Guide Jr., V.D.R. and R. Srivastava. Repairable inventory theory: Models and applications. *European Journal of Operational Research*, 102:1–20, 1997.

Guide Jr, V.D.R., R. Srivastava, and M.E. Kraus. Priority scheduling policies for repair shops. *International Journal of Production Research*, 38(4):929–950, 2000.

Gupta, A. and S.C. Albright. Steady-state approximations for a multi-echelon, multi-indentured repairable item-inventory system. *European Journal of Operational Research*, 62:340–353, 1992.

Hariharan, R. and P. Zipkin. Customer-order information, leadtimes, and inventories. *Management Science*, 41(10):1599–1607, 1995.

Hausman, W.H. and G.D. Scudder. Priority scheduling rules for repairable inventory systems. *Management Science*, 28(11):1215–1232, 1982.

Heng, A., S. Zhang, A.C.C. Tan, and J. Mathew. Rotating machinery prognostics: State of the art, challenges and opportunities. *Mechanical Systems and Signal Processing*, 23:724739, 2009.

Hopp, W.J. and M.L. Spearman. *Factory physics*. McGraw-Hill, 2001.

Hua, Z.S., B. Zhang, J. Yang, and D.S. Tan. A new approach of forecasting intermittent demand for spare parts inventories in the process industries. *Journal of the Operational Research Society*, 58:52–61, 2007.

Huiskonen, J. Maintenance spare parts logistics: Special characteristics and strategic choices. *International Journal of Production Economics*, 71:125–133, 2001.

Iglehart, D.L. Limiting diffusion approximations for the many server queue and the repairman problem. *Journal of Applied Probability*, 2:429–441, 1965.

Keizers, J., I. Adan, and J. van der Wal. A queueing model for due date control in a multiserver repair shop. *Naval Research Logisitics*, 48:281–292, 2001.

Kennedy, W.J., J.D. Patterson, and L.D. Fredendall. An overview of recent literature on spare parts inventories. *International Journal of Production Economics*, 76: 201–215, 2002.

Knotts, M.H. Civil aircraft maintenance and support fault diagnosis from a business perspective. *Journal of Quality in Maintenance Engineering*, 5(4):335–348, 1999.

Kohler-Gudum, C. and A.G. de Kok. A safety stock adjustment procedure to enable target service levels in simulation of generic inventory systems. *Beta working paper 71*, 2002.

Kranenburg, A.A. and G.J. Van Houtum. Effect of commonality on spare part provisioning costs for capital goods. *International Journal of Production Economics*, 108:221–227, 2007.

Kranenburg, A.A. and G.J. Van Houtum. A new partial pooling structure for spare parts networks. *European Journal of Operational research*, 199(3):908–921, 2009.

Lee, H.L. A multi-echelon inventory model for repairable items woth emergency lateral transshipments. *Management Science*, 33(10):1302–1316, 1987.

Lu, Y., J.S. Song, and Y. Zhao. No-holdback allocation rules for continuous-time assemble-to-order systems. *Operations Research*, 58(3):691–705, 2010.

MIL, STD-1390D. Level Of Repair Analysis (LORA). Technical report, Military Standard, United States Department of Defense: MIL-STD-1390D, 1993.

Muckstadt, J.A. A model for a multi-item, multi-echelon, multi-indenture inventory system. *Management Science*, 20(4):472–481, 1973.

Muckstadt, J.A. *Analysis and algorithms for service part supply chains*. Springer: Berlin, 2005.

Nelson, W. *Applied life data analysis*. John Wiley & Sons, 1982.

Öner, K.B., G.P. Kiesmüller, and G.J. Van Houtum. Optimization of component reliability in the design phase of capital goods. *European Journal of Operational Research*, 205:615–624, 2010.

Paterson, P., G.P. Kiesmüller, R. Teunter, and K. Glazebrook. Inventory models with lateral transshipments: A review. *European Journal of Operational Research*, 210 (2):125 – 136, 2011.

Pinedo, M.L. *Planning and scheduling in manufacturing and services.* Springer, 2009.

Porteus, E.L. *Foundations of Stochastic Inventory Theory.* Stanford University Press, Stanford, 2002.

Rustenburg, W.D. *A system approach to budget-constrained spare parts management.* PhD thesis, Eindhoven University of Technology, 2000.

Rustenburg, W.D., G.J. van Houtum, and W.H.M. Zijm. Exact and approximate analysis of multi-echelon, multi-indenture spare parts systems with commonality. In Shanthikumar, J.G., D.D. Yao, and W.H.M. Zijm, editors, *Stochastic Modeling and Optimization of Manufacturing Systems and Supply Chains*, pages 143–176. Kluwer Academic, Boston/Dordrecht/London, 2003.

Schmeiser, B.W. Batch size effects in the analysis of simulation output. *Operations Research*, 30(3):556–568, 1982.

Scudder, G.D. An evaluation of overtime policies for a repair shop. *Journal of Operations Management*, 6(1):87–98, 1985.

Scudder, G.D. Scheduling and labour assignment policies for a dual-constrained repair shop. *Intenational Journal of Production Research*, 24(3):623–634, 1986.

Scudder, G.D. and R.C.H. Chua. Determining overtime policies for a repair shop. *Omega*, 15(3):197–206, 1987.

Sherbrooke, C.C. METRIC: A multi-echelon technique for recoverable item control. *Operations Research*, 16(1):122–141, 1968.

Sherbrooke, C.C. VARI-METRIC: Improved approximations for multi-indenture, multi-echelon availability models. *Operations Research*, 34(2):311–319, 1986.

Sherbrooke, C.C. *Optimal inventory modeling of systems: Multi-echelon techniques.* Wiley, 2 edition, 2004.

Silver, E.A., D.F. Pyke, and R. Peterson. *Inventory management and production planning and scheduling.* John Wiley & Sons, 1998.

Sleptchenko, A., M.C. van der Heijden, and A. van Harten. Effects of finite repair capacity in multi-echelon, multi-indenture service part supply systems. *International Journal of Production Economics*, 79:209–230, 2002.

Sleptchenko, A., M.C. van der Heijden, and A. van Harten. Trade-off between inventory and repair capacity in spare part networks. *Journal of the Operational Research Society*, 54(3):263–272, 2003.

Sleptchenko, A., M.C. van der Heijden, and A. van Harten. Using repair priorities to reduce stock investment in spare part networks. *European Journal of Operational Research*, 163:733–750, 2005.

Sleptchenko, A., H.H. Turan, S. Pokharel, and T.Y. ElMekkawy. Cross-training policies for repair shops with spare part inventories. *International Journal of Production Economics*, in press, 2017.

Stamatis, D.H. *Failure mode and effect analysis.* ASQC Quality press, 1995.

Steiger, N.M. and J.R. Wilson. Convergence properties of the batch means method for simulation output analysis. *INFORMS Journal on Computing*, 13(4):277–293, 2001.

Suh, N.P. *The principles of design.* Oxford University Press, 1990.

Teunter, R. and L. Duncan. Forecasting intermittent demand: a comparative study. *Journal of the Operational Research Society*, 60:321–329, 2009.

Teunter, R.H. and L. Fortuin. End-of-life service. *International Journal of Production Economics*, 59:487–497, 1999.

Teunter, R.H. and W.K. Klein Haneveld. The 'final order' problem. *European Journal of Operational Research*, 107:35–44, 1998.

Tiemessen, H.G.H. *Dynamic allocation in multi-dimensional inventory models.* PhD thesis, Eindhoven University of Technology, School of Industrial Engineering, 2014.

Tiemessen, H.G.H. and G.J. van Houtum. Reducing costs of repairable inventory supply systems via dynamic scheduling. *International Journal of Production Economics*, 143(2):478–488, 2013.

Tiemessen, H.G.H., M. Fleischmann, and G.J. van Houtum. Dynamic control in multi-item production/inventory systems. *OR spectrum*, 39(1):165–191, 2017.

Van Houtum, G.J. and A.A. Kranenburg. *Spare Parts Inventory Control under System Availability Constraints.* New York: Springer, 1st edition, 2015.

Van Jaarsveld, W. and A. Scheller-Wolf. Optimization of industrial-scale assemble-to-order systems. *INFORMS Journal on Computing*, 27(3):544–560, 2015.

Van Jaarsveld, W., T. Dollevoet, and R. Dekker. Improving spare parts inventory control at a repair shop. *Omega*, 57:217–229, 2015.

Van Kooten, J.P.J. and T. Tan. The final order problem for repairable spare parts under condemnation. *Journal of the Operational Research Society*, 60:1449–1461, 2009.

Van Weele, A.J. *Purchasing and supply chain management.* London: Cengage, 5th edition, 2010.

Verrijdt, J., I. Adan, and T. de Kok. A trade off between emergency repair and inventory investment. *IIE Transactions*, 30:119–132, 1998.

Wagner, S.M. and E. Lindemann. A case study-based analysis of spare parts management in the engineering industry. *Production Planning & Control*, 19(4): 397–407, 2008.

Wang, W. and A.A. Syntetos. Spare parts demand: Linking forecasting to equipment maintenance. *Transportation Research Part E: Logistics and Transportation Review*, 47(6):1194–1209, 2011.

Wiers, V.C.S. The relationship between shop floor autonomy and aps implementation success: evidence from two cases. *Production Planning & Control*, 20(7):576–585, 2009.

Willemain, T.R., C.N. Smart, and H.F. Schwarz. A new approach to forecasting intermittent demand for service part inventories. *International Journal of Forecasting*, 20:375–387, 2004.

Yin, R.K. *Applications of Case Study Research.* Sage Publications, Inc., 2003.

Zijm, W.H.M. Manufacturing and logistic system analysis, planning and control. Unpublished manuscript, 2002.

Zijm, W.H.M. and Z.M. Avşar. Capacitated two-indenture models for repairable item systems. *International Journal of Production Economics*, 81-82:573–588, 2003.

Zijm, W.H.M. and J.M.J. Schutten. Production planning and shopfloor scheduling. 2017. Internal report of the University of Twente, to appear in Zijm, W.H.M. and Klump, M. and Heragu, S. and Regatierri, A., editors, *Operations, Logistics and Supply Chain Management.* Springer Publishing Company, New York - Berlin, 2018.

Summary

"Ik ben er nog steeds van overtuigd dat zoals ik het doe je het moet doen want anders zou ik het niet doen."

Johan Cruijff

In many industries, companies depend on the availability of high-value capital assets to provide their services or to manufacture their products. Examples are manifold, e.g. airlines, public and private transport organizations, hospitals, defense organizations, high-tech manufacturing companies and process industries. Periods in which the asset is not available for use may among others result in lost revenues, customer dissatisfaction and possible claims or a public safety risk. For optimal functioning and availability of the assets, timely maintenance of the assets is essential.

When maintenance of an asset is executed, parts may be taken out and replaced by ready-for-use parts based on their condition (the "repair-by-replacement" principle). These parts are called Line Replaceable Units (LRUs), and sometimes they are entire modules consisting of multiple parts. Modules that are taken out and appear to be irrepairable (or are not worth the effort) are discarded. In the other case, they are returned to a repair shop and further disassembled to find the malfunctioning part, which then is replaced by a well-functioning one after which the module is assembled and added to a ready-for-use inventory. The replaced part, called a Shop Replaceable Unit (SRU), may be processed similarly. In general, we say that an asset may consist of several indenture levels (module, part, part-component, etc.).

To maintain the high-value capital assets, a service network is set up that includes amongst others maintenance and repair facilities, resources and service tools required

to conduct the maintenance, and stock locations to store spare parts. Such a network may be installed by either the original equipment manufacturer or the asset owner, or may be outsourced to specialized maintenance service providers. In this dissertation, we look at companies that both use and maintain high-value capital assets. In such companies a Maintenance Logistics Organization (MLO) is responsible for planning and control of spare parts.

This dissertation presents control concepts for MLOs that aim at a high asset availability at affordable costs of both spare part inventories and repair shops. These concepts relate to multiple decisions that have to be made by the MLOs. To ensure the practical relevance of these concepts, case studies have been used to extensively describe real-life practices and, through a confrontation with the scientific literature, to identify open research topics. This dissertation consists of three parts, which are summarized hereafter.

Part I: Spare parts planning & control in MLOs

The planning and control of spare parts has been extensively researched, resulting in a large number of scientific publications. Nevertheless, in practice we observe solutions/approaches to decision making in MLOs that we do not understand. Therefore, we take a broad view on spare parts planning and control in Chapter 2 by presenting a normative decision framework for MLOs that is based on an extensive literature review. The relevant decision functions in MLOs are related to (generic) supply chain management processes (e.g. forecasting and inventory control) and to repair shop control. The framework is used to cluster the available literature, and to identify open research topics by comparing the framework to practice in six case study companies.

Chapter 3 presents the findings from these case studies. First, the case studies revealed that the framework is complete and that each decision function in the framework has been found in at least two out of the six case study companies. Second, decisions related to repair shop control (capacity dimensioning and scheduling of repair work) and inventory control (setting replenishment policy parameters for LRUs and SRUs) have been found to be implemented differently in the various companies, and are most relevant for MLOs in terms of asset availability and costs of spare parts. These decisions relate to decisions in *repairable parts inventory systems*: systems in which a repair shop returns multiple (failed) repairable parts to the stock of ready-for-use parts.

Decisions in such systems are in practice separated over two decision makers: a repair shop and an inventory control manager. Together, the repair shop and the inventory control manager are responsible for optimally balancing the efficiency and effectiveness of repairable parts inventory systems. The efficiency of a repairable parts inventory system is defined by the utilization of the shop's resources, the work-in-process that resides in the shop and the inventory investments in LRUs and SRUs. The effectiveness of the system is measured by the minimum or average number of backorders for (repairable) LRUs per asset type.

Decisions in repairable parts inventory systems clearly interrelate. A control concept is required to optimally balance the decisions in these systems, such that control and responsibilities of both decision makers are aligned and can be settled in an 'interface agreement'. The case studies have revealed that a "one-size-fits-all" control concept as currently applied in all case study companies, is not optimal and that instead tailor-made solutions are needed that respect the specific characteristics of each repair shop. Part II of this dissertation focuses on the design of control concepts for repairable parts inventory systems.

Part II: Control concepts for repairable parts inventory systems

How to design control concepts for repairable parts inventory systems is the topic of Chapters 4 to 6. The approach to arrive at such concepts is comparable to the approach in Part I of this dissertation. Chapter 4 reviews three streams of literature and serves as a theoretical background for the case studies, presented in Chapter 5. These case studies are used to extensively describe the repair process and repair shop characteristics observed at 22 repairable parts inventory systems in three case study companies. In Chapter 6, findings from both the literature review and the case studies are combined into control concept designs for repairable parts inventory systems.

The literature review of Chapter 4 presents two ways to design the interface between inventory and repair shop control. In the first stream of literature, inventory and repair shop control are decomposed by means of a 'lead time agreement', and decisions related to inventory control and repair shop control are modeled and solved independently. In the second and third stream of literature, these decisions are supposed to be integrated, and hence no interface is needed.

The literature review provides three important principles to account for when designing control concepts for repairable parts inventory systems. First, integrated inventory control of LRUs and SRUs significantly reduces inventory investments, especially when SRUs are relatively expensive and have relatively low demand rates. Second, integrating repair priority setting and LRU stocking decisions substantially reduces the inventory investments. This is in particular important if there exist large differences in the demand rates and prices of the LRUs. Third, workload mechanisms should be applied at the stage level of a repair shop, and hence repair jobs should be released based upon a cap on aggregate flows per resource.

Chapter 5 maps the model assumptions and control principles to the data we gathered from the case studies. It is concluded that there is no repair shop to which the models and control principles can be applied directly, because in each repair shop at least one important model assumption is not satisfied, while a varying set of repair shops complies to a subset of assumptions and control principles. The latter can be explained by the clear differences that have been observed between various repair shops, especially when it comes to the complexity and uncertainty of the materials and capacities required in the repair process.

Chapter 6 presents a typology of repair shops that is based on two dimensions: *capacity complexity* and *material uncertainty*. The capacity complexity concerns the requirement of specialized skills of repair men to execute a repair job. The material uncertainty concerns the extent to which repair jobs for the same item require different sets of materials, i.e. the extent to which it is predictable which materials will be required in the repair process. The value that both dimensions can take on are low and high, and hence we obtain a two-by-two classification and four stereotype repair shops.

Each type of repair shop has been found in the case study companies. For each type of repair shop a control concept design is presented, and design choices are qualitatively motivated. For each control concept design, the interface agreement between the repair shop and the inventory control manager is described, and each decision function is assigned to one of the two decision makers.

Part III: Repair capacity assignment under high material uncertainty

Chapter 7 quantitatively studies the proposed control concept for a repair shop with low capacity complexity and high material uncertainty. For this repair shop type, the design choice is made to decouple the inspection and repair phase of a repair job. Inspection and repair both take a significant amount of time, they are executed by a group of repair men, and repair times depend on the time that elapses between inspection and the actual repair of a part. The repair time gets higher once repair jobs are interrupted, because the inspected part needs to be put aside and the repair man needs to recall the problem essentials when the actual repair starts.

These aspects raise the question whether priority should be given to inspection of failed parts (to reveal which SRUs are required and speed up sourcing of unavailable SRUs) or to completion of the repair of parts for which all required SRUs are available (to increase the availability of LRUs). Therefore, we study a repair shop in which a single repair man faces two queues: one of parts still to be inspected (and to determine what is needed to repair them) and one of parts already inspected for which materials needed are available in stock. Both queues are served on a first-come first-serve basis, the question is when to switch between the two queues.

A simulation model is set up and used to compare several workload-based priority policies. A base case is constructed and (the total costs for) these policies are compared for several scenarios to set the SRU base stock levels. From the evaluated scenarios, the scenario with a high aggregate SRU fill rate combined with a policy that prioritizes repair above inspection leads to the lowest total costs.

About the author

"Ervaring. Dat heb je of dat heb je niet."

Johan Cruijff

Maarten Alexander Driessen was born in Valkenswaard on April 22, 1983. He holds a master degree in Econometrics and Management Science from Tilburg University. In 2007, he started working for Gordian Logistic Experts BV, a company dedicated to designing and implementing solutions on service logistics and spare parts management.

In his work as a service supply chain consultant he gained much experience in the planning and control of spare parts. He did projects for companies that operate in, amongst others, the defense, aerospace, maritime, rail and high-tech industries. He is also involved in the innovation projects Maselma Bridge and Proselo Next.

From 2010 to 2018, he was a part-time PhD student at the Eindhoven University of Technology. He was supervised by prof.dr.ir. Geert-Jan van Houtum, prof.dr. Henk Zijm and dr.ir. Jan Willem Rustenburg. His research is conducted in close cooperation with industry. After defending his PhD thesis, Maarten will continue to work for Gordian Logistic Experts BV.

COMPANIES SUCH AS AIRLINES AND PUBLIC TRANSPORT ORGANIZATIONS DEPEND ON THE AVAILABILITY OF HIGH-VALUE CAPITAL ASSETS TO PROVIDE THEIR SERVICES. TO PREVENT OR SHORTEN THE TIME OF COSTLY DISRUPTIONS WHEN PART FAILURES OCCUR, IT IS CRUCIAL TO HAVE THE RIGHT NUMBER OF SPARE PARTS AVAILABLE FROM STOCK IN ORDER TO QUICKLY REPLACE THE FAILED PARTS. THE REMOVED MALFUNCTIONING PARTS ARE SENT TO A REPAIR SHOP, AND AFTER A SUCCESSFUL REPAIR STOCKED AS SPARE PARTS FOR FUTURE USE.

IN THIS DISSERTATION CONTROL CONCEPTS ARE PRESENTED THAT INTEGRATE DECISIONS RELATED TO CAPACITY PLANNING AT REPAIR SHOPS AND INVENTORY CONTROL FOR REPAIRABLE SPARE PARTS, AND AIM AT A HIGH ASSET AVAILABILITY AT AFFORDABLE COSTS OF BOTH SPARE PART INVENTORIES AND REPAIR SHOPS. CASE STUDIES AT SIX COMPANIES HAVE BEEN USED TO EXTENSIVELY DESCRIBE REAL-LIFE PRACTICES, AND TO IDENTIFY KEY ASPECTS THAT SHOULD BE TAKEN INTO ACCOUNT WHEN DESIGNING SUCH CONTROL CONCEPTS. DIFFERENT CONTROL CONCEPTS ARE PRESENTED FOR THE VARIOUS REPAIR SHOP TYPES ENCOUNTERED IN PRACTICE, AND FOCUS GROUPS AND A MATHEMATICAL MODEL ARE USED TO VALIDATE THESE CONCEPTS IN DETAIL.

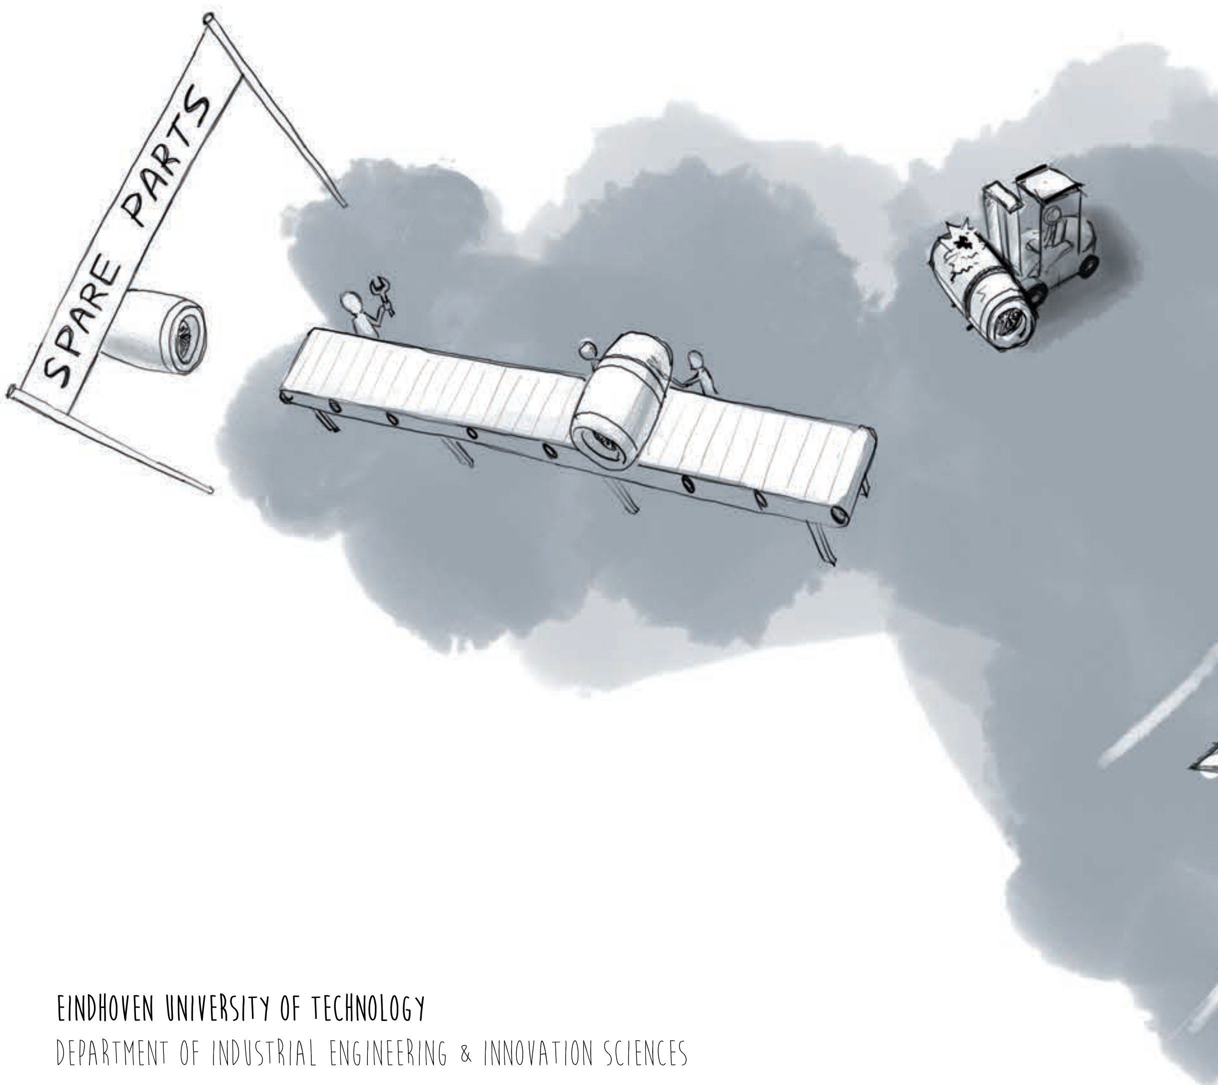

EINDHOVEN UNIVERSITY OF TECHNOLOGY
DEPARTMENT OF INDUSTRIAL ENGINEERING & INNOVATION SCIENCES